

An Infinitely Large Napkin

DRAFT (Last Updated February 15, 2017)

Evan Chen

<http://www.mit.edu/~evanchen/napkin.html>

For Brian and Lisa, who finally got me to write it.

*When introduced to a new idea, always ask why you should care.
Do not expect an answer right away, but demand one eventually.*

— Ravi Vakil [va15]

This version is currently a **first draft**.

There are likely to be numerous errors.

Send corrections, comments, pictures of kittens, etc. to chen.evan6@gmail.com.

© 2016 Evan Chen. All rights reserved. Personal use only.

Last updated February 15, 2017.

Preface: Why this exists

I'll be eating a quick lunch with some friends of mine who are still in high school. They'll ask me what I've been up to the last few weeks, and I'll tell them that I've been learning category theory. They'll ask me what category theory is about. I tell them it's about abstracting things by looking at just the structure-preserving morphisms between them, rather than the objects themselves. I'll try to give them the standard example $\mathbf{Grp}$, but then I'll realize that they don't know what a homomorphism is. So then I'll start trying to explain what a homomorphism is, but then I'll remember that they haven't learned what a group is. So then I'll start trying to explain what a group is, but by the time I finish writing the group axioms on my napkin, they've already forgotten why I was talking about groups in the first place. And then it's 1PM, people need to go places, and I can't help but think:

“Man, if I had forty hours instead of forty minutes, I bet I could actually have explained this all”.

This book is my attempt at those forty hours.

Olympians

What do you do if you're a talented high school student who wants to learn higher math?

To my knowledge, there aren't actually that many possible things to do. One obvious route is to try to negotiate with your school to let you take math classes from a local college. But this option isn't available to many people, and even when it is, more than likely you'll still be the best student in your class at said local college. Plus it's a huge logistical pain in the rare cases where it's at all possible.

Then you have another very popular route — the world of math contests. Suddenly you're connected to this peer group of other kids who are insanely smart. You start training for these contests, you get really good at solving hard problems, and before you know it you're in Lincoln, Nebraska, taking the IMO¹ Team Selection Test. Finally you're feeling challenged, you have clear goals to work towards, and a group of like-minded peers to support and encourage you. You gradually learn all about symmedians, quadratic reciprocity, Muirhead, ...

And yet math olympiads have a weird quirk to them: they restrict themselves to only elementary problems, leaving analysis, group theory, linear algebra, etc. all off limits.

So now suppose you're curious what category theory is all about. Not too many of your friends know, since it's not a contest topic. You could try taking a course from a local college, but that's just going back to square one again. You could try chats with friends or talks or whatever, but that's like the napkin at lunch again: I can tell you category theory is about looking at arrows instead of the objects themselves, but if I don't tell you what a functor is, or about the Yoneda lemma, or what a limit is, I haven't really showed you anything.

So you finally do what any sensible person would do — search “category theory” on Wikipedia. You scroll through the page, realize that you don't know what half the words are, give up, and go back to doing geometry problems from the IMO Shortlist.

¹IMO is short for “International Mathematical Olympiad”, the premier high school mathematical olympiad. See imo-official.org for more details.

Verbum sapienti satis est

Higher math for high school students typically comes in two flavors:

- Someone tells you about the hairy ball theorem in the form “you can’t comb the hair on a spherical cat” then doesn’t tell you anything about why it should be true, what it means to actually “comb the hair”, or any of the underlying theory, leaving you with just some vague notion in your head.
- You take a class and prove every result in full detail, and at some point you stop caring about what the professor is saying.

Presumably you already know how unsatisfying the first approach is. So the second approach seems to be the default, but in general I find it to be terribly inefficient.

I was talking to a friend of mine one day who described briefly what the Israel IMO training looked like. It turns out that rather than actually preparing for the IMO, the students would, say, get taught an entire semester’s worth of undergraduate algebra in the couple weeks. Seeing as a semester is twenty weeks or so, this is an improvement by a factor of ten.

This might sound ludicrous, but I find it totally believable. The glaring issue with classes is that they are not designed for the top students. Olympiad students have a huge advantage here — push them in the right direction, and they can figure out the rest of the details themselves without you having to spell out all the specificities. So it’s easily possible for these top students to learn subjects four or five or six times faster than the average student.

Along these lines, often classes like to prove things for the sake of proving them. I personally find that many proofs don’t really teach you anything, and that it is often better to say “you could work this out if you wanted to, but it’s not worth your time”. Unlike some of the classes you’ll have to take in college, it is not the purpose of this book to train you to solve exercises or write proofs,² but rather to just teach you interesting math. Indeed, most boring non-instructive proofs fall into two categories:

- (i) Short proofs (often “verify all details”) which one could easily work out themselves.
- (ii) Long, less easy proofs that no one remembers two weeks later anyways. (Sometimes so technical as to require lengthy detours.)

The things that are presented should be memorable, or something worth caring about.

In particular, I place a strong emphasis over explaining why a theorem *should* be true rather than writing down its proof. I find the former universally more enlightening. This is a recurrent theme of this book:

Natural explanations supersede proofs.

Presentation of specific topics

See Appendix A for some general comments on why I chose to present certain topics in the way that I did. At the end of said appendix are pointers to several references for further reading.

²Which is not to say problem-solving isn’t valuable; that’s why we do contest math. It’s just not the point of this book.

Obligatory disclaimer

Apparently this wasn't obvious, so let me add:

This book is not primarily intended as a replacement for all other books. The amount of depth is not comparable. (I probably have accidentally implied this in earlier versions of this draft, so let me set that record straight now.)

This is especially true if you plan on doing serious research in any of the subjects covered here. If you're just an curious ex-IMO medalist "what is algebraic NT all about?", then this Napkin may fit you quite well. If you're currently taking an algebraic number theory class and are confused about something, then this Napkin may also be a helpful second reference. But if you're an undergraduate at Harvard hoping that these notes will cover the equivalent of the graduate course Math 223a, you will probably be disappointed.

More succinctly, Napkin is meant to satisfy your curiosity rather than to give you a complete mastery of any topic.

Acknowledgements

I am indebted to countless people for this work. Here is a partial (surely incomplete) list.

Thanks to all my teachers and professors for teaching me much of the material covered in these notes, as well as the authors of all the references I have cited here. A special call-out to [Ga14], [Le14], [Sj05], [Ga03], [Ll15], [Et11], [Ko14], which were especially influential.

Thanks also to dozens of friends and strangers who read through preview copies of my draft, and pointed out errors and gave other suggestions. Special mention to Andrej Vuković and Alexander Chua for together catching over a thousand errors. Thanks also to Brian Gu and Tom Tseng for many corrections. I'd also like to express my gratitude for the many kind words I received during the development of this project; these generous comments led me to keep working on this.

Finally, a huge thanks to the math olympiad community. All the enthusiasm, encouragement, and thank-you notes I have received over the years led me to begin writing this in the first place. I otherwise would never have the arrogance to dream a project like this was at all possible. And of course I would be nowhere near where I am today were it not for the life-changing journey I took in chasing my dreams to the IMO. Forever TWN2!

0 Read this first

§0.1 Prerequisites

The main prerequisite is some amount of mathematical maturity: the ability to read/write proofs, and follow mathematical arguments, etc.


I also assume the reader is familiar with foundational notions such as sets and functions (e.g. “what is a bijection?”). If not, one should consult Appendix E.

§0.2 Graph of chapter dependencies

There is no need to read this book in linear order. Here is a plot of where the various chapters are relative to each other. In my opinion, the really cool stuff is the later chapters, which have been bolded below. Dependencies are indicated by arrows; dotted lines are optional dependencies. I suggest that you simply pick a chapter you find interesting, and then find the shortest path. With that in mind, I hope the length of the entire PDF is not intimidating.

Many of the later parts are subdivided into “Part I” and “Part II”. In general, the second part is substantially harder than the first part. If you intend to read a large portion of the text, one possible thing to do is to read all the Part I’s before reading all the Part II’s.

I have tried to keep the chapters short, and no more than 15 pages each. Several hints and solutions can be found in Appendices B and C.



§0.3 Questions, exercises, and problems

In this book, there are three hierarchies:

- A *question* is intended to be offensively easy. It is usually just a chance to help you internalize definitions. You really ought to answer almost all of these.
- An *exercise* is marginally harder. The difficulty is like something you might see on an AIME. Often I leave proofs of theorems and propositions as exercises if they are instructive and at least somewhat interesting.
- Each chapter features several *problems* at the end. Some of these are easy, while others are legitimately difficult olympiad-style problems. There are three types:
 - **Normal problems**, which are hopefully fun but non-central.
 - **Daggered problems**, which are (usually interesting) results that one should know, but won't be used directly later.
 - **Starred problems**, indicating results which will be used later on.¹

Several hints and solutions can be found in Appendices B and C.



Harder problems are marked with 🌶️'s, like this paragraph. For problems that have three chili's you should probably read the hint.




Image from [Go08]

¹A big nuisance in college for me was that, every so often, the professor would invoke some nontrivial result from the homework: “we are done by PSet 4, Problem 8”, as if I remembered what that was. I wish they could have told me in advance “please take note of this, we’ll use it later on”.

§0.4 Paper

At the risk of being blunt,

Read this book with pencil and paper.

Here's why:




Image from [Or]

You are not God. You cannot keep everything in your head.² If you've printed out a hard copy, then write in the margins. If you're trying to save paper, grab a notebook or something along with the ride. Somehow, some way, make sure you can write. Thanks.

§0.5 Examples

I am pathologically obsessed with examples. In this book, I place all examples in large boxes to draw emphasis to them, which leads to some pages of the book simply consisting of sequences of boxes one after another. I hope the reader doesn't mind.

I also often highlight a "prototypical example" for some sections, and reserve the color red for such a note. The philosophy is that any time the reader sees a definition or a theorem about such an object, they should test it against the prototypical example. If the example is a good prototype, it should be immediately clear why this definition is intuitive, or why the theorem should be true, or why the theorem is interesting, et cetera.

Let me tell you a secret. When I wrote a definition or a theorem in this book, I would have to recall the exact statement from my (quite poor) memory. So instead I have to consider the prototypical example, and then only after that do I remember what the definition or the theorem is. Incidentally, this is also how I learned all the definitions in the first place. I hope you'll find it useful as well.

² See also <https://usamo.wordpress.com/2015/03/14/writing/> and the source above.

§0.6 Topic choices

The appendix contains a list of resources I like, and explanations of pedagogical choices that I made for each chapter. I encourage you to check it out.

In particular, this is where you should go for further reading! There are some topics that should be covered in the Napkin, but are not, due to my own ignorance or laziness. The references provided in this appendix should hopefully help partially atone for my omissions.

§0.7 Conventions and notations

This part describes some of the less familiar notations and definitions and settles for once and for all some annoying issues (“is zero a natural number?”). Most of these are “remarks for experts”: if something doesn’t make sense, you probably don’t have to worry about it for now.

A full glossary of notation used can be found in the appendix.

Sets and equivalence relations

This is brief, intended as a reminder for experts. Consult Appendix E for full details.

The set $\mathbb{N}$ is the set of *positive* integers, not including 0.

An **equivalence relation** on a set X is a relation $\sim$ which is symmetric, reflexive, and transitive. An equivalence relation partitions X into several **equivalence classes**. We will denote this by $X/\sim$. An element of such an equivalence class is a **representative** of that equivalence class.

I always use $\cong$ for an “isomorphism”-style relation (formally: a relation which is an isomorphism in a reasonable category). The only time $\simeq$ is used in the Napkin is for homotopic paths.

I generally use $\subseteq$ and $\subsetneq$ since these are non-ambiguous, unlike $\subset$. I only use $\subset$ on rare occasions in which equality obviously does not hold yet pointing it out would be distracting. For example, I write $\mathbb{Q} \subset \mathbb{R}$ since “ $\mathbb{Q} \subsetneq \mathbb{R}$ ” is distracting.

I prefer $S \setminus T$ to $S - T$.

Functions

Same comments about Appendix E apply as previous section.

Let $X \xrightarrow{f} Y$ be a function.

- By $f^{\text{pre}}(T)$ I mean the **pre-image**

$$f^{\text{pre}}(T) \stackrel{\text{def}}{=} \{x \in X \mid f(x) \in T\}$$

in contrast to the usual $f^{-1}(T)$; I only use f^{-1} for an inverse *function*.

By abuse of notation, we may abbreviate $f^{\text{pre}}(\{y\})$ to $f^{\text{pre}}(y)$. We call $f^{\text{pre}}(y)$ a **fiber**.

- By $f^{\text{“}}(S)$ I mean the **image**

$$f^{\text{“}}(S) \stackrel{\text{def}}{=} \{f(x) \mid x \in S\}.$$

The notation “ is from set theory, and is meant to indicate “point-wise”. Most authors use $f(S)$, but this is abuse of notation, and I prefer $f^{\text{“}}(S)$ for emphasis. *This is probably the least standard notation in the whole Napkin.*

- If $S \subseteq X$, then the **restriction** of f to S is denoted $f|_S$, i.e. it is the function $f|_S : S \rightarrow Y$.
- Sometimes functions $f : X \rightarrow Y$ are *injective* or *surjective*; I may emphasize this sometimes by writing $f : X \hookrightarrow Y$ or $f : X \twoheadrightarrow Y$, respectively.

Rings

All rings have a multiplicative identity 1 unless otherwise specified. We allow $0 = 1$ in general rings but not in integral domains.

All rings are commutative unless otherwise specified. There is an elaborate scheme for naming rings which are not commutative, used only in the chapter on cohomology rings:

	Graded	Not Graded
1 not required	graded psuedo-ring	pseudo-ring
Anticommutative, 1 not required	anticommutative pseudo-ring	N/A
Has 1	graded ring	N/A
Anticommutative with 1	anticommutative ring	N/A
Commutative with 1	commutative graded ring	ring

On the other hand, an *algebra* always has 1, but it need not be commutative.

Choice

We accept the Axiom of Choice.

§0.8 Further reading

Recommendations for other reading can be found in Appendix A.

Contents

Preface	4
0 Read this first	7
0.1 Prerequisites	7
0.2 Graph of chapter dependencies	7
0.3 Questions, exercises, and problems	8
0.4 Paper	9
0.5 Examples	9
0.6 Topic choices	10
0.7 Conventions and notations	10
0.8 Further reading	11
I Basic Algebra and Topology	25
1 What is a group?	26
1.1 Definition and examples of groups	26
1.2 Properties of groups	30
1.3 Isomorphisms	31
1.4 Orders of groups, and Lagrange's theorem	33
1.5 Subgroups	34
1.6 Groups of small orders	36
1.7 Problems to think about	37
2 What is a space?	38
2.1 Definition and examples of metric spaces	38
2.2 Convergence in metric spaces	40
2.3 Continuous maps	41
2.4 Homeomorphisms	42
2.5 Open sets	43
2.6 Forgetting the metric	46
2.7 Closed sets	47
2.8 Common pitfalls	49
2.9 Problems to think about	49
3 Homomorphisms and quotient groups	51
3.1 Generators and group presentations	51
3.2 Homomorphisms	52
3.3 Cosets and modding out	54
3.4 (Optional) Proof of Lagrange's theorem	57
3.5 Eliminating the homomorphism	57
3.6 (Digression) The first isomorphism theorem	59
3.7 Problems to think about	60
4 Topological notions	61
4.1 Connected spaces	61
4.2 Path-connected spaces	62

4.3	Homotopy and simply connected spaces	62
4.4	Bases of spaces	64
4.5	Completeness	65
4.6	Subspacing	66
4.7	Hausdorff spaces	67
4.8	Problems to think about	68
5	Compactness	69
5.1	Definition of sequential compactness	69
5.2	Criteria for compactness	70
5.3	Compactness using open covers	71
5.4	Applications of compactness	73
5.5	(Optional) Equivalence of formulations of compactness	75
5.6	Problems to think about	76
II	Linear Algebra	77
6	What is a vector space?	78
6.1	The definitions of a ring and field	78
6.2	Modules and vector spaces	78
6.3	Direct sums	80
6.4	Linear independence, spans, and basis	82
6.5	Linear maps	84
6.6	What is a matrix?	85
6.7	Subspaces and picking convenient bases	86
6.8	A cute application: Lagrange interpolation	88
6.9	(Digression) Arrays of numbers are evil	88
6.10	A word on general modules	90
6.11	Problems to think about	90
7	Trace and determinant	92
7.1	Tensor product	92
7.2	Dual module	94
7.3	The trace	95
7.4	(Optional) Two more digressions on dual spaces	97
7.5	Wedge product	98
7.6	The determinant	100
7.7	Problems to think about	101
8	Spectral theory	103
8.1	Why you should care	103
8.2	Eigenvectors and eigenvalues	103
8.3	The Jordan form	104
8.4	Nilpotent maps	106
8.5	Reducing to the nilpotent case	107
8.6	(Optional) Proof of nilpotent Jordan	108
8.7	Characteristic polynomials, and Cayley-Hamilton	109
8.8	(Digression) Tensoring up	110
8.9	Problems to think about	111

9	Inner product spaces	112
9.1	Defining the norm	112
9.2	Norms	114
9.3	Orthogonality	115
9.4	Identifying with the dual space	117
9.5	The transpose of a matrix	117
9.6	Spectral theory of normal maps	118
9.7	Problems to think about	119
III	Groups, Rings, and More	121
10	Group actions overkill AIME problems	122
10.1	Definition of a group action	122
10.2	Stabilizers and orbits	123
10.3	Burnside's lemma	124
10.4	Conjugation of elements	125
10.5	Problems to think about	127
11	Find all groups	128
11.1	Sylow theorems	128
11.2	(Optional) Proving Sylow's theorem	129
11.3	(Optional) Simple groups and Jordan-Hölder	131
11.4	Problems to think about	132
12	Rings and ideals	134
12.1	Number theory motivation	134
12.2	Definition and examples of rings	134
12.3	Integral domains and fields	136
12.4	Ideals	137
12.5	Generating ideals	139
12.6	Principal ideal domains and Noetherian rings	140
12.7	Prime ideals	142
12.8	Maximal ideals	143
12.9	Problems	144
13	The PID structure theorem	145
13.1	Finitely generated abelian groups	145
13.2	Some ring theory prerequisites	146
13.3	The structure theorem	147
13.4	Reduction to maps of free R -modules	148
13.5	Smith normal form	149
13.6	Problems to think about	151
IV	Complex Analysis	152
14	Holomorphic functions	153
14.1	The nicest functions on earth	153
14.2	Complex differentiation	155
14.3	Contour integrals	156

14.4	Cauchy-Goursat theorem	158
14.5	Cauchy's integral theorem	159
14.6	Holomorphic functions are analytic	161
14.7	Problems to think about	163
15	Meromorphic functions	164
15.1	The second nicest functions on earth	164
15.2	Meromorphic functions	164
15.3	Winding numbers and the residue theorem	167
15.4	Argument principle	169
15.5	Philosophy: why are holomorphic functions so nice?	170
15.6	Problems to think about	170
16	Holomorphic square roots and logarithms	171
16.1	Motivation: square root of a complex number	171
16.2	Square roots of holomorphic functions	173
16.3	Covering projections	174
16.4	Complex logarithms	174
16.5	Some special cases	175
16.6	Problems to think about	176
V	Quantum Algorithms	177
17	Quantum states and measurements	178
17.1	Bra-ket notation	178
17.2	The state space	179
17.3	Observations	179
17.4	Entanglement	182
17.5	Problems to think about	185
18	Quantum circuits	186
18.1	Classical logic gates	186
18.2	Reversible classical logic	187
18.3	Quantum logic gates	189
18.4	Deutsch-Jozsa algorithm	191
18.5	Problems to think about	192
19	Shor's algorithm	194
19.1	The classical (inverse) Fourier transform	194
19.2	The quantum Fourier transform	195
19.3	Shor's algorithm	197
VI	Algebraic Topology I: Homotopy	199
20	Some topological constructions	200
20.1	Spheres	200
20.2	Quotient topology	200
20.3	Product topology	201
20.4	Disjoint union and wedge sum	202
20.5	CW complexes	203

20.6	The torus, Klein bottle, $\mathbb{RP}^n$, $\mathbb{CP}^n$	204
20.7	Problems to think about	210
21	Fundamental groups	211
21.1	Fusing paths together	211
21.2	Fundamental groups	212
21.3	Fundamental groups are functorial	216
21.4	Higher homotopy groups	217
21.5	Homotopy equivalent spaces	218
21.6	The pointed homotopy category	220
21.7	Problems to think about	221
22	Covering projections	222
22.1	Even coverings and covering projections	222
22.2	Lifting theorem	223
22.3	Lifting correspondence	225
22.4	Regular coverings	227
22.5	The algebra of fundamental groups	228
22.6	Problems to think about	230
VII	Category Theory	231
23	Objects and morphisms	232
23.1	Motivation: isomorphisms	232
23.2	Categories, and examples thereof	232
23.3	Special objects in categories	236
23.4	Binary products	237
23.5	Equalizers	241
23.6	Monic and epic maps	242
23.7	Problems to think about	243
24	Functors and natural transformations	244
24.1	Many examples of functors	244
24.2	Covariant functors	245
24.3	Contravariant functors	247
24.4	(Optional) Natural transformations	248
24.5	(Optional) The Yoneda lemma	251
24.6	Problems to think about	253
25	Abelian categories	254
25.1	Zero objects, kernels, cokernels, and images	254
25.2	Additive and abelian categories	255
25.3	Exact sequences	257
25.4	The Freyd-Mitchell embedding theorem	258
25.5	Breaking long exact sequences	259
25.6	Problems to think about	260

VIII Differential Geometry	262
26 Multivariable calculus done correctly	263
26.1 The total derivative	263
26.2 The projection principle	265
26.3 Total and partial derivatives	266
26.4 (Optional) A word on higher derivatives	268
26.5 Towards differential forms	269
26.6 Problems to think about	269
27 Differential forms	270
27.1 Pictures of differential forms	270
27.2 Pictures of exterior derivatives	272
27.3 Differential forms	273
27.4 Exterior derivatives	274
27.5 Closed and exact forms	276
27.6 Problems to think about	277
28 Integrating differential forms	278
28.1 Motivation: line integrals	278
28.2 Pullbacks	279
28.3 Cells	280
28.4 Boundaries	282
28.5 Stokes' theorem	284
28.6 Problems to think about	284
29 A bit of manifolds	285
29.1 Topological manifolds	285
29.2 Smooth manifolds	286
29.3 Differential forms on manifolds	287
29.4 Orientations	288
29.5 Stokes' theorem for manifolds	289
29.6 Problems to think about	289
IX Algebraic Topology II: Homology	290
30 Singular homology	291
30.1 Simplices and boundaries	291
30.2 The singular homology groups	293
30.3 The homology functor and chain complexes	296
30.4 More examples of chain complexes	300
30.5 Problems to think about	301
31 The long exact sequence	302
31.1 Short exact sequences and four examples	302
31.2 The long exact sequence of homology groups	304
31.3 The Mayer-Vietoris sequence	306
31.4 Problems to think about	311
32 Excision and relative homology	312
32.1 The long exact sequences	312

32.2	The category of pairs	313
32.3	Excision	314
32.4	Some applications	315
32.5	Invariance of dimension	316
32.6	Problems to think about	317
33	Bonus: Cellular homology	318
33.1	Degrees	318
33.2	Cellular chain complex	319
33.3	The cellular boundary formula	322
33.4	Problems to think about	324
34	Singular cohomology	325
34.1	Cochain complexes	325
34.2	Cohomology of spaces	326
34.3	Cohomology of spaces is functorial	327
34.4	Universal coefficient theorem	328
34.5	Example computation of cohomology groups	329
34.6	Relative cohomology groups	330
34.7	Problems to think about	331
35	Application of cohomology	332
35.1	Poincaré duality	332
35.2	de Rham cohomology	332
35.3	Graded rings	333
35.4	Cup products	335
35.5	Relative cohomology pseudo-rings	337
35.6	Wedge sums	337
35.7	Künneth formula	338
35.8	Problems to think about	340
X	Algebraic NT I: Rings of Integers	341
36	Algebraic integers	342
36.1	Motivation from high school algebra	342
36.2	Algebraic numbers and algebraic integers	343
36.3	Number fields	344
36.4	Norms and traces	345
36.5	The ring of integers	348
36.6	Primitive element theorem, and monogenic extensions	351
36.7	Problems to think about	352
37	Unique factorization (finally!)	354
37.1	Motivation	354
37.2	Ideal arithmetic	355
37.3	Field of fractions	356
37.4	Dedekind domains	356
37.5	Unique factorization works	357
37.6	The factoring algorithm	359
37.7	Fractional ideals	361

37.8	The ideal norm	363
37.9	Problems to think about	364
38	Minkowski bound and class groups	365
38.1	The class group	365
38.2	The discriminant of a number field	365
38.3	The signature of a number field	368
38.4	Minkowski's theorem	370
38.5	The trap box	371
38.6	The Minkowski bound	372
38.7	The class group is finite	373
38.8	Computation of class numbers	374
38.9	Problems to think about	377
39	More properties of the discriminant	378
39.1	Problems to think about	378
40	Bonus: Let's solve Pell's equation!	379
40.1	Units	379
40.2	Dirichlet's unit theorem	380
40.3	Finding fundamental units	381
40.4	Pell's equation	382
40.5	Problems to think about	383
XI	Algebraic NT II: Galois and Ramification Theory	384
41	Things Galois	385
41.1	Motivation	385
41.2	Field extensions, algebraic closures, and splitting fields	386
41.3	Embeddings into algebraic closures for number fields	387
41.4	Everyone hates characteristic 2: separable vs irreducible	388
41.5	Automorphism groups and Galois extensions	390
41.6	Fundamental theorem of Galois theory	392
41.7	Problems to think about	394
41.8	(Optional) Proof that Galois extensions are splitting	395
42	Finite fields	397
42.1	Example of a finite field	397
42.2	Finite fields have prime power order	398
42.3	All finite fields are isomorphic	399
42.4	The Galois theory of finite fields	400
42.5	Problems to think about	401
43	Ramification theory	402
43.1	Ramified / inert / split primes	402
43.2	Primes ramify if and only if they divide Δ_K	403
43.3	Inertial degrees	403
43.4	The magic of Galois extensions	404
43.5	(Optional) Decomposition and inertia groups	406
43.6	Tangential remark: more general Galois extensions	408

43.7	Problems to think about	408
44	The Frobenius element	409
44.1	Frobenius elements	409
44.2	Conjugacy classes	410
44.3	Chebotarev density theorem	412
44.4	Example: Frobenius elements of cyclotomic fields	412
44.5	Frobenius elements behave well with restriction	413
44.6	Application: Quadratic reciprocity	414
44.7	Frobenius elements control factorization	416
44.8	Example application: IMO 2003 problem 6	419
44.9	Problems to think about	420
45	Bonus: A Bit on Artin Reciprocity	421
45.1	Infinite primes	421
45.2	Modular arithmetic with infinite primes	421
45.3	Infinite primes in extensions	423
45.4	Frobenius element and Artin symbol	424
45.5	Artin reciprocity	426
45.6	Problems to think about	429
XII	Representation Theory	430
46	Representations of algebras	431
46.1	Algebras	431
46.2	Representations	432
46.3	Direct sums	434
46.4	Irreducible and indecomposable representations	435
46.5	Morphisms of representations	436
46.6	The representations of $\text{Mat}_d(k)$	438
46.7	Problems to think about	439
47	Semisimple algebras	441
47.1	Schur's lemma continued	441
47.2	Density theorem	442
47.3	Semisimple algebras	444
47.4	Maschke's theorem	445
47.5	Example: the representations of $\mathbb{C}[S_3]$	445
47.6	Problems to think about	446
48	Characters	447
48.1	Definitions	447
48.2	The dual space modulo the commutator	448
48.3	Orthogonality of characters	449
48.4	Examples of character tables	451
48.5	Problems to think about	452
49	Some applications	454
49.1	Frobenius divisibility	454
49.2	Burnside's theorem	455

49.3	Frobenius determinant	456
XIII	Algebraic Geometry I: Varieties	458
50	Affine varieties	459
50.1	Affine varieties	459
50.2	Naming affine varieties via ideals	460
50.3	Radical ideals and Hilbert's Nullstellensatz	461
50.4	Pictures of varieties in $\mathbb{A}^1$	462
50.5	Prime ideals correspond to irreducible affine varieties	463
50.6	Pictures in $\mathbb{A}^2$ and $\mathbb{A}^3$	464
50.7	Maximal ideals	465
50.8	Why schemes?	466
50.9	Problems to think about	466
51	Affine varieties as ringed spaces	467
51.1	Synopsis	467
51.2	The Zariski topology on $\mathbb{A}^n$	467
51.3	The Zariski topology on affine varieties	469
51.4	Coordinate rings	470
51.5	The sheaf of regular functions	471
51.6	Regular functions on distinguished open sets	472
51.7	Baby ringed spaces	474
51.8	Problems to think about	474
52	Projective varieties	475
52.1	Graded rings	475
52.2	The ambient space	476
52.3	Homogeneous ideals	477
52.4	As ringed spaces	478
52.5	Examples of regular functions	480
52.6	Problems to think about	481
53	Bonus: Bézout's theorem	482
53.1	Non-radical ideals	482
53.2	Hilbert functions of finitely many points	483
53.3	Hilbert polynomials	485
53.4	Bézout's theorem	487
53.5	Applications	487
53.6	Problems to think about	488
XIV	Algebraic Geometry II: Schemes	489
54	Morphisms of varieties	490
54.1	Defining morphisms of baby ringed spaces	490
54.2	Examples	491
54.3	Quasi-projective varieties	493
54.4	Some applications	494
54.5	Problems to think about	495

55 Sheaves and ringed spaces	496
55.1 Pre-sheaves	496
55.2 Sheaves	497
55.3 Stalks	498
55.4 Sections “are” sequences of germs	501
55.5 Sheafification	502
55.6 Morphisms of sheaves	503
55.7 Local rings, and locally ringed spaces	504
55.8 Morphisms of (locally) ringed spaces	505
55.9 Problems to think about	507
56 Schemes	508
56.1 The set of points	508
56.2 The Zariski topology of the spectrum	509
56.3 The structure sheaf	510
56.4 Example: fat points	512
56.5 Properties of affine schemes	513
56.6 Schemes	515
56.7 Projective scheme	515
56.8 Where to go from here	517
56.9 Problems to think about	517
XV Set Theory I: ZFC, Ordinals, and Cardinals	518
57 Bonus: Cauchy’s functional equation and Zorn’s lemma	519
57.1 Let’s construct a monster	519
57.2 Review of finite induction	520
57.3 Transfinite induction	520
57.4 Wrapping up functional equations	523
57.5 Zorn’s lemma	524
58 Zermelo-Fraenkel with choice	526
58.1 The ultimate functional equation	526
58.2 Cantor’s paradox	526
58.3 The language of set theory	527
58.4 The axioms of ZFC	528
58.5 Encoding	530
58.6 Choice and well-ordering	530
58.7 Sets vs classes	531
58.8 Problems to think about	532
59 Ordinals	533
59.1 Counting for preschoolers	533
59.2 Counting for set theorists	534
59.3 Definition of an ordinal	536
59.4 Ordinals are “tall”	537
59.5 Transfinite induction and recursion	538
59.6 Ordinal arithmetic	539
59.7 The hierarchy of sets	540
59.8 Problems to think about	542

60 Cardinals	543
60.1 Equinumerous sets and cardinals	543
60.2 Cardinalities	544
60.3 Aleph numbers	544
60.4 Cardinal arithmetic	545
60.5 Cardinal exponentiation	547
60.6 Cofinality	547
60.7 Inaccessible cardinals	549
60.8 Problems to think about	549
 XVI Set Theory II: Model Theory and Forcing	 550
61 Inner model theory	551
61.1 Models	551
61.2 Sentences and satisfaction	552
61.3 The Levy hierarchy	554
61.4 Substructures, and Tarski-Vaught	555
61.5 Obtaining the axioms of ZFC	556
61.6 Mostowski collapse	557
61.7 Adding an inaccessible	557
61.8 FAQ's on countable models	559
61.9 Picturing inner models	559
61.10 Problems to think about	561
 62 Forcing	 562
62.1 Setting up posets	563
62.2 More properties of posets	564
62.3 Names, and the generic extension	565
62.4 Fundamental theorem of forcing	568
62.5 (Optional) Defining the relation	568
62.6 The remaining axioms	570
62.7 Problems to think about	570
 63 Breaking the continuum hypothesis	 571
63.1 Adding in reals	571
63.2 The countable chain condition	572
63.3 Preserving cardinals	573
63.4 Infinite combinatorics	574
63.5 Problems to think about	575
 XVII Backmatter	 576
A Pedagogical comments and references	577
A.1 Basic algebra and topology	577
A.2 Second-year topics	578
A.3 Advanced topics	579
A.4 Further topics	580
 B Hints to selected problems	 581

C	Sketches of selected solutions	588
D	Glossary of notations	602
D.1	General	602
D.2	Functions and sets	602
D.3	Abstract and linear algebra	603
D.4	Quantum computation	604
D.5	Topology and (complex) analysis	604
D.6	Category theory	606
D.7	Differential geometry	606
D.8	Algebraic number theory	607
D.9	Representation theory	608
D.10	Algebraic geometry	608
D.11	Set theory	609
E	Terminology on sets and functions	611
E.1	Sets	611
E.2	Functions	612
E.3	Equivalence relations	614

DRAFT (Evan Chen)
Updated February 15, 2017

I

Basic Algebra and Topology

1	What is a group?	26
1.1	Definition and examples of groups	26
1.2	Properties of groups	30
1.3	Isomorphisms	31
1.4	Orders of groups, and Lagrange's theorem	33
1.5	Subgroups	34
1.6	Groups of small orders	36
1.7	Problems to think about	37
2	What is a space?	38
2.1	Definition and examples of metric spaces	38
2.2	Convergence in metric spaces	40
2.3	Continuous maps	41
2.4	Homeomorphisms	42
2.5	Open sets	43
2.6	Forgetting the metric	46
2.7	Closed sets	47
2.8	Common pitfalls	49
2.9	Problems to think about	49
3	Homomorphisms and quotient groups	51
3.1	Generators and group presentations	51
3.2	Homomorphisms	52
3.3	Cosets and modding out	54
3.4	(Optional) Proof of Lagrange's theorem	57
3.5	Eliminating the homomorphism	57
3.6	(Digression) The first isomorphism theorem	59
3.7	Problems to think about	60
4	Topological notions	61
4.1	Connected spaces	61
4.2	Path-connected spaces	62
4.3	Homotopy and simply connected spaces	62
4.4	Bases of spaces	64
4.5	Completeness	65
4.6	Subspacing	66
4.7	Hausdorff spaces	67
4.8	Problems to think about	68
5	Compactness	69
5.1	Definition of sequential compactness	69
5.2	Criteria for compactness	70
5.3	Compactness using open covers	71
5.4	Applications of compactness	73
5.5	(Optional) Equivalence of formulations of compactness	75
5.6	Problems to think about	76

1 What is a group?

A group is one of the most basic structures in higher mathematics. In this chapter I will tell you only the bare minimum: what a group is, and when two groups are the same.

§1.1 Definition and examples of groups

Prototypical example for this section: The additive group of integers $(\mathbb{Z}, +)$ and the cyclic group $\mathbb{Z}_m$. Just don't let yourself forget that most groups are non-commutative.

A group consists of two pieces of data: a set G , and an associative binary operation $\star$ with some properties. Before I write down the definition of a group, let me give two examples.

Example 1.1.1 (Additive integers)

The pair $(\mathbb{Z}, +)$ is a group: $\mathbb{Z} = \{\dots, -2, -1, 0, 1, 2, \dots\}$ is the set and the associative operation is *addition*. Note that

- The element $0 \in \mathbb{Z}$ is an *identity*: $a + 0 = 0 + a = a$ for any a .
- Every element $a \in \mathbb{Z}$ has an additive *inverse*: $a + (-a) = (-a) + a = 0$.

We call this group $\mathbb{Z}$.

Example 1.1.2 (Nonzero rationals)

Let $\mathbb{Q}^\times$ be the set of *nonzero rational numbers*. The pair $(\mathbb{Q}^\times, \cdot)$ is a group: the set is $\mathbb{Q}^\times$ and the associative operation is *multiplication*.

Again we see the same two nice properties.

- The element $1 \in \mathbb{Q}^\times$ is an *identity*: for any rational number, $a \cdot 1 = 1 \cdot a = a$.
- For any rational number $x \in \mathbb{Q}^\times$, we have an inverse x^{-1} , such that

$$x \cdot x^{-1} = x^{-1} \cdot x = 1.$$

From this I bet you can already guess what the definition of a group is.

Definition 1.1.3. A **group** is a pair $G = (G, \star)$ consisting of a set of elements G , and a binary operation $\star$ on G , such that:

- G has an **identity element**, usually denoted 1_G or just 1 , with the property that

$$1_G \star g = g \star 1_G = g \text{ for all } g \in G.$$

- The operation is **associative**, meaning $(a \star b) \star c = a \star (b \star c)$ for any $a, b, c \in G$. Consequently we generally don't write the parentheses.
- Each element $g \in G$ has an **inverse**, that is, an element $h \in G$ such that

$$g \star h = h \star g = 1_G.$$

Remark 1.1.4. Some authors like to add a “closure” axiom, i.e. to say explicitly that $g \star h \in G$. This is implied already by the fact that $\star$ is a binary operation on G , but is worth keeping in mind for the examples below.

Remark 1.1.5. It is not required that $\star$ is commutative ($a \star b = b \star a$). So we say that a group is **abelian** if the operation is commutative and **non-abelian** otherwise.

Example 1.1.6 (Non-Examples of groups)

- The pair $(\mathbb{Q}, \cdot)$ is NOT a group. (Here $\mathbb{Q}$ is rational numbers.) While there is an identity element, the element $0 \in \mathbb{Q}$ does not have an inverse.
- The pair $(\mathbb{Z}, \cdot)$ is also NOT a group. (Why?)
- Let $\text{Mat}_{2 \times 2}(\mathbb{R})$ be the set of 2×2 real matrices. Then $(\text{Mat}_{2 \times 2}(\mathbb{R}), \cdot)$ (where $\cdot$ is matrix multiplication) is NOT a group. Indeed, even though we have an identity matrix

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

we still run into the same issue as before: the zero matrix does not have a multiplicative inverse.

(Even if we delete the zero matrix from the set, the resulting structure is still not a group: those of you that know even a little linear algebra might recall that any matrix with determinant zero cannot have an inverse.)

Remark 1.1.7 (Digression). A common question is: why these axioms? For example, why associative but not commutative? This answer will likely not make sense until later, but here are some comments that may help.

One general heuristics is: Whenever you define a new type of general object, there’s always a balancing act going on. On the one hand, you want to include enough constraints that your objects are “nice”. On the other hand, if you include too many constraints, then your definition applies to too few objects. So, for example, we include “associative” because that makes our lives easier and most operations we run into are associative. We don’t include “commutative” though, because examples below show that there are lots of non-abelian groups we care about. (But we introduce the name “abelian” because we still want to keep track of it.)

Another comment: a good motivation for the inverse axioms is that you get a large amount of *symmetry*. The set of positive integers with addition is not a group, for example, because you can’t subtract 6 from 3: some elements are “larger” than others. By requiring an inverse element to exist, you get rid of this issue. (You also need identity for this; it’s hard to define inverses without it.)

Even more abstruse comment: Problem 1E[†] shows that groups are actually shadows of the so-called symmetric groups (defined later, also called permutation groups). This makes rigorous the notion that “groups are very symmetric”.

Let’s resume writing down examples. Here are some more **abelian examples** of groups:

Example 1.1.8 (Complex unit circle)

Let S^1 denote the set of complex numbers z with absolute value one; that is

$$S^1 \stackrel{\text{def}}{=} \{z \in \mathbb{C} \mid |z| = 1\}.$$

Then $(S^1, \times)$ is a group because

- The complex number $1 \in S^1$ serves as the identity, and
- Each complex number $z \in S^1$ has an inverse $\frac{1}{z}$ which is also in S^1 , since $|z^{-1}| = |z|^{-1} = 1$.

There is one thing I ought to also check: that $z_1 \times z_2$ is actually still in S^1 . But this follows from the fact that $|z_1 z_2| = |z_1| |z_2| = 1$.

Example 1.1.9 (Addition mod n)

Here is an example from number theory: Let $n > 1$ be an integer, and consider the residues (remainders) modulo n . These form a group under addition. We call this the **cyclic group of order n** , and abbreviate it as $\mathbb{Z}_n$, with elements $\bar{0}, \bar{1}, \dots$.

Example 1.1.10 (Multiplication mod p)

Let p be a prime. Consider the *nonzero residues modulo p* , which we denote by $\mathbb{Z}_p^\times$. Then $(\mathbb{Z}_p^\times, \times)$ is a group.

Question 1.1.11. Why do we need the fact that p is prime?

Here are some **non-abelian examples**:

Example 1.1.12 (General linear group)

Let n be a positive integer. Then $\text{GL}_n(\mathbb{R})$ is defined as the set of $n \times n$ real matrices which have nonzero determinant. It turns out that with this condition, every matrix does indeed have an inverse, so $(\text{GL}_n(\mathbb{R}), \times)$ is a group, called the **general linear group**.

(The fact that $\text{GL}_n(\mathbb{R})$ is closed under $\times$ follows from the linear algebra fact that $\det(AB) = \det A \det B$, proved in later chapters.)

Example 1.1.13 (Special linear group)

Following the example above, let $\text{SL}_n(\mathbb{R})$ denote the set of $n \times n$ matrices whose determinant is actually 1. Again, for linear algebra reasons it turns out that $(\text{SL}_n(\mathbb{R}), \times)$ is also a group, called the **special linear group**.

Example 1.1.14 (Symmetric groups)

Let S_n be the set of permutations of $\{1, \dots, n\}$. By viewing these permutations as functions from $\{1, \dots, n\}$ to itself, we can consider *compositions* of permutations. Then the pair $(S_n, \circ)$ (here $\circ$ is function composition) is also a group, because

- There is an identity permutation, and
- Each permutation has an inverse.

The group S_n is called the **symmetric group** on n elements.

Example 1.1.15 (Dihedral group)


The **dihedral group of order $2n$** , denoted D_{2n} , is the group of symmetries of a regular n -gon $A_1A_2 \dots A_n$, which includes rotations and reflections. It consists of the $2n$ elements

$$\{1, r, r^2, \dots, r^{n-1}, s, sr, sr^2, \dots, sr^{n-1}\}.$$

The element r corresponds to rotating the n -gon by $\frac{2\pi}{n}$, while s corresponds to reflecting it across the line OA_1 (here O is the center of the polygon). So rs mean “reflect then rotate” (like with function composition, we read from right to left).

In particular, $r^n = s^2 = 1$. You can also see that $r^k s = sr^{-k}$.

Here is a picture of some elements of D_{10} .



Trivia: the dihedral group D_{12} is my favorite example of a non-abelian group, and is the first group I try for any exam question of the form “find an example...”.

More examples:

Example 1.1.16 (Products of groups)

Let $(G, \star)$ and $(H, *)$ be groups. We can define a **product group** $(G \times H, \cdot)$, as follows. The elements of the group will be ordered pairs $(g, h) \in G \times H$. Then

$$(g_1, h_1) \cdot (g_2, h_2) = (g_1 \star g_2, h_1 * h_2) \in G \times H$$

is the group operation.

Question 1.1.17. What are the identity and inverses of the product group?

Example 1.1.18 (Trivial group)

The **trivial group**, often denoted 0 or 1, is the group with only an identity element. I will use the notation $\{1\}$.

Exercise 1.1.19. Which of these are groups?

- (a) Rational numbers with odd denominators (in simplest form), where the operation is addition. (This includes integers, written as $n/1$, and $0 = 0/1$).
- (b) The set of rational numbers with denominator at most 2, where the operation is addition.
- (c) The set of rational numbers with denominator at most 2, where the operation is multiplication.
- (d) The set of nonnegative integers, where the operation is addition.

§1.2 Properties of groups

Prototypical example for this section: $\mathbb{Z}_p^\times$ is possibly best.

Abuse of Notation 1.2.1. From now on, we'll often refer to a group $(G, \star)$ by just G . Moreover, we'll abbreviate $a \star b$ to just ab . Also, because the operation $\star$ is associative, we will omit unnecessary parentheses: $(ab)c = a(bc) = abc$.

Abuse of Notation 1.2.2. From now on, for any $g \in G$ and $n \in \mathbb{N}$ we abbreviate

$$g^n = \underbrace{g \star \cdots \star g}_{n \text{ times}}.$$

Moreover, we let g^{-1} denote the inverse of g , and $g^{-n} = (g^{-1})^n$.

If you did functional equations at the IMO, you might know that you can actually determine a lot about a function just by knowing a few properties of it. For example, if you know that $f : \mathbb{Q} \rightarrow \mathbb{R}$ satisfies $f(x + y) = f(x) + f(y)$, then you actually can show that $f(x) = cx$ for some $c \in \mathbb{R}$. (This is called Cauchy's functional equation.)

In the same vein, we can try to deduce some properties that a group must have just from Definition 1.1.3. (In fact, this really is a functional equation: $\star$ is just a function $G \times G \rightarrow G$.)

It is a law in Guam and 37 other states that I now state the following proposition.

Fact 1.2.3. Let G be a group.

- (a) The identity of a group is unique.
- (b) The inverse of any element is unique.
- (c) For any $g \in G$, $(g^{-1})^{-1} = g$.

Proof. This is mostly just some formal “functional-equation” style manipulations, and you needn't feel bad skipping it on a first read.

- (a) If 1 and $1'$ are identities, then $1 = 1 \star 1' = 1'$.
- (b) If h and h' are inverses to g , then $1_G = g \star h \implies h' = (h' \star g) \star h = 1_G \star h = h$.
- (c) Trivial; omitted. □

Now we state a slightly more useful proposition.

Proposition 1.2.4 (Inverse of products)

Let G be a group, and $a, b \in G$. Then $(ab)^{-1} = b^{-1}a^{-1}$.

Proof. Direct computation. We have

$$(ab)(b^{-1}a^{-1}) = a(bb^{-1})a^{-1} = aa^{-1} = 1_G.$$

Hence $(ab)^{-1} = b^{-1}a^{-1}$. Similarly, $(b^{-1}a^{-1})(ab) = 1_G$ as well. $\square$

Finally, we state a very important lemma about groups, which highlights why having an inverse is so valuable.

Lemma 1.2.5 (Left multiplication is a bijection)

Let G be a group, and pick a $g \in G$. Then the map $G \rightarrow G$ given by $x \mapsto gx$ is a bijection.

Exercise 1.2.6. Check this by showing injectivity and surjectivity directly. (If you don't know what these words mean, consult Appendix E.)

Example 1.2.7

Let $G = \mathbb{Z}_7^\times$ (as in Example 1.1.10) and pick $g = 3$. The above lemma states that the map $x \mapsto 3 \cdot x$ is a bijection, and we can see this explicitly:

$$\begin{aligned} 1 &\xrightarrow{\times 3} 3 \pmod{7} \\ 2 &\xrightarrow{\times 3} 6 \pmod{7} \\ 3 &\xrightarrow{\times 3} 2 \pmod{7} \\ 4 &\xrightarrow{\times 3} 5 \pmod{7} \\ 5 &\xrightarrow{\times 3} 1 \pmod{7} \\ 6 &\xrightarrow{\times 3} 4 \pmod{7}. \end{aligned}$$

The fact that the map is injective is often called the **cancellation law**. (Why do you think so?)

§1.3 Isomorphisms

Prototypical example for this section: $\mathbb{Z} \cong 10\mathbb{Z}$.

First, let me talk about what it means for groups to be isomorphic. Consider the two groups

- $\mathbb{Z} = (\{\dots, -2, -1, 0, 1, 2, \dots\}, +)$.
- $10\mathbb{Z} = (\{\dots, -20, -10, 0, 10, 20, \dots\}, +)$.

These groups are “different”, but only superficially so – you might even say they only differ in the names of the elements. Think about what this might mean formally for a moment.

Specifically the map

$$\phi : \mathbb{Z} \rightarrow 10\mathbb{Z} \text{ by } x \mapsto 10x$$

is a bijection of the underlying sets which respects the group action. In symbols,

$$\phi(x + y) = \phi(x) + \phi(y).$$

In other words, ϕ is a way of re-assigning names of the elements without changing the structure of the group. That’s all just formalism for capturing the obvious fact that $(\mathbb{Z}, +)$ and $(10\mathbb{Z}, +)$ are the same thing.

Now, let’s do the general definition.

Definition 1.3.1. Let $G = (G, \star)$ and $H = (H, *)$ be groups. A bijection $\phi : G \rightarrow H$ is called an **isomorphism** if

$$\phi(g_1 \star g_2) = \phi(g_1) * \phi(g_2) \quad \text{for all } g_1, g_2 \in G.$$

If there exists an isomorphism from G to H , then we say G and H are **isomorphic** and write $G \cong H$.

Note that in this definition, the left-hand side $\phi(g_1 \star g_2)$ uses the operation of G while the right-hand side uses the operation of H .

Example 1.3.2 (Examples of isomorphisms)

Let G and H be groups. We have the following isomorphisms.

(a) $\mathbb{Z} \cong 10\mathbb{Z}$, as above.

(b) There is an isomorphism

$$G \times H \cong H \times G$$

by the map $(g, h) \mapsto (h, g)$.

(c) The identity map $\text{id} : G \rightarrow G$ is an isomorphism, hence $G \cong G$.

(d) There is another isomorphism of $\mathbb{Z}$ to itself: send every x to $-x$.

Example 1.3.3 (Primitive roots modulo 7)

As a nontrivial example, we claim that $\mathbb{Z}_6 \cong \mathbb{Z}_7^\times$. The bijection is

$$\phi(a \bmod 6) = 3^a \bmod 7.$$

To check that this is an isomorphism, we need to verify several things.

- First, we need to check this map actually makes sense: why is it the case that if $a \equiv b \pmod{6}$, then $3^a \equiv 3^b \pmod{7}$? The reason is that Fermat's little theorem guarantees that $3^6 \equiv 1 \pmod{7}$.

- Next, we need to check that this map is a bijection. You can do this explicitly:

$$(3^1, 3^2, 3^3, 3^4, 3^5, 3^6) \equiv (3, 2, 6, 4, 5, 1) \pmod{7}.$$

- Finally, we need to verify that this map respects the group action. In other words, we want to see that $\phi(a+b) = \phi(a)\phi(b)$ since the operation of $\mathbb{Z}_6$ is addition while the operation of $\mathbb{Z}_7^\times$ is multiplication. That's just saying that $3^{a+b} \equiv 3^a 3^b \pmod{7}$, which is true.

Example 1.3.4 (Primitive roots)

More generally, for any prime p , there exists an element $g \in \mathbb{Z}_p^\times$ called a **primitive root** modulo p such that $1, g, g^2, \dots, g^{p-2}$ are all different modulo p . One can show by copying the above proof that

$$\mathbb{Z}_{p-1} \cong \mathbb{Z}_p^\times \text{ for all primes } p.$$

The example above was the special case $p = 7$ and $g = 3$.

Exercise 1.3.5. Assuming the existence of primitive roots, establish the isomorphism $\mathbb{Z}_{p-1} \cong \mathbb{Z}_p^\times$ as above.

It's not hard to see that $\cong$ is an equivalence relation (why?). Moreover, because we really only care about the structure of groups, we'll usually consider two groups to be the same when they are isomorphic. So phrases such as "find all groups" really mean "find all groups up to isomorphism".

§1.4 Orders of groups, and Lagrange's theorem

Prototypical example for this section: $\mathbb{Z}_p^\times$.

As is typical in math, we use the word "order" for way too many things. In groups, there are two notions of order.

Definition 1.4.1. The **order of a group** G is the number of elements of G . We denote this by $|G|$. Note that the order may not be finite, as in $\mathbb{Z}$. We say G is a **finite group** just to mean that $|G|$ is finite.

Example 1.4.2 (Orders of groups)

For a prime p , $|\mathbb{Z}_p^\times| = p - 1$. In other words, the order of $\mathbb{Z}_p^\times$ is $p - 1$. As another example, the order of the symmetric group S_n is $|S_n| = n!$ and the order of the dihedral group D_{2n} is $2n$.

Definition 1.4.3. The **order of an element** $g \in G$ is the smallest positive integer n such that $g^n = 1_G$, or ∞ if no such n exists. We denote this by $\text{ord } g$.

Example 1.4.4 (Examples of orders)

The order of -1 in $\mathbb{Q}^\times$ is 2, while the order of 1 in $\mathbb{Z}$ is infinite.

Question 1.4.5. Find the order of each of the six elements of $\mathbb{Z}_6$, the cyclic group on six elements. (See Example 1.1.9 if you've forgotten what $\mathbb{Z}_6$ means.)

Example 1.4.6 (Primitive roots)

If you know olympiad number theory, this coincides with the definition of an order of a residue mod p . That's why we use the term "order" there as well. In particular, a primitive root is precisely an element $g \in \mathbb{Z}_p^\times$ such that $\text{ord } g = p - 1$.

You might also know that if $x^n \equiv 1 \pmod{p}$, then the order of $x \pmod{p}$ must divide n . The same is true in a general group for exactly the same reason.

Fact 1.4.7. If $g^n = 1_G$ then $\text{ord } g$ divides n .

Also, you can show that any element of a finite group has a finite order. The proof is just an olympiad-style pigeonhole argument. Consider the infinite sequence $1_G, g, g^2, \dots$, and find two elements that are the same.

Fact 1.4.8. Let G be a finite group. For any $g \in G$, $\text{ord } g$ is finite.

What's the last property of $\mathbb{Z}_p^\times$ that you know from olympiad math? We have Fermat's little theorem: for any $a \in \mathbb{Z}_p^\times$, we have $a^{p-1} \equiv 1 \pmod{p}$. This is no coincidence: exactly the same thing is true in a more general setting.

Theorem 1.4.9 (Lagrange's theorem for orders)

Let G be any finite group. Then $x^{|G|} = 1_G$ for any $x \in G$.

Keep this result in mind! We'll prove it later in the generality of Theorem 3.4.1.

§1.5 Subgroups

Prototypical example for this section: $\text{SL}_n(\mathbb{R})$ is a subgroup of $\text{GL}_n(\mathbb{R})$.

Earlier we saw that $\text{GL}_n(\mathbb{R})$, the $n \times n$ matrices with nonzero determinant, formed a group under matrix multiplication. But we also saw that a subset of $\text{GL}_n(\mathbb{R})$, namely $\text{SL}_n(\mathbb{R})$, also formed a group with the same operation. For that reason we say that $\text{SL}_n(\mathbb{R})$ is a subgroup of $\text{GL}_n(\mathbb{R})$. And this definition generalizes in exactly the way you expect.

Definition 1.5.1. Let $G = (G, \star)$ be a group. A **subgroup** of G is exactly what you would expect it to be: a group $H = (H, \star)$ where H is a subset of G . It's a **proper subgroup** if $H \neq G$.

Remark 1.5.2. To specify a group G , I needed to tell you both what the set G was and the operation $\star$ was. But to specify a subgroup H of a given group G , I only need to tell you who its elements are: the operation of H is just inherited from the operation of G .

Example 1.5.3 (Examples of subgroups)

- (a) $2\mathbb{Z}$ is a subgroup of $\mathbb{Z}$, which is isomorphic to $\mathbb{Z}$ itself!
- (b) Consider again S_n , the symmetric group on n elements. Let T be the set of permutations $\tau : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ for which $\tau(n) = n$. Then T is a subgroup of S_n ; in fact, it is isomorphic to S_{n-1} .
- (c) Consider the group $G \times H$ (Example 1.1.16) and the elements $\{(g, 1_H) \mid g \in G\}$. This is a subgroup of $G \times H$ (why?). In fact, it is isomorphic to G by the isomorphism $(g, 1_H) \mapsto g$.

Example 1.5.4 (Stupid examples of subgroups)

For any group G , the trivial group $\{1_G\}$ and the entire group G are subgroups of G .

Next is an especially important example that we'll talk about more in later chapters.

Example 1.5.5 (Subgroup generated by an element)

Let x be an element of a group G . Consider the set

$$\langle x \rangle = \{\dots, x^{-2}, x^{-1}, 1, x, x^2, \dots\}.$$

This is also a subgroup of G , called the subgroup generated by x .

Exercise 1.5.6. If $\text{ord } x = 2015$, what is the above subgroup equal to? What if $\text{ord } x = \infty$?

Finally, we present some non-examples of subgroups.

Example 1.5.7 (Non-Examples of subgroups)

Consider the group $\mathbb{Z} = (\mathbb{Z}, +)$.

- (a) The set $\{0, 1, 2, \dots\}$ is not a subgroup of $\mathbb{Z}$ because it does not contain inverses.
- (b) The set $\{n^3 \mid n \in \mathbb{Z}\} = \{\dots, -8, -1, 0, 1, 8, \dots\}$ is not a subgroup because it is not closed under addition; the sum of two cubes is not in general a cube.
- (c) The empty set $\emptyset$ is not a subgroup of $\mathbb{Z}$ because it lacks an identity element.

§1.6 Groups of small orders

Just for fun, here is a list of all groups of order less than or equal to ten (up to isomorphism, of course).

1. The only group of order 1 is the trivial group.
2. The only group of order 2 is $\mathbb{Z}_2$.
3. The only group of order 3 is $\mathbb{Z}_3$.
4. The only groups of order 4 are
 - $\mathbb{Z}_4$, the cyclic group on four elements,
 - $\mathbb{Z}_2 \times \mathbb{Z}_2$, called the Klein Four Group.
5. The only group of order 5 is $\mathbb{Z}_5$.
6. The groups of order six are
 - $\mathbb{Z}_6$, the cyclic group on six elements.
 - S_3 , the permutation group of three elements. This is the first non-abelian group.

Some of you might wonder where $\mathbb{Z}_2 \times \mathbb{Z}_3$ is. All I have to say is: Chinese remainder theorem!

You might wonder where D_6 is in this list. It's actually isomorphic to S_3 .

7. The only group of order 7 is $\mathbb{Z}_7$.
8. The groups of order eight are more numerous.
 - $\mathbb{Z}_8$, the cyclic group on eight elements.
 - $\mathbb{Z}_4 \times \mathbb{Z}_2$.
 - $\mathbb{Z}_2 \times \mathbb{Z}_2 \times \mathbb{Z}_2$.
 - D_8 , the dihedral group with eight elements, which is not abelian.
 - A non-abelian group Q_8 , called the *quaternion group*. It consists of eight elements $\pm 1, \pm i, \pm j, \pm k$ with $i^2 = j^2 = k^2 = ijk = -1$.
9. The groups of order nine are
 - $\mathbb{Z}_9$, the cyclic group on nine elements.
 - $\mathbb{Z}_3 \times \mathbb{Z}_3$.
10. The groups of order 10 are
 - $\mathbb{Z}_{10} \cong \mathbb{Z}_5 \times \mathbb{Z}_2$ (again Chinese remainder theorem).
 - D_{10} , the dihedral group with 10 elements. This group is non-abelian.

§1.7 Problems to think about

Problem 1A. What is the joke in the following figure? (Source: [Ge].)



Problem 1B. Prove Lagrange's theorem for orders in the special case that G is a finite abelian group.

Problem 1C. Show that $D_6 \cong S_3$ but $D_{24} \not\cong S_4$.

Problem 1D*. Let p be a prime. Show that the only group of order p is $\mathbb{Z}_p$.

🌶️ **Problem 1E[†]** (Cayley's theorem). Let G be a finite group.¹ Show that there exists a positive integer n such that

- (a) (Cayley's theorem) G is isomorphic to some subgroup of the symmetric group S_n .
- (b) (Representation Theory) G is isomorphic to some subgroup of the general linear group $\text{GL}_n(\mathbb{R})$. (This is the group of invertible $n \times n$ matrices.)

🌶️ **Problem 1F** (IMO SL 2005 C5). There are n markers, each with one side white and the other side black. In the beginning, these n markers are aligned in a row so that their white sides are all up. In each step, if possible, we choose a marker whose white side is up (but not one of the outermost markers), remove it, and reverse the closest marker to the left of it and also reverse the closest marker to the right of it.

Prove that if $n \equiv 1 \pmod{3}$ it's impossible to reach a state with only two markers remaining. (In fact the converse is true as well.)

¹ In other words, permutation groups can be arbitrarily weird. I remember being highly unsettled by this theorem when I first heard of it, but in hindsight it is not so surprising.

2 What is a space?

At the time of writing, I'm convinced that metric topology is the morally correct way to motivate point-set topology as well as to generalize normal calculus. Also, "metric" is a fun word to say. So here is my best attempt.

The concept of a metric space is very "concrete", and lends itself easily to visualization. Hence throughout this chapter you should draw lots of pictures as you learn about new objects, like convergent sequences, open sets, closed sets, and so on.

§2.1 Definition and examples of metric spaces

Prototypical example for this section: $\mathbb{R}^2$, with the Euclidean metric.

Definition 2.1.1. A **metric space** is a pair (M, d) consisting of a set of points M and a **metric** $d : M \times M \rightarrow \mathbb{R}_{\geq 0}$. The distance function must obey:

- For any $x, y \in M$, we have $d(x, y) = d(y, x)$; i.e. d is symmetric.
- The function d must be **positive definite** which means that $d(x, y) \geq 0$ with equality if and only if $x = y$.
- The function d should satisfy the **triangle inequality**: for all $x, y, z \in M$,

$$d(x, z) + d(z, y) \geq d(x, y).$$

Abuse of Notation 2.1.2. Just like with groups, we will abbreviate (M, d) as just M .

Example 2.1.3 (Metric spaces of $\mathbb{R}$)

- (a) The real line $\mathbb{R}$ is a metric space under the metric $d(x, y) = |x - y|$.
- (b) The interval $[0, 1]$ is also a metric space with the same distance function.
- (c) In fact, any subset S of $\mathbb{R}$ can be made into a metric space in this way.

Example 2.1.4 (Metric spaces of $\mathbb{R}^2$)

- (a) We can make $\mathbb{R}^2$ into a metric space by imposing the Euclidean distance function

$$d((x_1, y_1), (x_2, y_2)) = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}.$$

- (b) Just like with the first example, any subset of $\mathbb{R}^2$ also becomes a metric space after we inherit it. The unit disk, unit circle, and the unit square $[0, 1]^2$ are special cases.

Example 2.1.5 (Taxicab on $\mathbb{R}^2$)

It is also possible to place the **taxicab distance** on $\mathbb{R}^2$:

$$d((x_1, y_1), (x_2, y_2)) = |x_1 - x_2| + |y_1 - y_2|.$$

For now, we will use the more natural Euclidean metric.

Example 2.1.6 (Metric spaces of $\mathbb{R}^n$)

We can generalize the above examples easily. Let n be a positive integer.

- (a) We let $\mathbb{R}^n$ be the metric space whose points are points in n -dimensional Euclidean space, and whose metric is the Euclidean metric

$$d((a_1, \dots, a_n), (b_1, \dots, b_n)) = \sqrt{(a_1 - b_1)^2 + \dots + (a_n - b_n)^2}.$$

This is the n -dimensional **Euclidean space**.

- (b) The open **unit ball** B^n is the subset of $\mathbb{R}^n$ consisting of those points $(x_1, \dots, x_n)$ such that $x_1^2 + \dots + x_n^2 < 1$.
- (c) The **unit sphere** S^{n-1} is the subset of $\mathbb{R}^n$ consisting of those points $(x_1, \dots, x_n)$ such that $x_1^2 + \dots + x_n^2 = 1$, with the inherited metric. (The superscript $n - 1$ indicates that S^{n-1} is an $n - 1$ dimensional space, even though it lives in n -dimensional space.) For example, $S^1 \subseteq \mathbb{R}^2$ is the unit circle, whose distance between two points is the length of the chord joining them. You can also think of it as the “boundary” of the unit ball B^n .

Example 2.1.7 (Function space)

We can let M be the space of integrable functions $f : [0, 1] \rightarrow \mathbb{R}$ and define the metric by $d(f, g) = \int_0^1 |f - g| \, dx$.

Here is a slightly more pathological example.

Example 2.1.8 (Discrete space)

Let S be any set of points (either finite or infinite). We can make S into a **discrete space** by declaring

$$d(x, y) = \begin{cases} 1 & \text{if } x \neq y \\ 0 & \text{if } x = y. \end{cases}$$

If $|S| = 4$ you might think of this space as the vertices of a regular tetrahedron, living in $\mathbb{R}^3$. But for larger S it's not so easy to visualize...

Example 2.1.9 (Graphs are metric spaces)

Any connected simple graph G can be made into a metric space by defining the distance between two vertices to be the graph-theoretic distance between them. (The discrete metric is the special case when G is the complete graph on S .)

Question 2.1.10. Check the conditions of a metric space for the metrics on the discrete space and for the connected graph.

Abuse of Notation 2.1.11. From now on, we will refer to $\mathbb{R}^n$ with the Euclidean metric by just $\mathbb{R}^n$. Moreover, if we wish to take the metric space for a subset $S \subseteq \mathbb{R}^n$ with the inherited metric, we will just write S .

§2.2 Convergence in metric spaces

Prototypical example for this section: The sequence $\frac{1}{n}$ (for $n = 1, 2, \dots$) in $\mathbb{R}$.

Since we can talk about the distance between two points, we can talk about what it means for a sequence of points to converge. This is the same as the typical epsilon-delta definition, with absolute values replaced by the distance function.

Definition 2.2.1. Let $(x_n)_{n \geq 1}$ be a sequence of points in a metric space M . We say that x_n **converges** to x if the following condition holds: for all $\varepsilon > 0$, there is an integer N (depending on ε) such that $d(x_n, x) < \varepsilon$ for each $n \geq N$. This is written

$$x_n \rightarrow x$$


or more verbosely as

$$\lim_{n \rightarrow \infty} x_n = x.$$

We say that a sequence converges in M if it converges to a point in M .

You should check that this definition coincides with your intuitive notion of “converges”.

Abuse of Notation 2.2.2. If the parent space M is understood, we will allow ourselves to abbreviate “converges in M ” to just “converges”. However, keep in mind that convergence is defined relative to the parent space; the “limit” of the space must actually be a point in M for a sequence to converge.



Example 2.2.3

Consider the sequence $x_1 = 1, x_2 = 1.4, x_3 = 1.41, x_4 = 1.414, \dots$

- (a) If we view this as a sequence in $\mathbb{R}$, it converges to $\sqrt{2}$.
- (b) However, even though each x_i is in $\mathbb{Q}$, this sequence does NOT converge when we view it as a sequence in $\mathbb{Q}$!

Question 2.2.4. What are the convergent sequences in a discrete metric space?

§2.3 Continuous maps

Abuse of Notation 2.3.1. For a function f and its argument x , we will begin abbreviating $f(x)$ to just fx when there is no risk of confusion.

In calculus you were also told (or have at least heard) of what it means for a function to be continuous. Probably something like

A function $f : \mathbb{R} \rightarrow \mathbb{R}$ is continuous at a point $p \in \mathbb{R}$ if for every $\varepsilon > 0$ there exists a $\delta > 0$ such that $|x - p| < \delta \implies |fx - fp| < \varepsilon$.

Question 2.3.2. Can you guess what the corresponding definition for metric spaces is?

All we have to do is replace the absolute values with the more general distance functions: this gives us a definition of continuity for any function $M \rightarrow N$.

Definition 2.3.3. Let $M = (M, d_M)$ and $N = (N, d_N)$ be metric spaces. A function $f : M \rightarrow N$ is **continuous** at a point $p \in M$ if for every $\varepsilon > 0$ there exists a $\delta > 0$ such that

$$d_M(x, p) < \delta \implies d_N(fx, fp) < \varepsilon.$$

Moreover, the entire function f is continuous if it is continuous at every point $p \in M$.

Notice that, just like in our definition of an isomorphism of a group, we use the metric of M for one condition and the metric of N for the other condition.

This generalization is nice because it tells us immediately how we could carry over continuity arguments in $\mathbb{R}$ to more general spaces (for example, replacing $\mathbb{R}$ with $\mathbb{C}$ to get complex analysis). Nonetheless, this definition is kind of cumbersome to work with, because it makes extensive use of the real numbers (epsilons and deltas). Here is an equivalent condition.

Theorem 2.3.4 (Sequential continuity)

A function $f : M \rightarrow N$ of metric spaces is continuous at a point $p \in M$ if and only if the following property holds: if $x_1, x_2, \dots$ is a sequence in M converging to p , then the sequence $f(x_1), f(x_2), \dots$ in N converges to $f(p)$.

Proof. It's not too hard to see that ε - δ continuity implies sequential continuity. The reverse direction is trickier and left as Problem 2C. $\square$

The next example illustrates why this criterion can often be much easier to work with.

Proposition 2.3.5 (Composition of continuous functions is continuous)

Let $f : M \rightarrow N$ and $g : N \rightarrow L$ be continuous maps of metric spaces. Then their composition $g \circ f$ is continuous.

Proof. Dead simple with sequences: Let $p \in M$ be arbitrary and let $x_n \rightarrow p$ in M . Then $fx_n \rightarrow fp$ in N and $gfx_n \rightarrow gfp$ in L , QED. $\square$

I hope you'll agree this is much cleaner than having to deal with ε 's and δ 's.

Question 2.3.6. Let M be any metric space and D a discrete space. When is a map $f : D \rightarrow M$ continuous?

§2.4 Homeomorphisms

When do we consider two groups to be the same? Answer: if there's a structure-preserving map between them which is also a bijection. For metric spaces, we do exactly the same thing, but replace “structure-preserving” with “continuous”.

Definition 2.4.1. Let M and N be metric spaces. A function $f : M \rightarrow N$ is a **homeomorphism** or **bi-continuous function** if it is a bijection, and both $f : M \rightarrow N$ and its inverse $f^{-1} : N \rightarrow M$ are continuous. We say M and N are **homeomorphic**.

Needless to say, homeomorphism is an equivalence relation.

You might be surprised that we require f^{-1} to also be continuous. Here's the reason: you can show that if ϕ is an isomorphism of groups, then ϕ^{-1} also preserves the group operation, hence ϕ^{-1} is itself an isomorphism. The same is not true for continuous bijections, which is why we need the new condition.


Example 2.4.2 (Homeomorphism $\neq$ continuous bijection)(a) There is a continuous bijection from $[0, 1]$ to the circle, but it has no continuous inverse.

(b) Let M be a discrete space with size $|\mathbb{R}|$. Then there is a continuous function $M \rightarrow \mathbb{R}$ which certainly has no continuous inverse.

Note that this is the topologist's definition of “same” – homeomorphisms are “continuous deformations”. Here are some examples.

Example 2.4.3 (Examples of homeomorphisms)

- (a) Any space M is homeomorphic to itself through the identity map.
- (b) The old saying: a doughnut (torus) is homeomorphic to a coffee cup. (Look this up if you haven't heard of it.)
- (c) The unit circle S^1 is homeomorphic to the boundary of the unit square. Here's one bijection between them, after an appropriate scaling:



Example 2.4.4 (Metrics on the unit circle)

It may have seemed strange that our metric function on S^1 was the one inherited from $\mathbb{R}^2$, meaning the distance between two points on the circle was defined to be the length of the chord. Wouldn't it have made more sense to use the circumference of the arc joining the two points? In fact, it doesn't matter: if we consider S^1 with the “chord” metric and the “arc” metric, we get two homeomorphic spaces.

The same goes for S^{n-1} for general n .

Example 2.4.5 (Homeomorphisms really don't preserve size)

Surprisingly, the open interval $(0, 1)$ is homeomorphic to the real line $\mathbb{R}$! This might come as a surprise, since $(0, 1)$ doesn't look that much like $\mathbb{R}$; the former is “bounded” while the latter is “unbounded”.

Exercise 2.4.6. Write down a homeomorphism from $(0, 1)$ to $\mathbb{R}$.

Example 2.4.7 (Product metrics)

Let $M = (M, d_M)$ and $N = (N, d_N)$ be metric spaces (say, $M = N = \mathbb{R}$). Let $p_i = (x_i, y_i) \in M \times N$ for $i = 1, 2$. Consider the following metrics on the set of points $M \times N$:

- $d_{\max}(p_1, p_2) = \max\{d_M(x_1, x_2), d_N(y_1, y_2)\}$.
- $d_{\text{Euclid}}(p_1, p_2) = \sqrt{d_M(x_1, x_2)^2 + d_N(y_1, y_2)^2}$.
- $d_{\text{taxicab}}(p_1, p_2) = d_M(x_1, x_2) + d_N(y_1, y_2)$.

It's easy to verify that

$$d_{\max}(p_1, p_2) \leq d_{\text{Euclid}}(p_1, p_2) \leq d_{\text{taxicab}}(p_1, p_2) \leq 2d_{\max}(p_1, p_2).$$

Using this you can show that

$$(M \times N, d_{\max}), \quad (M \times N, d_{\text{Euclid}}), \quad (M \times N, d_{\text{taxicab}})$$

are homeomorphic, with the homeomorphism being just the identity map. Hence we will usually simply refer to *the* metric on $M \times N$, called the **product metric**, and it will not be important which metric we select.

§2.5 Open sets


Prototypical example for this section: The open disk $x^2 + y^2 < r^2$ in $\mathbb{R}^2$.

Continuity is really about what happens “locally”: how a function behaves “close to a certain point p ”. One way to capture this notion of “closeness” is to use metrics as we've done above. In this way we can define a neighborhood of a point.

Definition 2.5.1. Let M be a metric space. For each real number $r > 0$ and point $p \in M$, we define

$$M_r(p) \stackrel{\text{def}}{=} \{x \in M : d(x, p) < r\}.$$

The set $M_r(p)$ is called an **r -neighborhood** of p .



We can rephrase convergence more succinctly in terms of r -neighborhoods. Specifically, a sequence (x_n) converges to x if for every r -neighborhood of x , all terms of x_n eventually stay within that r -neighborhood.

Let's try to do the same with functions.

Question 2.5.2. In terms of r -neighborhoods, what does it mean for a function $f : M \rightarrow N$ to be continuous at a point $p \in M$?

Essentially, we require that the pre-image of every ε -neighborhood has the property that some δ -neighborhood exists inside it. This motivates:

Definition 2.5.3. A set $U \subseteq M$ is *open* in M if for each $p \in U$, some r -neighborhood of p is contained inside U . In other words, there exists $r > 0$ such that $M_r(p) \subseteq U$.

Abuse of Notation 2.5.4. Note that a set being open is defined *relative to* the parent space M . However, if M is understood we can abbreviate “open in M ” to just “open”.




Figure 2.1: The set of points $x^2 + y^2 < 1$ in $\mathbb{R}^2$ is open in $\mathbb{R}^2$.

Example 2.5.5 (Examples of open sets)

- (a) Any r -neighborhood of a point is open.
- (b) Open intervals of $\mathbb{R}$ are open in $\mathbb{R}$, hence the name! This is the prototypical example to keep in mind.
- (c) The open unit ball B^n is open in $\mathbb{R}^n$ for the same reason.
- (d) In particular, the open interval $(0, 1)$ is open in $\mathbb{R}$. However, if we embed it in $\mathbb{R}^2$, it is no longer open!
- (e) The empty set $\emptyset$ and the whole set of points M are open in M .

Example 2.5.6 (Non-Examples of open sets)

- (a) The closed interval $[0, 1]$ is not open in $\mathbb{R}$. There is no neighborhood of the point 0 which is contained in $[0, 1]$.
- (b) The unit circle S^1 is not open in $\mathbb{R}^2$.

Question 2.5.7. What are the open sets of the discrete space?

Here are two quite important properties of open sets.

Proposition 2.5.8

- (a) The intersection of finitely many open sets is open.
- (b) The union of open sets is open, even if there are infinitely many.

Question 2.5.9. Convince yourself this is true.

Exercise 2.5.10. Exhibit an infinite collection of open sets in $\mathbb{R}$ whose intersection is the set $\{0\}$. This implies that infinite intersections of open sets are not necessarily open.

The whole upshot of this is:


Theorem 2.5.11 (Open set condition)

A function $f : M \rightarrow N$ of metric spaces is continuous if and only if the pre-image of every open set in N is open in M .

Proof. I'll just do one direction...

Exercise 2.5.12. Show that δ - ε continuity follows from the open set continuity.

Now assume f is continuous. First, suppose V is an open subset of the metric space N ; let $U = f^{\text{pre}}(V)$. Pick $x \in U$, so $y = f(x) \in V$; we want a neighborhood of x inside U .



As V is open, there is some small ε -neighborhood around y which is contained inside V . By continuity of f , we can find a δ such that the δ -neighborhood of x gets mapped by f into the ε -neighborhood in N , which in particular lives inside V . Thus the δ -neighborhood lives in V , as desired. $\square$

From this we can get a new definition of homeomorphism which makes it clear why open sets are good things to consider.

Theorem 2.5.13

A function $f : M \rightarrow N$ of metric spaces is a homeomorphism if

- (i) It is a bijection of the underlying points.
- (ii) It induces a bijection of the open sets of M and N : for any open set $U \subseteq M$ the set $f(U)$ is open, and for any open set $V \subseteq N$ the set $f^{\text{pre}}(V)$ is open.

This leads us to...

§2.6 Forgetting the metric

Notice something interesting about the previous theorem – it doesn't reference the metrics of M and N at all. Instead, it refers only to the open sets.

This leads us to consider: what if we could refer to spaces *only* by their open sets, forgetting about the fact that we had a metric to begin with? That's exactly what we do in “point-set topology”.

Definition 2.6.1. A **topological space** is a pair $(X, \mathcal{T})$, where X is a set of points, and $\mathcal{T}$ is the **topology**, which consists of several subsets of X , called the **open sets** of X . The topology must obey the following axioms.

- $\emptyset$ and X are both in $\mathcal{T}$.
- Finite intersections of open sets are also in $\mathcal{T}$.
- Arbitrary unions (possibly infinite) of open sets are also in $\mathcal{T}$.

So this time, the open sets are *given*. Rather than defining a metric and getting open sets from the metric, we instead start from just the open sets.


Abuse of Notation 2.6.2. We refer to the space $(X, \mathcal{T})$ by just X . (Do you see a pattern here?)

Example 2.6.3 (Examples of topologies)

- Given a metric space M , we can let $\mathcal{T}$ be the open sets in the metric sense. The point is that the axioms are satisfied.
- In particular, **discrete space** is a topological space in which every set is open. (Why?)
- Given X , we can let $\mathcal{T} = \{\emptyset, X\}$, the opposite extreme of the discrete space.

Now we can port over our metric definitions.

Definition 2.6.4. An **open neighborhood** of a point $x \in X$ is an open set U which contains x (see figure).



Abuse of Notation 2.6.5. Just to be perfectly clear: by a “open neighborhood” I mean *any* open set containing x . But by an “ r -neighborhood” I always mean the points with distance less than r from x , and so I can only use this term if my space is a metric space.

Abuse of Notation 2.6.6. There's another related term commonly used: a *neighborhood* V of x is a set which contains some open neighborhood of x (often V itself). Think of it as “open around x ”, though not always at other points. However, for most purposes, you should think of neighborhoods as just open neighborhoods.

Definition 2.6.7. A function $f : X \rightarrow Y$ of topological spaces is **continuous** at $p \in X$ if the pre-image of any open neighborhood of $f(p)$ is an open neighborhood of p . It is continuous if it is continuous at every point, meaning that the pre-image of any open set is open.

You can also port over the notion of sequences and convergent sequences. But I won't bother to do so because sequences lose most of the nice properties they had before.

Finally, what are the homeomorphisms? The same definition carries over: a bijection which is continuous in both directions.

Definition 2.6.8. A **homeomorphism** of topological spaces (X, τ_X) and (Y, τ_Y) is a bijection from X to Y which induces a bijection from τ_X to τ_Y : i.e. the bijection preserves open sets.

Therefore, any property defined only in terms of open sets is preserved by homeomorphism. Such a property is called a **topological property**. That's why $(0, 1)$ homeomorphic to $\mathbb{R}$ is not so surprising, because the notion of being “bounded” is not a notion which can be expressed in terms of open sets.

Remark 2.6.9. As you might have guessed, there exist topological spaces which cannot be realized as metric spaces (in other words, are not **metrizable**). One example is just to take $X = \{a, b, c\}$ and the topology $\tau_X = \{\emptyset, \{a, b, c\}\}$. This topology is fairly “stupid”: it can't tell apart any of the points a, b, c ! But any metric space can tell its points apart (because $d(x, y) > 0$ when $x \neq y$). We'll see less trivial examples later.

§2.7 Closed sets

Prototypical example for this section: The closed unit disk $x^2 + y^2 \leq r^2$ in $\mathbb{R}^2$.

It would be criminal for me to talk about open sets without talking about the dual concept, a closed set. The name “closed” comes from the definition in a metric space.

Definition 2.7.1. Let M be a metric space. A subset $S \subseteq M$ is **closed** in M if the following property holds: let $x_1, x_2, \dots$ be a sequence of points in S and suppose that x_n converges to x in M . Then $x \in S$ as well.

Abuse of Notation 2.7.2. Same caveat: we abbreviate “closed in M ” to just “closed” if the parent space M is understood.

Here's another way to phrase it. The **limit points** of a subset $S \subseteq M$ are defined by

$$\lim S \stackrel{\text{def}}{=} \{p \in M : \exists (x_n) \in S \text{ such that } x_n \rightarrow p\}.$$

Thus S is closed if and only if $S = \lim S$.

Exercise 2.7.3. Prove that $\lim S$ is closed even if S isn't closed. (Draw a picture.)

For this reason, $\lim S$ is also called the **closure** of S in M , and denoted $\bar{S}$. It is simply the smallest closed set which contains S .

Example 2.7.4 (Examples of closed sets)

- (a) The empty set $\emptyset$ is closed in M for vacuous reasons: there are no sequences of points with elements in $\emptyset$.
- (b) The entire space M is closed in M for tautological reasons. (Verify this!)
- (c) The closed interval $[0, 1]$ in $\mathbb{R}$ is closed in $\mathbb{R}$, hence the name. Like with open sets, this is the prototypical example of a closed set to keep in mind!
- (d) In fact, the closed interval $[0, 1]$ is even closed in $\mathbb{R}^2$.

Example 2.7.5 (Non-Examples of closed sets)

Let $S = (0, 1)$ denote the open interval. Then S is not closed in $\mathbb{R}$ because the sequence of points

$$\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \dots$$

converges to $0 \in \mathbb{R}$, but $0 \notin (0, 1)$.

In what sense are these concepts “dual”? Despite first impressions, most sets are neither open nor closed.

Example 2.7.6 (A set neither open nor closed)

The half-open interval $[0, 1)$ is neither open nor closed in $\mathbb{R}$.

Remark 2.7.7. It’s also possible for a set to be both open and closed; this will be discussed in Chapter 4.

Remarkably, though, the following *is* true.

Theorem 2.7.8 (Closed sets are complements of open sets)

Let M be a metric space, and $S \subseteq M$ any subset. Then the following are equivalent:

- The set S is closed in M .
- The complement $M \setminus S$ is open in M .

Exercise 2.7.9. Prove this theorem! You’ll want to draw a picture to make it clear what’s happening: for example, you might take $M = \mathbb{R}^2$ and S to be the closed unit disk.

This leads us to a definition for a general topological space.

Definition 2.7.10. In a general topological space X , we say that $S \subseteq X$ is **closed** in X if the complement $X \setminus S$ is open in X .

Hence, for general topological spaces, open and closed sets carry the same information, and it is entirely a matter of taste whether we define everything in terms of open sets or closed sets. In particular,

Question 2.7.11. Show that the (possibly infinite) intersection of closed sets is closed while the union of finitely many closed sets is closed. (Hint: just look at complements.)

Question 2.7.12. Show that a function is continuous if and only if the pre-image of every closed set is closed.

Mathematicians seem to have agreed that they like open sets better.

§2.8 Common pitfalls

An important distinction to keep in mind is that

Convergence and open/closed sets are defined *relative to a parent space*. Therefore, it makes no sense to ask a question like “is $[0, 1]$ open?”.

For example, here are some gotchas:

- Consider the sequence $1, 1.4, 1.41, 1.414, \dots$. Viewed as a sequence in $\mathbb{R}$, it converges to $\sqrt{2}$. But if viewed as a sequence in $\mathbb{Q}$, this sequence does *not* converge! For a sequence to converge one has to be able to write down where it approaches, but $\sqrt{2} \notin \mathbb{Q}$. Similarly, the sequence $0.9, 0.99, 0.999, 0.9999$ does not converge in the space $(0, 1)$, although it does converge in $[0, 1]$. In general it doesn’t make sense to ask “does (x_n) converge” without specifying which space one is discussing.

The fact that these sequences fail to converge even though they “ought to” is weird and bad, and so in Chapter 4 we’ll define a *complete* metric space as one where sequences that “should” converge (called *Cauchy*) actually do.

- In general, it makes no sense to ask a question like “is $[0, 1]$ open?”. The questions “is $[0, 1]$ open in $\mathbb{R}$?” and “is $[0, 1]$ open in $[0, 1]$?” do make sense, however. The answer to the first question is “no” but the answer to the second question is “yes”; indeed, every space is open in itself. Similarly, $[0, \frac{1}{2})$ is an open set in the space $M = [0, 1]$ because it is the ball *in* M of radius $\frac{1}{2}$ centered at 0.
- Dually, it doesn’t make sense to ask “is $[0, 1]$ closed”? It is closed *in* $\mathbb{R}$ and *in itself* (but every space is closed in itself, anyways).

To make sure you understand the above, you should convince yourself that:

Exercise 2.8.1. Let $M = [0, 1] \cup (2, 3)$. Show that $[0, 1]$ and $(2, 3)$ are both open and closed in M .

This illustrates a third point: a nontrivial set can be both open and closed (apologies for the terminology). (As we’ll see in Chapter 4, this implies the space is disconnected; i.e. the only examples look quite like the one we’ve given above.)

§2.9 Problems to think about

Problem 2A. Exhibit a function $f : \mathbb{R} \rightarrow \mathbb{R}$ such that f is continuous at $x \in \mathbb{R}$ if and only if $x = 0$.

Problem 2B (Furstenberg). We declare a subset of $\mathbb{Z}$ to be open if it's the union (possibly empty or infinite) of arithmetic sequences $\{a + nd \mid n \in \mathbb{Z}\}$, where a and d are positive integers.

(a) Verify this forms a topology on $\mathbb{Z}$, called the **evenly spaced integer topology**.

(b) Prove there are infinitely many primes by considering $\bigcup_p p\mathbb{Z}$.

Problem 2C. Show that sequentially continuous at p implies ε - δ continuous at p , as in Theorem 2.3.4.



Problem 2D. Prove that a function $f : \mathbb{R} \rightarrow \mathbb{R}$ which is strictly increasing must be continuous at some point.



Problem 2E. Prove that the evenly spaced integer topology on $\mathbb{Z}$ is metrizable. In other words, show that one can impose a metric $d : \mathbb{Z}^2 \rightarrow \mathbb{R}$ which makes $\mathbb{Z}$ into a metric space whose open sets are those described above.

DRAFT (Evan Chen)
Updated February 15, 2017

3 Homomorphisms and quotient groups

§3.1 Generators and group presentations

Let G be a group. Recall that for some element $x \in G$, we could consider the subgroup

$$\{\dots, x^{-2}, x^{-1}, 1, x, x^2, \dots\}$$

of G . Here's a more pictorial version of what we did: **put x in a box, seal it tightly, and shake vigorously**. Using just the element x , we get a pretty explosion that produces the subgroup above.

What happens if we put two elements x, y in the box? Among the elements that get produced are things like

$$xyxyx, \quad x^2y^9x^{-5}y^3, \quad y^{-2015}, \quad \dots$$

Essentially, I can create any finite product of x, y, x^{-1}, y^{-1} . This leads us to define:

Definition 3.1.1. Let S be a subset of G . The subgroup **generated** by S , denoted $\langle S \rangle$, is the set of elements which can be written as a finite product of elements in S (and their inverses). If $\langle S \rangle = G$ then we say S is a set of **generators** for G , as the elements of S together create all of G .

Exercise 3.1.2. Why is the condition “and their inverses” not necessary if G is a finite group? (As usual, assume Lagrange's theorem.)

Example 3.1.3 ($\mathbb{Z}$ is the infinite cyclic group)

Consider 1 as an element of $\mathbb{Z} = (\mathbb{Z}, +)$. We see $\langle 1 \rangle = \mathbb{Z}$, meaning $\{1\}$ generates $\mathbb{Z}$. It's important that -1 , the inverse of 1 is also allowed: we need it to write all integers as the sum of 1 and -1 .

This gives us an idea for a way to try and express groups compactly. Why not just write down a list of generators for the groups? For example, we could write

$$\mathbb{Z} \cong \langle a \rangle$$

meaning that $\mathbb{Z}$ is just the group generated by one element.

There's one issue: the generators usually satisfy certain properties. For example, consider $\mathbb{Z}_{100}$. It's also generated by a single element x , but this x has the additional property that $x^{100} = 1$. This motivates us to write

$$\mathbb{Z}_{100} = \langle x \mid x^{100} = 1 \rangle.$$

I'm sure you can see where this is going. All we have to do is specify a set of generators and **relations** between the generators, and say that two elements are equal if and only if you can get from one to the other using relations. Such an expression is appropriately called a **group presentation**.

Example 3.1.4 (Dihedral group)

The dihedral group of order $2n$ has a presentation

$$D_{2n} = \langle r, s \mid r^n = s^2 = 1, rs = sr^{-1} \rangle.$$

Thus each element of D_{2n} can be written uniquely in the form r^α or sr^α , where $\alpha = 0, 1, \dots, n-1$.

Example 3.1.5 (Klein four group)

The **Klein four group**, isomorphic to $\mathbb{Z}_2 \times \mathbb{Z}_2$, is given by the presentation

$$\langle a, b \mid a^2 = b^2 = 1, ab = ba \rangle.$$

Example 3.1.6 (Free group)

The **free group on n elements** is the group whose presentation has n generators and no relations at all. It is denoted F_n , so

$$F_n = \langle x_1, x_2, \dots, x_n \rangle.$$

In other words, $F_2 = \langle a, b \rangle$ is the set of strings formed by appending finitely many copies of a, b, a^{-1}, b^{-1} together.

Question 3.1.7. Notice that $F_1 \cong \mathbb{Z}$.

Abuse of Notation 3.1.8. One might unfortunately notice that “subgroup generated by a and b ” has exactly the same notation as the free group $\langle a, b \rangle$. We’ll try to be clear based on context which one we mean.

Presentations are nice because they provide a compact way to write down groups. They do have some shortcomings, though.¹

Example 3.1.9 (Presentations can look very different)

The same group can have very different presentations. For instance consider

$$D_{2n} = \langle x, y \mid x^2 = y^2 = 1, (xy)^n = 1 \rangle.$$

(To see why this is equivalent, set $x = s, y = rs$.)

§3.2 Homomorphisms

Prototypical example for this section: The “mod out by 100” map, $\mathbb{Z} \rightarrow \mathbb{Z}_{100}$.

How can groups talk to each other?

¹ Actually, determining whether two elements of a presentation are isomorphic is undecidable. In fact, it is undecidable to even determine if a group is finite from its presentation.

Two groups are “the same” if we can write an isomorphism between them. And as we saw, two metric spaces are “the same” if we can write a homeomorphism between them. But what’s the group analogy of a continuous map? We simply drop the “bijection” condition.

Definition 3.2.1. Let $G = (G, \star)$ and $H = (H, *)$ be groups. A **group homomorphism** is a map $\phi : G \rightarrow H$ such that for any $g_1, g_2 \in G$ we have

$$\phi(g_1 \star g_2) = \phi(g_1) * \phi(g_2).$$

Example 3.2.2 (Examples of homomorphisms)

Let G and H be groups.

- (a) Any isomorphism $G \rightarrow H$ is a homomorphism. In particular, the identity map $G \rightarrow G$ is a homomorphism.
- (b) The **trivial homomorphism** $G \rightarrow H$ sends everything to 1_H .
- (c) There is a homomorphism from $\mathbb{Z}$ to $\mathbb{Z}_{100}$ by sending each integer to its residue modulo 100.
- (d) There is a homomorphism from $\mathbb{Z}$ to itself by $x \mapsto 10x$ which is injective but not surjective.
- (e) There is a homomorphism from S_n to S_{n+1} by “embedding”: every permutation on $\{1, \dots, n\}$ can be thought of as a permutation on $\{1, \dots, n+1\}$ if we simply let $n+1$ be a fixed point.
- (f) A homomorphism $\phi : D_{12} \rightarrow D_6$ is given by $s_{12} \mapsto s_6$ and $r_{12} \mapsto r_6$.
- (g) Specifying a homomorphism $\mathbb{Z} \rightarrow G$ is the same as specifying just the image of the element $1 \in \mathbb{Z}$. Why?

The last two examples illustrates something: suppose we have a presentation of G . To specify a homomorphism $G \rightarrow H$, we only have to specify where each generator of G goes, in such a way that the relations are all satisfied.

Important remark: the right way to think about an isomorphism is as a “bijective homomorphism”. To be explicit,

Exercise 3.2.3. Show that $G \cong H$ if and only if there exist homomorphisms $\phi : G \rightarrow H$ and $\psi : H \rightarrow G$ such that $\phi \circ \psi = \text{id}_H$ and $\psi \circ \phi = \text{id}_G$.

So the definitions of homeomorphism of metric spaces and isomorphism of groups are not too different

Some obvious properties of homomorphisms follow.

Fact 3.2.4. Let $\phi : G \rightarrow H$ be a homomorphism. Then $\phi(1_G) = 1_H$ and $\phi(g^{-1}) = \phi(g)^{-1}$.

Proof. Boring, and I’m sure you could do it yourself if you wanted to. $\square$

Now let me define a very important property of a homomorphism.

Definition 3.2.5. The **kernel** of a homomorphism $\phi : G \rightarrow H$ is defined by

$$\ker \phi \stackrel{\text{def}}{=} \{g \in G : \phi(g) = 1_H\}.$$

It is a *subgroup* of G (in particular, $1_G \in \ker \phi$ for obvious reasons).

Question 3.2.6. Verify that $\ker \phi$ is in fact a subgroup of G .

We also have the following important fact, which we also encourage the reader to verify.

Proposition 3.2.7 (Kernel determines injectivity)

The map ϕ is injective if and only if $\ker \phi = \{1_G\}$.

To make this concrete, let's compute the kernel of each of our examples.

Example 3.2.8 (Examples of kernels)

- (a) The kernel of any isomorphism $G \rightarrow H$ is trivial, since an isomorphism is injective. In particular, the kernel of the identity map $G \rightarrow G$ is $\{1_G\}$.
- (b) The kernel of the trivial homomorphism $G \rightarrow H$ (by $g \mapsto 1_H$) is all of G .
- (c) The kernel of the homomorphism $\mathbb{Z} \rightarrow \mathbb{Z}_{100}$ by $n \mapsto \bar{n}$ is precisely

$$100\mathbb{Z} = \{\dots, -200, -100, 0, 100, 200, \dots\}.$$

- (d) The kernel of the map $\mathbb{Z} \rightarrow \mathbb{Z}$ by $x \mapsto 10x$ is trivial: $\{0\}$.
- (e) There is a homomorphism from S_n to S_{n+1} by “embedding”, but it also has trivial kernel because it is injective.
- (f) A homomorphism $\phi : D_{12} \rightarrow D_6$ is given by $s_{12} \mapsto s_6$ and $r_{12} \mapsto r_6$. You can check that

$$\ker \phi = \{1, r_{12}^6\} \cong \mathbb{Z}_2.$$
- (g) Exercise below.

Exercise 3.2.9. Fix any $g \in G$. Suppose we have a homomorphism $\mathbb{Z} \rightarrow G$ by $n \mapsto g^n$. What is the kernel?

Question 3.2.10. Show that for any homomorphism $\phi : G \rightarrow H$, the image $\phi(G)$ is a subgroup of H . Hence, we'll be especially interested in the case where ϕ is surjective.

§3.3 Cosets and modding out

Prototypical example for this section: Modding out by n : $\mathbb{Z}/(n \cdot \mathbb{Z}) \cong \mathbb{Z}_n$.

The next few sections are a bit dense. If this exposition doesn't work for you, try [Go11]. Let G and Q be groups, and suppose there exists a *surjective* homomorphism

$$\phi : G \twoheadrightarrow Q.$$

In other words, if ϕ is injective then $\phi : G \rightarrow Q$ is a bijection, and hence an isomorphism. But suppose we're not so lucky and $\ker \phi$ is bigger than just $\{1_G\}$. What is the correct interpretation of a more general homomorphism?

Let's look at the special case where $\phi : \mathbb{Z} \rightarrow \mathbb{Z}_{100}$ is “modding out by 100”. We already saw that the kernel of this map is

$$\ker \phi = 100\mathbb{Z} = \{\dots, -200, -100, 0, 100, 200, \dots\}.$$

Recall now that $\ker \phi$ is a subgroup of G . What this means is that ϕ is **indifferent to the subgroup $100\mathbb{Z}$ of $\mathbb{Z}$** :

$$\phi(15) = \phi(2000 + 15) = \phi(-300 + 15) = \phi(700 + 15) = \dots$$

So $\mathbb{Z}_{100}$ is what we get when we “mod out by 100”. Cool.

In other words, let G be a group and $\phi : G \rightarrow Q$ be a surjective homomorphism with kernel $N \subseteq G$.

We claim that Q should be thought of as the quotient of G by N .

To formalize this, we will define a so-called **quotient group** G/N in terms of G and N only (without referencing Q) which will be naturally isomorphic to Q .

For motivation, let's give a concrete description of Q using just ϕ and G . Continuing our previous example, let $N = 100\mathbb{Z}$ be our subgroup of G . Consider the sets

$$\begin{aligned} N &= \{\dots, -200, -100, 0, 100, 200, \dots\} \\ 1 + N &= \{\dots, -199, -99, 1, 101, 201, \dots\} \\ 2 + N &= \{\dots, -198, -98, 2, 102, 202, \dots\} \\ &\vdots \\ 99 + N &= \{\dots, -101, -1, 99, 199, 299, \dots\}. \end{aligned}$$

The elements of each set all have the same image when we apply ϕ , and moreover any two elements in different sets have different images. Then the main idea is to notice that

We can think of Q as the group whose *elements* are the *sets* above.

Thus, given ϕ we define an equivalence relation $\sim_N$ on G by saying $x \sim_N y$ for $\phi(x) = \phi(y)$. This $\sim_N$ divides G into several equivalence classes in G which are in obvious bijection with Q , as above. Now we claim that we can write these equivalence classes very explicitly.

Exercise 3.3.1. Show that $x \sim_N y$ if and only if $x = yn$ for some $n \in N$ (in the mod 100 example, this means they “differ by some multiple of 100”). Thus for any $g \in G$, the equivalence class of $\sim_N$ which contains g is given explicitly by

$$gN \stackrel{\text{def}}{=} \{gn \mid n \in N\}.$$

Here's the word that describes the types of sets we're running into now.

Definition 3.3.2. Let H be any subgroup of G (not necessarily the kernel of some homomorphism). A set of the form gH is called a **left coset** of H .

Remark 3.3.3. Although the notation might not suggest it, keep in mind that g_1N is often equal to g_2N even if $g_1 \neq g_2$. In the “mod 100” example, $3 + N = 103 + N$. In other words, these cosets are *sets*.

This means that if I write “let gH be a coset” without telling you what g is, you can’t figure out which g I chose from just the coset itself. If you don’t believe me, here’s an example of what I mean:

$$x + 100\mathbb{Z} = \{\dots, -97, 3, 103, 203, \dots\} \implies x = ?.$$


There’s no reason to think I picked $x = 3$. (I actually picked $x = -13597$.)

Remark 3.3.4. Given cosets g_1H and g_2H , you can check that the map $x \mapsto g_2g_1^{-1}x$ is a bijection between them. So actually, all cosets have the same cardinality.

So, long story short,

Elements of the group Q are naturally identified with left cosets of G .

In practice, people often still prefer to picture elements of Q as single points (for example it’s easier to think of $\mathbb{Z}_2$ as $\{0, 1\}$ rather than $\{\dots, -2, 0, 2, \dots\}, \{\dots, -1, 1, 3, \dots\}$). If you like this picture, then you might then draw G as a bunch of equally tall fibers (the cosets), which are then “collapsed” onto Q .



Now that we’ve done this, we can give an *intrinsic* definition for the quotient group we alluded to earlier.

Definition 3.3.5. A subgroup N of G is called **normal** if it is the kernel of some homomorphism. We write this as $N \trianglelefteq G$.

Definition 3.3.6. Let $N \trianglelefteq G$. Then the **quotient group**, denoted G/N , is the group defined as follows.

- The elements of G/N will be the left cosets of N .
- We want to define the product of two cosets C_1 and C_2 in G/N . Recall that the cosets are in bijection with elements of Q . So let q_1 be the value associated to the coset C_1 , and q_2 the one for C_2 . Then we can take the product to be the coset corresponding to q_1q_2 .

Quite importantly, **we can also do this in terms of representatives of the cosets**. Let $g_1 \in C_1$ and $g_2 \in C_2$, so $C_1 = g_1N$ and $C_2 = g_2N$. Then $C_1 \cdot C_2$

should be the coset which contains g_1g_2 . This is the same as the above definition since $\phi(g_1g_2) = \phi(g_1)\phi(g_2) = q_1q_2$; all we've done is define the product in terms of elements of G , rather than values in H .

Using the gN notation, and with Remark 3.3.3 in mind, we can write this even more succinctly:

$$(g_1N) \cdot (g_2N) \stackrel{\text{def}}{=} (g_1g_2)N.$$

And now you know why the integers modulo n are often written $\mathbb{Z}/n\mathbb{Z}$!

Question 3.3.7. Take a moment to digest the above definition.

By the way we've built it, the resulting group G/N is isomorphic to Q . In a sense we think of G/N as “ G modulo the condition that $n = 1$ for all $n \in N$ ”.

§3.4 (Optional) Proof of Lagrange's theorem

As an aside, with the language of cosets we can now show Lagrange's theorem in the general case.

Theorem 3.4.1 (Lagrange's theorem)

Let G be a finite group, and let H be any subgroup. Then $|H|$ divides $|G|$.

The proof is very simple: note that the cosets of H all have the same size and form a partition of G (even when H is not necessarily normal). Hence if n is the number of cosets, then $n \cdot |H| = |G|$.

Question 3.4.2. Conclude that $x^{|G|} = 1$ by taking $H = \langle x \rangle \subseteq G$.

Remark 3.4.3. It should be mentioned at this point that in general, if G is a finite group and N is normal, then $|G/N| = |G|/|N|$.

§3.5 Eliminating the homomorphism

Prototypical example for this section: Again $\mathbb{Z}/n\mathbb{Z} \cong \mathbb{Z}_n$.

Let's look at the last definition of G/N we provided. The short version is:

- The elements of G/N are cosets gN , which you can think of as equivalence classes of a relation $\sim_N$ (where $g_1 \sim_N g_2$ if $g_1 = g_2n$ for some $n \in N$).
- Given cosets g_1N and g_2N the group operation is

$$g_1N \cdot g_2N \stackrel{\text{def}}{=} (g_1g_2)N.$$

Question: where do we actually use the fact that N is normal? We don't talk about ϕ or Q anywhere in this definition.

The answer is in Remark 3.3.3. The group operation takes in two cosets, so it doesn't know what g_1 and g_2 are. But behind the scenes, **the normal condition guarantees that the group operation can pick any g_1 and g_2 it wants and still end up with the same coset.** If we didn't have this property, then it would be hard to define

the product of two cosets C_1 and C_2 because it might make a difference which $g_1 \in C_1$ and $g_2 \in C_2$ we picked. The fact that N came from a homomorphism meant we could pick any representatives g_1 and g_2 of the cosets we wanted, because they all had the same ϕ -value.

We want some conditions which force this to be true without referencing ϕ at all. Suppose $\phi : G \rightarrow K$ is a homomorphism of groups with $H = \ker \phi$. Aside from the fact H is a group, we can get an “obvious” property:

Question 3.5.1. Show that if $h \in H$, $g \in G$, then $ghg^{-1} \in H$. (Check $\phi(ghg^{-1}) = 1_H$.)

Example 3.5.2 (Example of a non-normal subgroup)

Let $D_{12} = \langle r, s \mid r^6 = s^2 = 1, rs = sr^{-1} \rangle$. Consider the subgroup of order two $H = \{1, s\}$ and notice that

$$rsr^{-1} = r(sr^{-1}) = r(rs) = r^2s \notin H.$$

Hence H is not normal, and cannot be the kernel of any homomorphism.

Well, duh – so what? Amazingly it turns out that that this is the *sufficient* condition we want. Specifically, it makes the nice “coset multiplication” we wanted work out.

Remark 3.5.3 (For math contest enthusiasts). This coincidence is really a lot like functional equations at the IMO. We all know that normal subgroups H satisfy $ghg^{-1} \in H$; the surprise is that from the latter seemingly weaker condition, we can deduce H is normal.

Thus we have a new criterion for “normal” subgroups which does not make any external references to ϕ .

Theorem 3.5.4 (Algebraic condition for normal subgroups)

Let H be a subgroup of G . Then the following are equivalent:

- $H \trianglelefteq G$.
- For every $g \in G$ and $h \in H$, $ghg^{-1} \in H$.

Proof. We already showed one direction.

For the other direction, we need to build a homomorphism with kernel H . So we simply *define* the group G/H as the cosets. To put a group operation, we need to verify:

Claim 3.5.5. If $g'_1 \sim_H g_1$ and $g'_2 \sim_H g_2$ then $g'_1 g'_2 \sim_H g_1 g_2$.

Proof. Boring algebraic manipulation (again functional equation style). Let $g'_1 = g_1 h_1$ and $g'_2 = g_2 h_2$, so we want to show that $g_1 h_1 g_2 h_2 \sim_H g_1 g_2$. Since H has the property, $g_2^{-1} h_1 g_2$ is some element of H , say h_3 . Thus $h_1 g_2 = g_2 h_3$, and the left-hand side becomes $g_1 g_2 (h_3 h_2)$, which is fine since $h_3 h_2 \in H$. ■

With that settled we can just *define* the product of two cosets (of normal subgroups) by

$$(g_1 H) \cdot (g_2 H) = (g_1 g_2) H.$$

Thus the claim above shows that this multiplication is well-defined (this verification is the “content” of the theorem). So G/H is indeed a group! Moreover there is an obvious “projection” homomorphism $G \rightarrow G/H$ (with kernel H), by $g \mapsto gH$. $\square$

Example 3.5.6 (Modding out in the product group)

Consider again the product group $G \times H$. Earlier we identified a subgroup

$$G' = \{(g, 1_H) \mid g \in G\} \cong G.$$

You can easily see that $G' \trianglelefteq G \times H$. (Easy calculation.)

Moreover, just as the notation would imply, you can check that

$$(G \times H)/(G') \cong H.$$

Indeed, we have $(g, h) \sim_{G'} (1_G, h)$ for all $g \in G$ and $h \in H$.

Example 3.5.7 (Another explicit computation)

Let $\phi : D_8 \rightarrow \mathbb{Z}_4$ be defined by

$$r \mapsto \bar{2}, \quad s \mapsto \bar{2}.$$

The kernel of this map is $N = \{1, r^2, sr, sr^3\}$.

We can do a quick computation of all the elements of D_8 to get

$$\phi(1) = \phi(r^2) = \phi(sr) = \phi(sr^3) = \bar{0} \text{ and } \phi(r) = \phi(r^3) = \phi(s) = \phi(sr^2) = \bar{2}.$$

The two relevant fibers are

$$\phi^{\text{pre}}(\bar{0}) = 1N = r^2N = srN = sr^3N = \{1, r^2, sr, sr^3\}$$

and

$$\phi^{\text{pre}}(\bar{2}) = rN = r^3N = sN = sr^2N = \{r, r^3, s, sr^2\}.$$

So we see that $|D_8/N| = 2$ is a group of order two, or $\mathbb{Z}_2$. Indeed, the image of ϕ is

$$\{\bar{0}, \bar{2}\} \cong \mathbb{Z}_2.$$

Question 3.5.8. Suppose G is abelian. Why does it follow that any subgroup of G is normal?

Finally here’s some food for thought: suppose one has a group presentation for a group G that uses n generators. Can you write it as a quotient of the form F_n/N , where N is a normal subgroup of F_n ?

§3.6 (Digression) The first isomorphism theorem

One quick word about what other sources usually say.

Most textbooks actually *define* normal using the $ghg^{-1} \in H$ property. Then they define G/H for normal H in the way I did above, using the coset definition

$$(g_1H) \cdot (g_2H) = g_1g_2H.$$

Using purely algebraic manipulations (like I did) this is well-defined, and so now you have this group G/H or something. The underlying homomorphism isn't mentioned at all, or is just mentioned in passing.

I think this is incredibly dumb. The normal condition looks like it gets pulled out of thin air and no one has any clue what's going on, because no one has any clue what a normal subgroup actually should look like.

Other sources like to also write the so-called first isomorphism theorem.² It goes like this.

Theorem 3.6.1 (First isomorphism theorem)

Let $\phi : G \rightarrow H$ be a homomorphism. Then $G/\ker \phi$ is isomorphic to $\phi(G)$.

To me, this is just a clumsier way of stating the same idea.

About the only merit this claim has is that if ϕ is injective, then the image $\phi(G)$ is an *isomorphic copy* of G inside the group H . (Try to see this directly!) This is a pattern we'll often see in other branches of mathematics: whenever we have an *injective structure-preserving map*, often the image of this map will be some "copy" of G . (Here "structure" refers to the group multiplication, but we'll see some more other examples of "types of objects" later!)

In that sense an injective homomorphism $\phi : G \hookrightarrow H$ is an *embedding* of G into H .

§3.7 Problems to think about

Problem 3A (18.701 at MIT). Determine all groups G for which the map $\phi : G \rightarrow G$ defined by

$$\phi(g) = g^2$$

is a homomorphism.

Problem 3B. Let G and H be finite groups, where $|G| = 1000$ and $|H| = 999$. Show that a homomorphism $G \rightarrow H$ must be trivial.

Problem 3C. Let $\mathbb{C}^\times$ denote the nonzero complex numbers under multiplication. Show that there are five homomorphisms $\mathbb{Z}_5 \rightarrow \mathbb{C}^\times$ but only two homomorphisms $D_{10} \rightarrow \mathbb{C}^\times$, even though $\mathbb{Z}_5$ is a subgroup of D_{10} .



Problem 3D. Let p be a prime and $F_1 = F_2 = 1$, $F_{n+2} = F_{n+1} + F_n$ be the Fibonacci sequence. Show that $F_{2p(p^2-1)}$ is divisible by p .



Problem 3E. Find a non-abelian group G such that every subgroup of G is normal. (These groups are called **Hamiltonian**.)



Problem 3F (Homophony group). The homophony group (of English) is the group with 26 generators $a, b, \dots, z$ and one relation for every pair of English words which sound the same. For example $knight = night$ (and hence $k = 1$). Prove that the group is trivial.

² There is a second and third isomorphism theorem. But four years after learning about them, I *still* don't know what they are. So I'm guessing they weren't very important.

4 Topological notions

Here I'm just talking about some various properties of general topological spaces. The chapter is admittedly a bit disconnected (pun not intended, but hey, why not).

I should say a little about metric spaces vs general topological spaces here. In most of the sections we only really talk about continuous maps or open sets, so it doesn't make a difference whether we restrict our attention to metric spaces or not. But when we talk about completeness, we consider *sequences* of points, and in particular the distances between them. This notion doesn't make sense for general spaces, so for that section one must keep in mind we are only working in metric spaces.

The most important topological notion is missing from this chapter: that of a *compact* space. It is so important that I have dedicated a separate chapter just for it.

Note that in contrast to the warning on open/closed sets I gave earlier,

The adjectives in this chapter will be used to describe *spaces*.

§4.1 Connected spaces

It is possible for a set to be both open and closed (or **clopen**) in a topological space X ; for example $\emptyset$ and the entire space are examples of clopen sets. In fact, the presence of a nontrivial clopen set other than these two leads to a so-called *disconnected* space.

Question 4.1.1. Show that a space X has a nontrivial clopen set (one other than $\emptyset$ and X) if and only if X can be written as a disjoint union $U \sqcup V$ where U and V are both open and nonempty. (Use the definition that X is the complement of an open set.)

We say X is **disconnected** if there are nontrivial clopen sets, and **connected** otherwise. To see why this should be a reasonable definition, it might help to solve Problem 4A[†].

Example 4.1.2 (Disconnected and connected spaces)

(a) The metric space

$$\{(x, y) \mid x^2 + y^2 \leq 1\} \cup \{(x, y) \mid (x - 4)^2 + y^2 \leq 1\} \subseteq \mathbb{R}^2$$

is disconnected (it consists of two disks).

(b) A discrete space on more than one point is disconnected, since *every* set is clopen in the discrete space.

(c) Convince yourself that the set

$$\{x \in \mathbb{Q} : x^2 < 2014\}$$

is a clopen subset of $\mathbb{Q}$. Hence $\mathbb{Q}$ is disconnected too – it has *gaps*.

(d) $[0, 1]$ is connected.

§4.2 Path-connected spaces

Prototypical example for this section: Walking around in $\mathbb{C}$.


A stronger and perhaps more intuitive notion of a connected space is a *path-connected* space. The short description: “walk around in the space”.

Definition 4.2.1. A **path** in the space X is a continuous function

$$\gamma : [0, 1] \rightarrow X.$$

Its **endpoints** are the two points $\gamma(0)$ and $\gamma(1)$.

You can think of $[0, 1]$ as measuring “time”, and so we’ll often write $\gamma(t)$ for $t \in [0, 1]$ (with t standing for “time”). Here’s a picture of a path.



Question 4.2.2. Why does this agree with your intuitive notion of what a “path” and “loop” is?

Definition 4.2.3. A space X is **path-connected** if any two points in it are connected by some path.

Exercise 4.2.4. Let $X = U \sqcup V$ be a disconnected space. Show that there is no path from a point of U to point V . (If $\gamma : [0, 1] \rightarrow X$, then we get $[0, 1] = \gamma^{\text{pre}}(U) \sqcup \gamma^{\text{pre}}(V)$, but $[0, 1]$ is connected.)

In fact, there exist connected spaces which are not path-connected (one is called the *topologist’s sine curve*). This shows that path-connected is stronger than connected.


Example 4.2.5 (Examples of path-connected spaces)

- $\mathbb{R}^2$ is path-connected, since we can “connect” any two points with a straight line.
- The unit circle S^1 is path-connected, since we can just draw the major or minor arc to connect two points.


§4.3 Homotopy and simply connected spaces

Prototypical example for this section: $\mathbb{C}$ and $\mathbb{C} \setminus \{0\}$.

Now let’s motivate the idea of homotopy. Consider the example of the complex plane $\mathbb{C}$ (which you can think of just as $\mathbb{R}^2$) with two points p and q . There’s a whole bunch of paths from p to q but somehow they’re not very different from one another. If I told you “walk from p to q ” you wouldn’t have too many questions.



So we're living happily in $\mathbb{C}$ until a meteor strikes the origin, blowing it out of existence. Then suddenly to get from p to q , people might tell you two different things: "go left around the meteor" or "go right around the meteor".



So what's happening? In the first picture, the red, green, and blue paths somehow all looked the same: if you imagine them as pieces of elastic string pinned down at p and q , you can stretch each one to any other one.

But in the second picture, you can't move the red string to match with the blue string: there's a meteor in the way. The paths are actually different.¹

The formal notion we'll use to capture this is *homotopy equivalence*. We want to write a definition such that in the first picture, the three paths are all *homotopic*, but the two paths in the second picture are somehow not homotopic. And the idea is just continuous deformation.

Definition 4.3.1. Let α and β be paths in X whose endpoints coincide. A (path) **homotopy** from α to β is a continuous function $F : [0, 1]^2 \rightarrow X$, which we'll write $F_s(t)$ for $s, t \in [0, 1]$, such that

$$F_0(t) = \alpha(t) \text{ and } F_1(t) = \beta(t) \text{ for all } t \in [0, 1]$$

and moreover

$$\alpha(0) = \beta(0) = F_s(0) \text{ and } \alpha(1) = \beta(1) = F_s(1) \text{ for all } s \in [0, 1].$$


If a path homotopy exists, we say α and β are path **homotopic** and write $\alpha \simeq \beta$.

Abuse of Notation 4.3.2. While I strictly should say "path homotopy" to describe this relation between two paths, I will shorten this to just "homotopy" instead. Similarly I will shorten "path homotopic" to "homotopic".

¹If you know about winding numbers, you might feel this is familiar. We'll talk more about this in the chapter on homotopy groups.

Picture: <https://commons.wikimedia.org/wiki/File:HomotopySmall.gif>. Needless to say, $\simeq$ is an equivalence relation.

What this definition is doing is taking α and “continuously deforming” it to β , while keeping the endpoints fixed. Note that for each particular s , F_s is itself a function. So s represents time as we deform α to β : it goes from 0 to 1, starting at α and ending at β .



Question 4.3.3. Convince yourself the above definition is right. What goes wrong when the meteor strikes?

So now I can tell you what makes $\mathbb{C}$ special:

Definition 4.3.4. A space X is **simply connected** if it’s path-connected and for any points p and q , all paths from p to q are homotopic.

That’s why you don’t ask questions when walking from p to q in $\mathbb{C}$: there’s really only one way to walk. Hence the term “simply” connected.

Question 4.3.5. Convince yourself that $\mathbb{R}^n$ is simply connected for all n .

§4.4 Bases of spaces

Prototypical example for this section: $\mathbb{R}$ has a basis of open intervals, and $\mathbb{R}^2$ has a basis of open disks.

You might have noticed that the open sets of $\mathbb{R}$ are a little annoying to describe: the prototypical example of an open set is $(0, 1)$, but there are other open sets like

$$(0, 1) \cup \left(1, \frac{3}{2}\right) \cup \left(2, \frac{7}{3}\right) \cup (2014, 2015).$$

Question 4.4.1. Check this is an open set.

But okay, this isn’t *that* different. All I’ve done is taken a bunch of my prototypes and threw a bunch of $\cup$ signs at it. And that’s the idea behind a basis.

Definition 4.4.2. A **basis** for a topological space X is a subset $\mathcal{B}$ of the open sets such that every open set in X is a union of some (possibly infinite) number of elements in $\mathcal{B}$.

And all we’re doing is saying

Example 4.4.3 (Basis of $\mathbb{R}$)

The open intervals form a basis of $\mathbb{R}$.

In fact, more generally we have:

Theorem 4.4.4 (Basis of metric spaces)

The r -neighborhoods form a basis of any metric space M .

Proof. Kind of silly – given an open set U draw a r_p -neighborhood U_p contained entirely inside U . Then $\bigcup_p U_p$ is contained in U and covers every point inside it. $\square$

Hence, an open set in $\mathbb{R}^2$ is nothing more than a union of a bunch of open disks, and so on. The point is that in a metric space, the only open sets you really ever have to worry too much about are the r -neighborhoods.

§4.5 Completeness

Prototypical example for this section: $\mathbb{R}$ is complete, but $\mathbb{Q}$ and $(0, 1)$ are not.

Completeness is a property of *metric spaces*.

So far we can only talk about sequences converging if they have a limit. But consider the sequence $x_1 = 1, x_2 = 1.4, x_3 = 1.41, x_4 = 1.414, \dots$. It converges to $\sqrt{2}$ in $\mathbb{R}$, of course.

But it fails to converge in $\mathbb{Q}$. And so somehow, if we didn't know about the existence of $\mathbb{R}$, we would have *no idea* that the sequence (x_n) is “approaching” something.

That seems to be a shame. Let's set up a new definition to describe these sequences whose terms get close to each other, even if they don't approach any particular point in the space. Thus, we only want to mention the given points in the definition. This is hard to do in a general topological space, but such a notion does exist for metric spaces.

Definition 4.5.1. Let $x_1, x_2, \dots$ be a sequence which lives in a metric space $M = (M, d_M)$. We say the sequence is **Cauchy** if for any $\varepsilon > 0$, we have

$$d_M(x_m, x_n) < \varepsilon$$

for all sufficiently large m and n .

Question 4.5.2. Show that a sequence which converges is automatically Cauchy. (Draw a picture.)

Now we can define:

Definition 4.5.3. A metric space M is **complete** if every Cauchy sequence converges.

Example 4.5.4 (Examples of complete spaces)

- (a) $\mathbb{R}$ is complete. (Depending on your definition of $\mathbb{R}$, this either follows by definition, or requires some work. We won't go through this here.)
- (b) The discrete space is complete, as the only Cauchy sequences are eventually constant.
- (c) The closed interval $[0, 1]$ is complete.
- (d) $\mathbb{R}^n$ is complete as well. (You're welcome to prove this by induction on n .)

Example 4.5.5 (Non-Examples of complete spaces)

- (a) The rationals $\mathbb{Q}$ are not complete.
- (b) The open interval $(0, 1)$ is not complete, as the sequence $x_n = \frac{1}{n}$ is Cauchy but does not converge.

So, metric spaces need not be complete, like $\mathbb{Q}$. But it turns out that every metric space can be *completed* by “filling in the gaps”, resulting in a space called the **completion** of the metric space. For example, the completion of $\mathbb{Q}$ is $\mathbb{R}$ (in fact, this is often taken as the definition of $\mathbb{R}$).

§4.6 Subspacing

Earlier in this chapter, I declared that all the adjectives we defined were used to describe spaces. However, I now want to be more flexible and describe how they can be used to define subsets of spaces as well.

We've talked about some properties that a space can have; say, a space is *path-connected*, or *simply connected*, or *complete*. If you think about it, it would make sense to use the same adjectives on sets. You'd say a set S is path-connected if there's a path between every pair of points in S through points in S . And so on.

The reason you can do this is that for a metric space M ,

Every subset $S \subseteq M$ is a metric space in its own right .

To see this, we just use the same distance function; this is called the **subspace topology**. For example, you can think of a circle S^1 as a connected set by viewing it as a subspace of $\mathbb{R}^2$. Thus,

Abuse of Notation 4.6.1. Any adjective used to describe a space can equally be used for a subset of a space, and I'll thus be letting these mix pretty promiscuously. So in what follows, I might refer to subsets being complete, even when I only defined completeness for spaces. This is what I mean.

So to be perfectly clear:


- The statement “[$0, 1$] is complete” makes sense (and is true); it says that $[0, 1]$ is a complete metric space.
- The statement “[$0, 1$] is a complete subset of $\mathbb{R}$ ” is valid; it says that the subspace $[0, 1]$ of $\mathbb{R}$ is a complete metric space, which is of course true.

- The statement “ $[0, 1]$ is a closed subset of $\mathbb{R}$ ” makes sense; it says that the set of points $[0, 1]$ form a closed subset of a parent space $\mathbb{R}$.
- The statement “ $[0, 1]$ is closed” does *not* make sense. Closed sets are only defined relative to parent spaces.

To make sure you understand this:

Question 4.6.2. Let M be a complete metric space and let $S \subseteq M$. Prove that S is complete if and only if it is closed in M . (This is obvious once you figure out what the question is asking.) In particular, $[0, 1]$ is complete.

One can also define the subspace topology in a general topological space X . Given a subset $S \subseteq X$, the open sets of the subspace S are those of the form $S \cap U$, where U is an open set of X . So for example, if we view S^1 as a subspace of $\mathbb{R}^2$, then any open arc is an open set, because you can view it as the intersection of an open disk with S^1 .



Needless to say, for metric spaces it doesn't matter which of these definitions I choose. (Proving this turns out to be surprisingly annoying, so I won't do so.)

§4.7 Hausdorff spaces

Prototypical example for this section: Every space that's not the Zariski topology (defined much later).

In analysis, almost all topological spaces we encounter satisfy some additional hypotheses about their points. There is a whole hierarchy of such axioms, labelled T_n for integers n (with $n = 0$ being the weakest and $n = 6$ the strongest). These axioms are called **separation axioms**.

By far the most common hypothesis is the T_2 axiom, which bears a special name.

Definition 4.7.1. A topological space X is **Hausdorff** if for any two disjoint points p and q in X , there exists a neighborhood U of p and a neighborhood V of q such that

$$U \cap V = \emptyset.$$

In other words, around any two distinct points we should be able to draw disjoint neighborhoods.

Question 4.7.2. Important special case: all metric spaces are Hausdorff.

I just want to define this here so that I can use this word later. In any case, basically any space we will encounter other than the Zariski topology is Hausdorff.

§4.8 Problems to think about

Problem 4A[†]. Let X be a topological space. Show that there exists a nonconstant continuous function $X \rightarrow \{0, 1\}$ if and only if X is disconnected (here $\{0, 1\}$ is given the discrete topology).

Problem 4B^{*}. Let $f : X \rightarrow Y$ be a continuous function.

- (a) Show that if X is connected then so is $f(X)$.
- (b) Show that if X is path-connected then so is $f(X)$.

Problem 4C[†] (Banach fixed point theorem). Let $M = (M, d)$ be a complete metric space. Suppose $T : M \rightarrow M$ is a continuous map such that for any $p, q \in M$,

$$d(T(p), T(q)) < 0.999d(p, q).$$

(We call T a **contraction**.) Show that T has a unique fixed point.



Problem 4D. We know that any open set $U \subseteq \mathbb{R}$ is a union of open intervals (allowing $\pm\infty$ as endpoints). One can show that it's actually possible to write U as the union of *pairwise disjoint* open intervals.² Prove that there exists such a disjoint union with at most *countably many* intervals in it.



Problem 4E. From the plane $\mathbb{R}^2$ we delete two distinct points p and q . Is it possible to partition the remaining points into non-overlapping circumferences of nonzero radius?



Problem 4F[†] (Completion). Let M be a metric space. Show that we can add some points to M to obtain a metric space $\bar{M}$ which is complete, in such a way that every open set of $\bar{M}$ contains a point of M (meaning M is **dense** in $\bar{M}$).

²You are invited to try and prove this, but I personally found the proof quite boring.

5 Compactness

One of the most important notions of topological spaces is that of *compactness*. It generalizes the notion of “closed and bounded” in Euclidean space to any topological space (e.g. see Problem 5D[†]).

For metric spaces, there are two equivalent ways of formulating compactness:

- A “natural” definition using *sequences*, called sequential compactness.
- A less natural definition using open covers.

As I alluded to earlier, sequences in metric spaces are super nice, but sequences in general topological spaces *suck* (to the point where I didn’t bother to define convergence of general sequences). So it’s the second definition that will be used for general spaces.

§5.1 Definition of sequential compactness

Prototypical example for this section: $[0, 1]$ is compact, but $(0, 1)$ is not.

To emphasize, compactness is one of the *best* possible properties that a metric space can have.

Definition 5.1.1. A **subsequence** of an infinite sequence $x_1, x_2, \dots$ is exactly what it sounds like: a sequence $x_{i_1}, x_{i_2}, \dots$ where $i_1 < i_2 < \dots$ are positive integers. Note that the sequence is required to be infinite.

Another way to think about this is “selecting infinitely many terms” or “deleting some terms” of the sequence, depending on whether your glass is half empty or half full.

Definition 5.1.2. A metric space M is **sequentially compact** if every sequence has a subsequence which converges.

This time, let me give some non-examples before the examples.

Example 5.1.3 (Non-Examples of compact metric spaces)

- (a) The space $\mathbb{R}$ is not compact: consider the sequence $1, 2, 3, 4, \dots$. Any subsequence explodes, hence $\mathbb{R}$ cannot possibly be compact.
- (b) More generally, if a space is not **bounded** it cannot be compact. By bounded I mean that there exists a constant L such that $d_M(x, y) < L$ for all $x, y \in M$.
- (c) The open interval $(0, 1)$ is bounded but not compact: consider the sequence $\frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \dots$. No subsequence can converge to a point in $(0, 1)$ because the sequence “converges to 0”.
- (d) More generally, any space which is not complete cannot be compact.

Now for the examples!

Question 5.1.4. Show that a finite set is compact. (Pigeonhole Principle.)

Example 5.1.5 (Examples of compact spaces)

Here are some more examples of compact spaces. I'll prove they're compact in just a moment; for now just convince yourself they are.

- (a) $[0, 1]$ is compact. Convince yourself of this! Imagine having a large number of dots in the unit interval. . .
- (b) The surface of a sphere, $S^2 = \{(x, y, z) \mid x^2 + y^2 + z^2 = 1\}$ is compact.
- (c) The unit ball $B^2 = \{(x, y) \mid x^2 + y^2 = 1\}$ is compact.
- (d) The **Hawaiian earring** living in $\mathbb{R}^2$ is compact: it consists of mutually tangent circles of radius $\frac{1}{n}$ for each n , as in Figure 5.1.




Figure 5.1: Hawaiian Earring.

To aid in generating more examples, we remark:

Proposition 5.1.6 (Closed subsets of compacts)

Closed subsets of sequentially compact sets are compact.

Question 5.1.7. Prove this. (It should follow easily from definitions.)

We need to do a bit more work for these examples, which we do in the next section.

§5.2 Criteria for compactness

Theorem 5.2.1 (Tychonoff's theorem)

If X and Y are compact spaces, then so is $X \times Y$.

Proof. Problem 5C. □

We also have:

Theorem 5.2.2 (The interval is compact)

$[0, 1]$ is compact.

Proof. Killed by Problem 5D[†]; however, here is a sketch of a direct proof. Split $[0, 1]$ into $[0, \frac{1}{2}] \cup [\frac{1}{2}, 1]$. By Pigeonhole, infinitely many terms of the sequence lie in the left half (say); let x_1 be the first one and then keep only the terms in the left half after x_1 . Now split $[0, \frac{1}{2}]$ into $[0, \frac{1}{4}] \cup [\frac{1}{4}, \frac{1}{2}]$. Again, by Pigeonhole, infinitely many terms fall in some half; pick one of them, call it x_2 . Rinse and repeat. In this way we generate a sequence $x_1, x_2, \dots$ which is Cauchy, implying that it converges since $[0, 1]$ is complete. $\square$

Now we can prove the main theorem about Euclidean space: in $\mathbb{R}^n$, compactness is equivalent to being “closed and bounded”.

Theorem 5.2.3 (Bolzano-Weierstraß)

A subset of $\mathbb{R}^n$ is compact if and only if it closed and bounded.

Question 5.2.4. Why does this imply the spaces in our examples are compact?

Proof. Well, look at a closed and bounded $S \subseteq \mathbb{R}^n$. Since it’s bounded, it lives inside some box $[a_1, b_1] \times [a_2, b_2] \times \dots \times [a_n, b_n]$. By Tychonoff’s theorem, since each $[a_i, b_i]$ is compact the entire box is. Since S is a closed subset of this compact box, we’re done. $\square$

One really has to work in $\mathbb{R}^n$ for this to be true! In other spaces, this criteria can easily fail.

Example 5.2.5 (Closed and bounded but not compact)

Let $S = \{s_1, s_2, \dots\}$ be any infinite set equipped with the discrete metric. Then S is closed (since all convergent sequences are constant sequences) and S is bounded (all points are a distance 1 from each other) but it’s certainly not compact since the sequence $s_1, s_2, \dots$ doesn’t converge.

The Heine-Borel theorem, which is Problem 5D[†], tells you exactly which sets are compact in metric spaces in a geometric way.

§5.3 Compactness using open covers

Prototypical example for this section: $[0, 1]$ is compact.

There’s a second related notion of compactness which I’ll now define. The following definitions might appear very unmotivated, but bear with me.

Definition 5.3.1. An open cover of a metric space M is a collection of open sets $\{U_\alpha\}$ (possibly infinite or uncountable) which *cover* it: every point in M lies in at least one of the U_α , so that

$$M = \bigcup U_\alpha.$$

Such a cover is called an **open cover**. A **subcover** is exactly what it sounds like: it takes only some of the U_α , while ensuring that M remains covered.


Definition 5.3.2. A metric space M is **compact** if *every* open cover has a finite subcover.

What does this mean? Here's an example:

Example 5.3.3 (Example of a finite subcover)

Suppose we cover the unit square $M = [0, 1]^2$ by putting an open disk of diameter 1 at every point (trimming any overflow). This is clearly an open cover because, well, every point lies in *many* of the open sets, and in particular is the center of one.

But this is way overkill – we only need about four of these circles to cover the whole square. That's what is meant by a “finite subcover”.



Why do we care? Because of this:

Theorem 5.3.4 (Sequentially compact $\iff$ compact)

A metric space M is sequentially compact if and only if it is compact.

We defer the proof to the last section.

Example 5.3.5 (An example of non-compactness)

The space $X = [0, 1)$ is not compact in either sense. We can already see it is not sequentially compact, because it is not even complete (look at $x_n = 1 - \frac{1}{n}$). To see it is not compact under the covering definition, consider the sets

$$U_m = \left[0, 1 - \frac{1}{m+1}\right)$$

for $m = 1, 2, \dots$. Then $X = \bigcup U_i$; hence the U_i are indeed a cover. But no finite collection of the U_i 's will cover X .

Question 5.3.6. Convince yourself that $[0, 1]$ is compact; this is a little less intuitive than it being sequentially compact.

Abuse of Notation 5.3.7. Thus, we'll never call a metric space “sequentially compact” again – we'll just say “compact”. (Indeed, I kind of already did this in the previous few sections.)

Open covers also have the nice property that, again, they don't refer to metrics or sequences. So rather than port the sequential definition, we port this definition over for a general space:

Definition 5.3.8. A topological space X is **compact** if every open cover has finite subcover.

§5.4 Applications of compactness

Compactness lets us reduce *infinite* open covers to finite ones. Actually, it lets us do this even if the open covers are *blithely stupid*. Very often one takes an open cover consisting of a neighborhood of $x \in X$ for every single point x in the space; this is a huge number of open sets, and yet compactness lets us reduce to a finite set.

To give an example of a typical usage:

Proposition 5.4.1 (Compact $\implies$ totally bounded)

Let M be compact. Then M is **totally bounded**, meaning that for any $\varepsilon > 0$ we can cover M with finitely many ε -neighborhoods.

(To picture this: in $\mathbb{R}^2$, any compact set can be covered by unit disks.)

Proof Using Covers. For every point $p \in M$, take an ε -neighborhood of p , say U_p . These cover M for the horrendously stupid reason that each point is at the very least covered by the neighborhood U_p . Compactness then lets us take a finite subcover. $\square$

Next, an important result about maps between compact spaces.

Theorem 5.4.2 (Images of compacts are compact)

Let $f : X \rightarrow Y$ be a continuous function, where X is compact. Then the image

$$f(X) \subseteq Y$$

is compact.

Proof Using Covers. Take any open cover $\{V_\alpha\}$ in Y of $f(X)$. By continuity of f , it pulls back to an open cover $\{U_\alpha\}$ of X . Thus some finite subcover of this covers X . The corresponding V 's cover $f(X)$. $\square$

Question 5.4.3. Solve this using the sequential definition of continuity too. (This is even easier.)

Some nice corollaries of this:

Corollary 5.4.4 (Extreme value theorem)

Let X be compact and consider a continuous function $f : X \rightarrow \mathbb{R}$. Then f achieves a *maximum value* at some point, i.e. there is a point $p \in X$ such that $f(p) \geq f(q)$ for any other $q \in X$.

Corollary 5.4.5 (Intermediate value theorem)

Consider a continuous function $f : [0, 1] \rightarrow \mathbb{R}$. Then the image of f is of the form $[a, b]$ for some real numbers $a \leq b$.

The typical counterexample given for $(0, 1) \rightarrow \mathbb{R}$ is the function $\frac{1}{x}$.

Sketch of Proof. The point is that the image of f is compact in $\mathbb{R}$, and hence closed and bounded. You can convince yourself that the closed sets are just unions of closed intervals. That implies the extreme value theorem.

When $X = [0, 1]$, the image is also connected, so there should only be one closed interval in $f([0, 1])$. Since the image is bounded, we then know it's of the form $[a, b]$. (To give a full proof, you would use the so-called *least upper bound* property, but that's a little involved for a bedtime story; also, I think $\mathbb{R}$ is boring.) $\square$

One last application: if M is a compact metric space, then continuous functions $f : M \rightarrow N$ are continuous in an especially “nice” way:

Definition 5.4.6. A function $f : M \rightarrow N$ of metric spaces is called **uniformly continuous** if for any $\varepsilon > 0$, there exists a $\delta > 0$ (depending only on ε) such that whenever $d_M(x, y) < \delta$ we also have $d_N(fx, fy) < \varepsilon$.

The name means that for $\varepsilon > 0$, we need a δ that works for *every point* of M .

Example 5.4.7 (Uniform continuity)

- (a) The functions $\mathbb{R} \rightarrow \mathbb{R}$ of the form $x \mapsto ax + b$ are all uniformly continuous, since one can always take $\delta = \varepsilon/|a|$ (or $\delta = 1$ if $a = 0$).
- (b) Actually, it is true that a differentiable function $\mathbb{R} \rightarrow \mathbb{R}$ with a bounded derivative is uniformly continuous. (The converse is false for the reason that uniformly continuous doesn't imply differentiable at all.)
- (c) The function $f : \mathbb{R} \rightarrow \mathbb{R}$ by $x \mapsto x^2$ is *not* uniformly continuous, since for large x , tiny δ changes to x lead to fairly large changes in x^2 . (If you like, you can try to prove this formally now.)
Think $f(2017.01) - f(2017) > 40$; even when $\delta = 0.01$, one can still cause large changes in f .
- (d) However, when restricted to $(0, 1)$ or $[0, 1]$ the function $x \mapsto x^2$ becomes uniformly continuous. (For $\varepsilon > 0$ one can now pick for example $\delta = \min(1, \varepsilon)/3$.)
- (e) The function $(0, 1) \rightarrow \mathbb{R}$ by $x \mapsto 1/x$ is *not* uniformly continuous (same reason as before).

Now, as promised:

Proposition 5.4.8 (Continuous on compact $\implies$ uniformly continuous)

If M is compact and $f : M \rightarrow N$ is continuous, then f is uniformly continuous.

Proof Using Sequences. Fix $\varepsilon > 0$, and assume for contradiction that for every $\delta = 1/k$ there exists points x_k and y_k within δ of each other but with images $\varepsilon > 0$ apart. By compactness, take a convergent subsequence $x_{i_k} \rightarrow p$. Then $y_{i_k} \rightarrow p$ as well, since the x_k 's and y_k 's are close to each other. So both sequences $f(x_{i_k})$ and $f(y_{i_k})$ should converge to $f(p)$ by sequential continuity, but this can't be true since the two sequences are always ε apart. $\square$

§5.5 (Optional) Equivalence of formulations of compactness

We'll prove that for a metric space M , the following are equivalent:

- (i) Every sequence has a convergent subsequence,
- (ii) The space M is complete and totally bounded, and
- (iii) Every open cover has a finite subcover.

We leave the proof that (i) $\iff$ (ii) as Problem 5D[†]; the idea of the proof is much in the spirit of Theorem 5.2.2.

Proof that (i) and (ii) $\implies$ (iii). We prove the following lemma, which is interesting in its own right.

Lemma 5.5.1 (Lebesgue number lemma)

Let M be a compact metric space and $\{U_\alpha\}$ an open cover. Then there exists a real number $\delta > 0$, called a **Lebesgue number** for that covering, such that the δ -neighborhood of any point p lies entirely in some U_α .

Proof of lemma. Assume for contradiction that for every $\delta = 1/k$ there is a point $x_k \in M$ such that its $1/k$ -neighborhood isn't contained in any U_α . In this way we construct a sequence $x_1, x_2, \dots$; thus we're allowed to take a subsequence which converges to some x . Then for every $\varepsilon > 0$ we can find an integer n such that $d(x_n, x) + 1/n < \varepsilon$; thus the ε -neighborhood at x isn't contained in any U_α for every $\varepsilon > 0$. This is impossible, because we assumed x was covered by some open set. $\blacksquare$

Now, take a Lebesgue number δ for the covering. Since M is totally bounded, finitely many δ -neighborhoods cover the space, so finitely many U_α do as well. $\square$

Proof that (iii) $\implies$ (ii). One step is immediate:

Question 5.5.2. Show that the covering condition $\implies$ totally bounded.

The tricky part is showing M is complete. Assume for contradiction it isn't and thus that the sequence (x_k) is Cauchy, but it doesn't converge to any particular point.

Question 5.5.3. Show that this implies for each $p \in M$, there is an ε_p -neighborhood U_p which contains at most finitely many of the points of the sequence (x_k) . (You will have to use the fact that $x_k \not\rightarrow p$ and (x_k) is Cauchy.)

Now if we consider $M = \bigcup_p U_p$ we get a finite subcover of these neighborhoods; but this finite subcover can only cover finitely many points of the sequence, by contradiction. $\square$

§5.6 Problems to think about

The later problems are pretty hard; some have the flavor of IMO 3/6-style constructions. It's important to draw lots of pictures so one can tell what's happening. Of these Problem 5D[†] is definitely my favorite.

Problem 5A. Show that $[0, 1]$ and $(0, 1)$ are not homeomorphic.

Problem 5B (Cantor's intersection theorem). Let X be a compact space, and suppose

$$X = K_0 \supseteq K_1 \supseteq K_2 \supseteq \dots$$

is an infinite sequence of nested nonempty closed subsets. Show that $\bigcap_{n \geq 0} K_n \neq \emptyset$.

Problem 5C (Tychonoff's theorem). Let X and Y be compact metric spaces. Show that $X \times Y$ is compact. (This is also true for general topological spaces, but the proof is surprisingly hard.)



Problem 5D[†] (Heine-Borel theorem). Prove that a metric space M is (sequentially) compact if and only if it is complete and totally bounded.



Problem 5E (Almost Arzelà-Ascoli theorem). Let $f_1, f_2, \dots : [0, 1] \rightarrow [-100, 100]$ be an **equicontinuous** sequence of functions, meaning

$$\forall \varepsilon > 0 \quad \exists \delta > 0 \quad \forall n \quad \forall x, y \quad (|x - y| < \delta \implies |f_n(x) - f_n(y)| < \varepsilon)$$

Show that we can extract a subsequence $f_{i_1}, f_{i_2}, \dots$ of these functions such that for every $x \in [0, 1]$, the sequence $f_{i_1}(x), f_{i_2}(x), \dots$ converges.

II

Linear Algebra

6	What is a vector space?	78
6.1	The definitions of a ring and field	78
6.2	Modules and vector spaces	78
6.3	Direct sums	80
6.4	Linear independence, spans, and basis	82
6.5	Linear maps	84
6.6	What is a matrix?	85
6.7	Subspaces and picking convenient bases	86
6.8	A cute application: Lagrange interpolation	88
6.9	(Digression) Arrays of numbers are evil	88
6.10	A word on general modules	90
6.11	Problems to think about	90
7	Trace and determinant	92
7.1	Tensor product	92
7.2	Dual module	94
7.3	The trace	95
7.4	(Optional) Two more digressions on dual spaces	97
7.5	Wedge product	98
7.6	The determinant	100
7.7	Problems to think about	101
8	Spectral theory	103
8.1	Why you should care	103
8.2	Eigenvectors and eigenvalues	103
8.3	The Jordan form	104
8.4	Nilpotent maps	106
8.5	Reducing to the nilpotent case	107
8.6	(Optional) Proof of nilpotent Jordan	108
8.7	Characteristic polynomials, and Cayley-Hamilton	109
8.8	(Digression) Tensoring up	110
8.9	Problems to think about	111
9	Inner product spaces	112
9.1	Defining the norm	112
9.2	Norms	114
9.3	Orthogonality	115
9.4	Identifying with the dual space	117
9.5	The transpose of a matrix	117
9.6	Spectral theory of normal maps	118
9.7	Problems to think about	119

6 What is a vector space?

This is a pretty light chapter. The point of it is to just define what a vector space and a basis are. These are very intuitive concepts that you likely already know.

§6.1 The definitions of a ring and field

Prototypical example for this section: $\mathbb{Z}$, $\mathbb{R}$, or $\mathbb{C}$ are rings; the latter two are fields.

I'll very informally define a ring/field here; if you're interested in the actual definition, you can consult the first part of the chapter on ideals. For now you can just remember:

- A **commutative ring** is a structure with a *commutative* addition and multiplication, as well as subtraction, like $\mathbb{Z}$. It also has an additive identity 0 and multiplicative identity 1. (Yes, I'm considering only commutative rings.)
- If the multiplication is invertible like in $\mathbb{R}$ or $\mathbb{C}$, (meaning $\frac{1}{x}$ makes sense for any $x \neq 0$), then the ring is called a **field**.

In fact, if you replace “field” by “ $\mathbb{R}$ ” everywhere in what follows, you probably won't lose much. It's customary to use the letter R for rings, and k or K for fields.

Finally, in case you skipped the chapter on groups, I should also mention:

- An **additive abelian group** is a structure with a commutative addition, as well as subtraction, plus an additive identity 0. It doesn't have to have multiplication. A good example is $\mathbb{R}^3$ (with addition componentwise).

§6.2 Modules and vector spaces

Prototypical example for this section: Polynomials of degree at most n .

You intuitively know already that $\mathbb{R}^n$ is a “vector space”: its elements can be added together, and there's some scaling by real numbers. Let's develop this more generally.

Fix a commutative ring R . Then informally,

An R -module is any structure where you can add two elements and scale by elements of R .

You can think of the R -module as consisting of soldiers being commanded by the ring R . Moreover, a **vector space** is just a module whose commanding ring is actually a field. I'll give you the full definition in a moment, but first, examples...

Example 6.2.1 (Quadratic polynomials, aka my favorite example)

My favorite example of an $\mathbb{R}$ -vector space is the set of polynomials of degree at most two, namely

$$\{ax^2 + bx + c \mid a, b, c \in \mathbb{R}\}.$$

Indeed, you can add any two quadratics, and multiply by constants. You can't multiply two quadratics to get a quadratic, but that's irrelevant – in a vector space there need not be a notion of multiplying two vectors together.

In a sense we'll define later, this vector space has dimension 3 (as expected!).

Example 6.2.2 (All polynomials)

The set of *all* polynomials with real coefficients is an $\mathbb{R}$ -vector space, because you can *add any two polynomials* and *scale by constants*.

Example 6.2.3 (Euclidean space)

- (a) The complex numbers

$$\{a + bi \mid a, b \in \mathbb{R}\}$$

form a real vector space. As we'll see later, it has "dimension 2".

- (b) The real numbers $\mathbb{R}$ form a real vector space of dimension 1.

- (c) The set of 3D vectors

$$\{(x, y, z) \mid x, y, z \in \mathbb{R}\}$$

forms a real vector space, because you can add any two triples component-wise. Again, we'll later explain why it has "dimension 3".

Example 6.2.4 (More examples of vector spaces)

- (a) The set

$$\mathbb{Q}[\sqrt{2}] = \{a + b\sqrt{2}, a, b \in \mathbb{Q}\}$$

has a structure of a $\mathbb{Q}$ -vector space in the obvious fashion: one can add any two elements, and scale by rational numbers. (It is not a real vector space – why?)

- (b) The set

$$\{(x, y, z) \mid x + y + z = 0 \text{ and } x, y, z \in \mathbb{R}\}$$

is a 2-dimensional real vector space.

- (c) The set of all functions $f : \mathbb{R} \rightarrow \mathbb{R}$ is also a real vector space (since the notions $f + g$ and $c \cdot f$ both make sense for $c \in \mathbb{R}$).

Now let me write the actual rules for how this multiplication behaves.

Definition 6.2.5. Let R be a commutative ring. An R -**module** is an additive abelian group $M = (M, +)$ equipped with a left multiplication by elements of R . This multiplication must satisfy the following properties for every $r_1, r_2 \in R$ and $m \in M$:

- (i) $r_1 \cdot (r_2 m) = (r_1 r_2) \cdot m$.
- (ii) Multiplication is distributive, meaning

$$(r_1 + r_2) \cdot m = r_1 \cdot m + r_2 \cdot m \text{ and } r \cdot (m_1 + m_2) = r \cdot m_1 + r \cdot m_2.$$

- (iii) $1_R \cdot m = m$.
- (iv) $0_R \cdot m = 0_M$. (This is actually extraneous; one can deduce it from the first three.)

If R is a field we say M is an R -**vector space**; its elements are called **vectors** and the members of R are called **scalars**.

Abuse of Notation 6.2.6. In the above, we’re using the same symbol $+$ for the addition of M and the addition of R . Sorry about that, but it’s kind of hard to avoid, and the point of the axioms is that these additions should be related. I’ll try to remember to put $r \cdot m$ for the multiplication of the module and just $r_1 r_2$ for the multiplication of R .

Question 6.2.7. In Example 6.2.1, I was careful to say “degree at most 2” instead of “degree 2”. What’s the reason for this? In other words, why is

$$\{ax^2 + bx + c \mid a, b, c \in \mathbb{R}, a \neq 0\}$$

not an $\mathbb{R}$ -vector space?

A couple less intuitive but somewhat important examples...

Example 6.2.8 (Abelian groups are $\mathbb{Z}$ -modules)

(Skip this example if you’re not comfortable with groups.)

(a) The example of real polynomials

$$\{ax^2 + bx + c \mid a, b, c \in \mathbb{R}\}$$

is also a $\mathbb{Z}$ -module! Indeed, we can add any two such polynomials, and we can scale them by integers.

(b) The set of integers modulo 100, say $\mathbb{Z}/100\mathbb{Z}$, is a $\mathbb{Z}$ -module as well. Can you see how?

(c) In fact, *any* abelian group $G = (G, +)$ is a $\mathbb{Z}$ -module. The multiplication can be defined by

$$n \cdot g = \underbrace{g + \cdots + g}_{n \text{ times}} \quad (-n) \cdot g = n \cdot (-g)$$

for $n \geq 0$. (Here $-g$ is the additive inverse of g .)

Example 6.2.9 (Every ring is its own module)

(a) $\mathbb{R}$ can be thought of as an $\mathbb{R}$ -vector space over itself. Can you see why?

(b) By the same reasoning, we see that *any* commutative ring R can be thought of as an R -module over itself.

§6.3 Direct sums

Prototypical example for this section: $\{ax^2 + bx + c\} = \mathbb{R} \oplus x\mathbb{R} \oplus x^2\mathbb{R}$, and $\mathbb{R}^2$ is the sum of its axes.

Let’s return to Example 6.2.1, and consider

$$V = \{ax^2 + bx + c \mid a, b, c \in \mathbb{R}\}.$$

Even though I haven’t told you what a dimension is, you can probably see that this vector space “should have” dimension 3. We’ll get to that in a moment.

The other thing you may have noticed is that somehow the x^2 , x and 1 terms don't "talk to each other". They're totally unrelated. In other words, we can consider the three sets

$$\begin{aligned} x^2\mathbb{R} &\stackrel{\text{def}}{=} \{ax^2 \mid a \in \mathbb{R}\} \\ x\mathbb{R} &\stackrel{\text{def}}{=} \{bx \mid b \in \mathbb{R}\} \\ \mathbb{R} &\stackrel{\text{def}}{=} \{c \mid c \in \mathbb{R}\}. \end{aligned}$$

In an obvious way, each of these can be thought of as a "copy" of $\mathbb{R}$.

Then V quite literally consists of the "sums of these sets". Specifically, every element of V can be written *uniquely* as the sum of one element from each of these sets. This motivates us to write

$$V = x^2\mathbb{R} \oplus x\mathbb{R} \oplus \mathbb{R}.$$

The notion which captures this formally is the **direct sum**.

Definition 6.3.1. Let M be an R -module. Let M_1 and M_2 be subsets of M which are themselves R -modules. Then we write $M = M_1 \oplus M_2$ and say M is a **direct sum** of M_1 and M_2 if every element from M can be written uniquely as the sum of an element from M_1 and M_2 .

Example 6.3.2 (Euclidean plane)

Take the vector space $\mathbb{R}^2 = \{(x, y) \mid x \in \mathbb{R}, y \in \mathbb{R}\}$. We can consider it as a direct sum of its x -axis and y -axis:

$$X = \{(x, 0) \mid x \in \mathbb{R}\} \text{ and } Y = \{(0, y) \mid y \in \mathbb{R}\}.$$

Then $\mathbb{R}^2 = X \oplus Y$.

This gives us a "top-down" way to break down modules into some disconnected components.

By applying this idea in reverse, we can also construct new vector spaces as follows. In a very unfortunate accident, the two names and notations for technically distinct things are exactly the same.

Definition 6.3.3. Let M and N be R -modules. We define the **direct sum** $M \oplus N$ to be the R -module whose elements are pairs $(m, n) \in M \times N$. The operations are given by

$$(m_1, n_1) + (m_2, n_2) = (m_1 + m_2, n_1 + n_2).$$

and

$$r \cdot (m, n) = (r \cdot m, r \cdot n).$$

For example, while we technically wrote $\mathbb{R}^2 = X \oplus Y$, since each of X and Y is just a copy of $\mathbb{R}$, we may as well have written $\mathbb{R}^2 \cong \mathbb{R} \oplus \mathbb{R}$.

Abuse of Notation 6.3.4. The above illustrates an abuse of notation in the way we write a direct sum. The symbol $\oplus$ has two meanings.

- If V is a *given* space and W_1 and W_2 are subspaces, then $V = W_1 \oplus W_2$ means that " V *splits* as a direct sum $W_1 \oplus W_2$ " in the way we defined above.

- If W_1 and W_2 are two *unrelated* spaces, then $W_1 \oplus W_2$ is *defined* as the vector space whose *elements* are pairs $(w_1, w_2) \in W_1 \times W_2$.

You can see that these definitions “kind of” coincide.

In this way, you can see that V should be isomorphic to $\mathbb{R} \oplus \mathbb{R} \oplus \mathbb{R}$; we had $V = x^2\mathbb{R} \oplus x\mathbb{R} \oplus \mathbb{R}$, but the $1, x, x^2$ don’t really talk to each other and each of the summands is really just a copy of $\mathbb{R}$ at heart.

Definition 6.3.5. We can also define, for every positive integer n , the module

$$M^{\oplus n} \stackrel{\text{def}}{=} \underbrace{M \oplus M \oplus \cdots \oplus M}_{n \text{ times}}.$$

§6.4 Linear independence, spans, and basis

Prototypical example for this section: $\{1, x, x^2\}$ is a basis of $\{ax^2 + bx + c \mid a, b, c \in \mathbb{R}\}$.

The idea of a basis, the topic of this section, gives us another way to capture the notion that

$$V = \{ax^2 + bx + c \mid a, b, c \in \mathbb{R}\}$$

is just sums of copies of $\{1, x, x^2\}$. This section should be very very intuitive, if technical. If you can’t see why the theorems here “should” be true, you’re doing it wrong.

Let M be an R -module now. We define three very classical notions that you likely are already familiar with. If not, fall upon your notion of Euclidean space or V above.

Definition 6.4.1. A **linear combination** of some vectors $v_1, \dots, v_n$ is a sum of the form $r_1v_1 + \cdots + r_nv_n$, where $r_1, \dots, r_n \in R$. The linear combination is called **trivial** if $r_1 = r_2 = \cdots = r_n = 0_R$, and **nontrivial** otherwise.

Definition 6.4.2. Consider a finite set of vectors $v_1, \dots, v_n$ in a module M .

- It is called **linearly independent** if there is no nontrivial linear combination with value 0_M . (Observe that $0_M = 0 \cdot v_1 + 0 \cdot v_2 + \cdots + 0 \cdot v_n$ is always true – the assertion is that there is no other way to express 0_M in this form.)
- It is called a **generating set** if every $v \in M$ can be written as a linear combination of the $\{v_i\}$. If M is a vector space we say it is **spanning** instead.
- It is called a **basis** (plural **bases**) if every $v \in M$ can be written *uniquely* as a linear combination of the $\{v_i\}$.

The same definitions apply for an infinite set, with the proviso that all sums must be finite.

So by definition, $\{1, x, x^2\}$ is a basis for V . It’s not the only one: $\{2, x, x^2\}$ and $\{x + 4, x - 2, x^2 + x\}$ are other examples of bases, though not as natural. However, the set $S = \{3 + x^2, x + 1, 5 + 2x + x^2\}$ is not a basis; it fails for two reasons:

- Note that $0 = (3 + x^2) + 2(x + 1) - (5 + 2x + x^2)$. So the set S is not linearly independent.
- It’s not possible to write x^2 as a sum of elements of S . So S fails to be spanning.

With these new terms, we can just say a basis is a linearly independent and spanning set.

Example 6.4.3 (More example of bases)

- (a) Regard $\mathbb{Q}[\sqrt{2}]$ as a $\mathbb{Q}$ -vector space. Then $\{1, \sqrt{2}\}$ is a basis.
- (b) If V is the set of all real polynomials, there is an infinite basis $\{1, x, x^2, \dots\}$. The condition that we only use finitely many terms just says that the polynomials must have finite degree (which is good).
- (c) Let $V = \{(x, y, z) \mid x + y + z = 0 \text{ and } x, y, z \in \mathbb{R}\}$. Then we expect there to be a basis of size 2, but unlike previous examples there is no immediately “obvious” choice. Some working examples include:
- $(1, -1, 0)$ and $(1, 0, -1)$,
 - $(0, 1, -1)$ and $(1, 0, -1)$,
 - $(5, 3, -8)$ and $(2, -1, -1)$.

Question 6.4.4. Show that a set of vectors is a basis if and only if it is linearly independent and spanning. (Think about the polynomial example if you get stuck.)

Now we state a few results which assert that bases in vector spaces behave as nicely as possible.

Theorem 6.4.5 (Maximality and minimality of bases)

Let V be a vector space over some field k and take $e_1, \dots, e_n \in V$. The following are equivalent:

- (a) The e_i form a basis.
- (b) The e_i are spanning, but no proper subset is spanning.
- (c) The e_i are linearly independent, but adding any other element of V makes them not linearly independent.

Remark 6.4.6. If we replace V by a general module M over a commutative ring R , then (a) $\implies$ (b) and (a) $\implies$ (c) but not conversely.

Proof. Straightforward, do it yourself if you like. The key point to notice is that you need to divide by scalars for the converse direction, hence V is required to be a vector space instead of just a module for the implications (b) $\implies$ (a) and (c) $\implies$ (a). $\square$

Theorem 6.4.7 (Dimension theorem for vector spaces)

Any two bases of a vector space have the same size.

Proof. We prove something stronger: Let V be a vector space, and assume $v_1, \dots, v_n$ is a spanning set while $w_1, \dots, w_m$ is linearly independent. We claim that $n \geq m$.

Question 6.4.8. Show that this claim is enough to imply the theorem.

Let $A_0 = \{v_1, \dots, v_n\}$ be the spanning set. Throw in w_1 : by the spanning condition, so $w_1 = c_1 v_1 + \dots + c_n v_n$. There's some nonzero coefficient, say c_n . Thus

$$v_n = \frac{1}{c_n} w_1 - \frac{c_1}{c_n} v_1 - \frac{c_2}{c_n} v_2 - \dots$$

Thus $A_1 = \{v_1, \dots, v_{n-1}, w_1\}$ is spanning. Now do the same thing, throwing in w_2 , and deleting some element of the v_i as before to get A_2 ; the condition that the w_i are linearly independent ensures that some v_i coefficient must always not be zero. Since we can eventually get to A_m , we have $n \geq m$. $\square$

Remark 6.4.9. In fact, this is true for modules over any commutative ring. Interestingly, the proof for the general case proceeds by reducing to the case of a vector space.

The dimension theorem, true to its name, lets us define the **dimension** of a vector space as the size of any finite basis, if one exists. When it does exist we say V is **finite-dimensional**. So for example,

$$V = \{ax^2 + bx + c \mid a, b, c \in \mathbb{R}\}$$

has dimension three, because $\{1, x, x^2\}$ is a basis. That's not the only basis: we could as well have written

$$\{a(x^2 - 4x) + b(x + 2) + c \mid a, b, c \in \mathbb{R}\}$$

and gotten the exact same vector space. But the beauty of the theorem is that no matter how we try to contrive the generating set, we always will get exactly three elements. That's why it makes sense to say V has dimension three.

On the other hand, the set of all polynomials is *infinite-dimensional* (which should be intuitively clear).

A basis $e_1, \dots, e_n$ of V is really cool because it means that to specify $v \in V$, I just have to specify $a_1, \dots, a_n \in k$, and then let $v = a_1 e_1 + \dots + a_n e_n$. You can even think of v as just $(a_1, \dots, a_n)$. To put it another way, if V is a k -vector space we always have

$$V = e_1 k \oplus e_2 k \oplus \dots \oplus e_n k.$$

For the remainder of this chapter, **assume all vector spaces are finite-dimensional unless otherwise specified**. Infinite-dimensional vector spaces make my eyes glaze over.

§6.5 Linear maps

We've seen homomorphisms and continuous maps. Now we're about to see linear maps, the structure preserving maps between vector spaces. Can you guess the definition?

Definition 6.5.1. Let V and W be vector spaces over the same field k . A **linear map** is a map $T : V \rightarrow W$ such that:

- (i) We have $T(v_1 + v_2) = T(v_1) + T(v_2)$ for $v_1, v_2 \in V$ (in group language, T is a homomorphism $(V, +) \rightarrow (W, +)$).
- (ii) For any $a \in k$ and $v \in V$, $T(a \cdot v) = a \cdot T(v)$.

If this map is a bijection (equivalently, if it has an inverse), it is an **isomorphism**. We say V and W are **isomorphic** vector spaces and write $V \cong W$.

Example 6.5.2 (Examples of linear maps)

- (a) For any vector spaces V and W there is a trivial linear map sending everything to $0_W \in W$.
- (b) For any vector space V , there is the identity isomorphism $\text{id} : V \rightarrow V$.
- (c) The map $\mathbb{R}^3 \rightarrow \mathbb{R}$ by $(a, b, c) \mapsto 4a + 2b + c$ is a linear map.
- (d) Let V be the set of real polynomials of degree at most 2. The map $\mathbb{R}^3 \rightarrow V$ by $(a, b, c) \mapsto ax^2 + bx + c$ is an *isomorphism*.
- (e) Let W be the set of functions $\mathbb{R} \rightarrow \mathbb{R}$. The evaluation map $W \rightarrow \mathbb{R}$ by $f \mapsto f(0)$ is a linear map.
- (f) There is a map of $\mathbb{Q}$ -vector spaces $\mathbb{Q}[\sqrt{2}] \rightarrow \mathbb{Q}[\sqrt{2}]$ called “multiply by $\sqrt{2}$ ”; this map sends $a + b\sqrt{2} \mapsto 2b + a\sqrt{2}$. This map is an isomorphism, because it has an inverse “multiply by $1/\sqrt{2}$ ”.

In the expression $T(a \cdot v) = a \cdot T(v)$, note that the first $\cdot$ is the multiplication of V and the second $\cdot$ is the multiplication of W . Note that this notion of isomorphism really only cares about the size of the basis:

Proposition 6.5.3 (n -dimensional vector spaces are isomorphic)

If V is an n -dimensional vector space, then $V \cong k^{\oplus n}$.

Question 6.5.4. Let $e_1, \dots, e_n$ be a basis for V . What is the isomorphism? (Your first guess is probably right.)

Remark 6.5.5. You could technically say that all finite-dimensional vector spaces are just $k^{\oplus n}$ and that no other space is worth caring about. But this seems kind of rude. Spaces often are more than just triples: $ax^2 + bx + c$ is a polynomial, and so it has some “essence” to it that you’d lose if you just compressed it into (a, b, c) .

Moreover, a lot of spaces, like the set of vectors (x, y, z) with $x + y + z = 0$, do not have an obvious choice of basis. Thus to cast such a space into $k^{\oplus n}$ would require you to make arbitrary decisions.

§6.6 What is a matrix?

Suppose we have a vector space V with basis $e_1, \dots, e_m$ and a vector space W with basis $w_1, \dots, w_n$. I also have a map $T : V \rightarrow W$ and I want to tell you what T is. It would be awfully inconsiderate of me to try and tell you what $T(v)$ is at every point v . In fact, I only have to tell you what $T(e_1), \dots, T(e_m)$ are, because from there you can work out $T(a_1e_1 + \dots + a_me_m)$ for yourself:

$$T(a_1e_1 + \dots + a_me_m) = a_1T(e_1) + \dots + a_mT(e_m).$$

Since the e_i are a basis, that tells you all you need to know about T .

Example 6.6.1

Let $V = \{ax^2 + bx + c \mid a, b, c \in \mathbb{R}\}$. Then $T(ax^2 + bx + c) = aT(x^2) + bT(x) + cT(1)$.

Now I can even be more concrete. I could tell you what $T(e_1)$ is, but seeing as I have a basis of W , I can actually just tell you what $T(e_1)$ is in terms of this basis. Specifically, there are unique $a_{11}, a_{21}, \dots, a_{n1} \in k$ such that

$$T(e_1) = a_{11}w_1 + a_{21}w_2 + \dots + a_{n1}w_n.$$

So rather than telling you the value of $T(e_1)$ in some abstract space W , I could just tell you what $a_{11}, a_{21}, \dots, a_{n1}$ were. Then I'd just repeat this for $T(e_2)$, $T(e_3)$, all the way up to $T(e_m)$, and that would tell you everything you need to know about T .

That's where the matrix T comes from! It's just a concise way of writing down all mn numbers I need to tell you. To be explicit, the matrix for T is defined as the array

$$T = \left[\underbrace{\begin{array}{c|c|c|c} | & | & & | \\ T(e_1) & T(e_2) & \dots & T(e_m) \\ | & | & & | \end{array}}_m \right] \Bigg\}_n = \begin{bmatrix} a_{11} & a_{12} & \dots & a_{1m} \\ a_{21} & a_{22} & \dots & a_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \dots & a_{nm} \end{bmatrix}.$$

From here you can actually work out for yourself what it means to multiply two matrices. Suppose we have picked a basis for three spaces U, V, W . Given maps $T : U \rightarrow V$ and $S : V \rightarrow W$, we can consider their composition $S \circ T$, i.e.

$$U \xrightarrow{T} V \xrightarrow{S} W.$$

Matrix multiplication is defined exactly so that the matrix ST is the same thing we get from interpreting the composed function $S \circ T$ as a matrix.

Exercise 6.6.2. Check this for yourself! For a concrete example let $\mathbb{R}^2 \xrightarrow{T} \mathbb{R}^2 \xrightarrow{S} T$ by $T(e_1) = 2e_1 + 3e_2$ and $T(e_2) = 4e_1 + 5e_2$, $S(e_1) = 6e_1 + 7e_2$ and $S(e_2) = 8e_1 + 9e_2$. Compute $S(T(e_1))$ and $S(T(e_2))$ and see how it compares to multiplying the matrices associated to S and T .

In particular, since function composition is associative, it follows that matrix multiplication is as well. To drive this point home,

A matrix is the laziest possible way to specify a linear map from V to W .

This means you can define concepts like the determinant or the trace of a matrix both in terms of an “intrinsic” map $T : V \rightarrow W$ and in terms of the entries of the matrix. Since the map T itself doesn't refer to any basis, the abstract definition will imply that the numerical definition doesn't depend on the choice of a basis.

§6.7 Subspaces and picking convenient bases

Prototypical example for this section: Any two linearly independent vectors in $\mathbb{R}^3$.

Definition 6.7.1. Let M be a left R -module. A **submodule** N of M is a module N such that every element of N is also an element of M . If M is a vector space then N is called a **subspace**.

Example 6.7.2 (Kernels)

The **kernel** of a map $T : V \rightarrow W$ is the set of $v \in V$ such that $T(v) = 0_W$. It is a subspace of V , since it's closed under addition and scaling (why?).

Example 6.7.3 (Spans)

Let V be a vector space and $v_1, \dots, v_m$ be any vectors of V . The **span** of these vectors is defined as the set

$$\{a_1v_1 + \dots + a_mv_m \mid a_1, \dots, a_m \in k\}.$$

Note that it is a subspace of V as well!

Question 6.7.4. Why is 0_V an element of each of the above examples? In general, why must any subspace contain 0_V ?

Subspaces behave nicely with respect to bases.

Theorem 6.7.5 (Basis completion)

Let V be an n -dimensional space, and V' a subspace of V . Then

- (a) V' is also finite-dimensional.
- (b) If $e_1, \dots, e_m$ is a basis of V' , then there exist $e_{m+1}, \dots, e_n$ in V such that $e_1, \dots, e_n$ is a basis of V .

Proof. Boring and intuitive. □

This means that given V' a subspace of V , one can pick a basis of V such that V' is the span of some subset of that basis

The same goes for maps, using $\ker T$:

Theorem 6.7.6 (Picking a basis for linear maps)

Let $T : V \rightarrow W$ be a map of finite-dimensional vector spaces. Then there exists a basis $v_1, \dots, v_n$ of V and a basis $w_1, \dots, w_m$ of W , as well as a nonnegative integer k , such that

$$T(v_i) = \begin{cases} w_i & \text{if } i \leq k \\ 0_W & \text{if } i > k. \end{cases}$$

For example, if $V = k^{\oplus 3}$ and $W = k^{\oplus 99}$, T might be the map which sends $e_3 \in V$ to $0 \in W$ and $e_1, e_2 \in V$ to f_{13}, f_{37} in W .

Sketch of Proof. You might like to try this one yourself before reading on.

Let $\ker T$ have dimension $n - k$. We can pick $v_{k+1}, \dots, v_n$ a basis of $\ker T$. Then extend it to a basis $v_1, \dots, v_k$ of V . The map T is injective over the span of $v_1, \dots, v_k$ (since only 0_V is in the kernel) so its images in W are linearly independent. Setting $w_i = T(v_i)$ for each i , we get some linearly independent set in W . Then extend it again to a basis of W . □

The above theorem is quite important: you should keep it in your head. It gives you a picture of how a linear map behaves: in particular, for $T : V \rightarrow W$, one can write $V = \ker T \oplus V'$, so that T annihilates its kernel while sending V' to an isomorphic copy in W .

A corollary of this (which you should have expected anyways) is the so called rank-nullity theorem, which is just the analog of the first isomorphism theorem.

Theorem 6.7.7 (Rank-nullity theorem)

If $T : V \rightarrow W$, then

$$\dim V = \dim \ker T + \dim \operatorname{im} T.$$

§6.8 A cute application: Lagrange interpolation

Here's a cute application¹ of linear algebra to a theorem from high school we all know.

Theorem 6.8.1 (Lagrange interpolation)

Let $x_1, \dots, x_{n+1}$ be distinct real numbers and $y_1, \dots, y_{n+1}$ any real numbers. Then there exists a *unique* polynomial P of degree at most n such that

$$P(x_i) = y_i$$

for every i .

Proof. The idea is to consider the vector space V of polynomials with degree at most n , as well as the vector space $W = \mathbb{R}^{n+1}$.

Question 6.8.2. Check that $\dim V = \dim W$.

Then consider the linear map $T : V \rightarrow W$ given by

$$P \mapsto (P(x_1), \dots, P(x_{n+1})).$$

This is indeed a linear map because, well, $T(P + Q) = T(P) + T(Q)$ and $T(cP) = cT(P)$. It also happens to be injective: if $P \in \ker T$, then $P(x_1) = \dots = P(x_{n+1}) = 0$, but $\deg P = n$ and so P can only be the zero polynomial.

So T is an injective map between vector spaces of the same dimension. Thus it is actually a bijection, which is exactly what we wanted. $\square$

§6.9 (Digression) Arrays of numbers are evil

As I'll stress repeatedly, a matrix represents a *linear map between two vector spaces*. Writing it in the form of an $m \times n$ matrix is merely a very convenient way to see the map concretely. But it obfuscates the fact that this map is, well, a map, not an array of numbers.

If you took high school precalculus, you'll see everything done in terms of matrices. To any typical high school student, a matrix is an array of numbers. No one is sure what

¹Source: Communicated to me by Joe Harris at the first Harvard-MIT Undergraduate Math Symposium.

exactly these numbers represent, but they're told how to magically multiply these arrays to get more arrays. They're told that the matrix

$$\begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \end{bmatrix}$$

is an “identity matrix”, because when you multiply by another matrix it doesn't change. Then they're told that the determinant is some magical combination of these numbers formed by this weird multiplication rule. No one knows what this determinant does, other than the fact that $\det(AB) = \det A \det B$, and something about areas and row operations and Cramer's rule.

Then you go into linear algebra in college, and you do more magic with these arrays of numbers. You're told that two matrices T_1 and T_2 are similar if

$$T_2 = ST_1S^{-1}$$

for some invertible matrix S . You're told that the trace of a matrix $\text{Tr } T$ is the sum of the diagonal entries. Somehow this doesn't change if you look at a similar matrix, but you're not sure why. Then you define the characteristic polynomial as

$$p_T = \det(XI - T).$$

Somehow this also doesn't change if you take a similar matrix, but now you really don't know why. And then you have the Cayley-Hamilton theorem in all its black magic: $p_T(T)$ is the zero map. Out of curiosity you Google the proof, and you find some ad-hoc procedure which still leaves you with no idea why it's true.

This is terrible. Who gives a damn about $T_2 = ST_1S^{-1}$? Only if you know that the matrices are linear maps does this make sense: T_2 is just T_1 rewritten with a different choice of basis.

In my eyes, this is *evil*. Linear algebra is the study of *linear maps*, but it is taught as the study of *arrays of numbers*, and no one knows what these numbers mean. And for a good reason: the numbers are meaningless. They are a highly convenient way of encoding the matrix, but they are not the main objects of study, any more than the dates of events are the main objects of study in history.

The other huge downside is that people get the impression that the only (real) vector space in existence is $\mathbb{R}^{\oplus n}$. As explained in Remark 6.5.5, while you *can* work this way if you're a soulless robot, it's very unnatural for humans to do so.

When I took Math 55a as a freshman at Harvard, I got the exact opposite treatment: we did all of linear algebra without writing down a single matrix. During all this time I was quite confused. What's wrong with a basis? I didn't appreciate until later that this approach was the morally correct way to treat the subject: it made it clear what was happening.

Throughout this project, I've tried to strike a balance between these two approaches, using matrices when appropriate to illustrate the maps and to simplify proofs, but ultimately writing theorems and definitions in their *morally correct* form. I hope that this has both the advantage of giving the “right” definitions while being concrete enough to be digested. But I would just like to say for the record that, if I had to pick between the high school approach and the 55a approach, I would pick 55a in a heartbeat.

§6.10 A word on general modules

Prototypical example for this section: $\mathbb{Z}[\sqrt{2}]$ is a $\mathbb{Z}$ -module of rank two.

I focused mostly on vector fields (aka modules over a field) in this chapter for simplicity, so I want to make a few remarks about modules over a general commutative ring R before concluding.

Firstly, recall that for general modules, we say “generating set” instead of “spanning set”. Shrug.

The main issue with rings is that our key theorem Theorem 6.4.5 fails in spectacular ways. For example, consider $\mathbb{Z}$ as a $\mathbb{Z}$ -module over itself. Then $\{2\}$ is linearly independent, but it cannot be extended to a basis. Similarly, $\{2, 3\}$ is spanning, but one cannot cut it down to a basis. You can see why defining dimension is going to be difficult.

Nonetheless, there are still analogs of some of the definitions above.

Definition 6.10.1. An R -module M is called **finitely generated** if it has a finite generating set.

Definition 6.10.2. An R -module M is called **free** if it has a basis. As said before, the analogue of the dimension theorem holds, and we use the word **rank** to denote the size of the basis. As before, there’s an isomorphism $M \cong R^{\oplus n}$ where n is the rank.

Example 6.10.3 (An example of a $\mathbb{Z}$ -module)

The $\mathbb{Z}$ -module

$$\mathbb{Z}[\sqrt{2}] = \{a + b\sqrt{2} \mid a, b \in \mathbb{Z}\}$$

has a basis $\{1, \sqrt{2}\}$, so we say it is a free $\mathbb{Z}$ -module of rank 2.

Abuse of Notation 6.10.4 (Notation for groups). Recall that an abelian group can be viewed a $\mathbb{Z}$ -module (and in fact vice-versa!), so we can (and will) apply these words to abelian groups. We’ll use the notation $G \oplus H$ for two abelian groups G and H for their Cartesian product, emphasizing the fact that G and H are abelian. This will happen when we study algebraic number theory and homology groups.

§6.11 Problems to think about

Problem 6A (TSTST 2014). Let $P(x)$ and $Q(x)$ be arbitrary polynomials with real coefficients, and let d be the degree of $P(x)$. Assume that $P(x)$ is not the zero polynomial. Prove that there exist polynomials $A(x)$ and $B(x)$ such that

- (i) Both A and B have degree at most $d/2$,
- (ii) At most one of A and B is the zero polynomial,
- (iii) P divides $A + Q \cdot B$.

Problem 6B (Putnam 2003). Do there exist polynomials $a(x)$, $b(x)$, $c(y)$, $d(y)$ such that

$$1 + xy + (xy)^2 = a(x)c(y) + b(x)d(y)$$

holds identically?

Problem 6C* (Idempotents are projection maps). Let $P : V \rightarrow V$ be a linear map, where V is a vector space (not necessarily finite-dimensional). Suppose P is **idempotent**, meaning $P \circ P = P$, or equivalently P is the identity on its image. Prove that $V = \ker P \oplus \operatorname{im} P$. Thus we can think of P as *projection* onto the subspace $\operatorname{im} P$.

Problem 6D*. Let V be a finite dimensional vector space. Let $T : V \rightarrow V$ be a linear map, and let $T^n : V \rightarrow V$ denote T applied n times. Prove that there exists an integer N such that

$$V = \ker T^N \oplus \operatorname{im} T^N.$$

DRAFT (Evan Chen)
Updated February 15, 2017

7 Trace and determinant

You may have learned in high school that given a matrix

$$\begin{bmatrix} a & c \\ b & d \end{bmatrix}$$

the trace is the sum along the diagonals $a + d$ and the determinant is $ad - bc$. But we know that a matrix is somehow just encoding a linear map using a choice of basis. Why would these random formulas somehow not depend on the choice of a basis?

The goal of this chapter is to give the basis-free definition of a determinant and trace of a matrix: that is, we're going to define $\det T$ and $\text{Tr } T$ for $T : V \rightarrow V$ without making reference to the encoding for T . This will make it obvious the determinant and trace of a matrix do not depend on the choice of basis, and that several properties are vacuously true (e.g. that the determinant is multiplicative).

Along the way, we'll meet the tensor product, the dual module, and the wedge product. Fair warning that this makes this one of the harder chapters.

In what follows, all vector spaces are finite-dimensional.

§7.1 Tensor product

Note that $\dim(V \oplus W) = \dim V + \dim W$, even though as sets $V \oplus W$ looks like $V \times W$. What if we wanted a real “product” of spaces, with multiplication of dimensions?

For example, let's pull out (for hopefully the last time) my favorite example of a real vector space, namely

$$V = \{ax^2 + bx + c \mid a, b, c \in \mathbb{R}\}.$$

Here's another space, a little smaller:

$$W = \{dy + e \mid d, e \in \mathbb{R}\}.$$

If we take the direct sum, then we would get some rather unnatural vector space of dimension five (whose elements can be thought of as pairs $(ax^2 + bx + c, dy + e)$). But suppose we want a vector space whose elements are *products* of polynomials in V and W ; it would contain elements like $4x^2y + 5xy + y + 3$. In particular, the basis would be

$$\{x^2y, x^2, xy, x, y, 1\}$$

and thus have dimension six.

For this we resort to the *tensor product*. It does exactly this, except that the “multiplication” is done by a scary¹ symbol $\otimes$: think of it as a “wall” that separates the elements between the two vector spaces. For example, the above example might be written as

$$4x^2 \otimes y + 5x \otimes y + 1 \otimes y + 3 \otimes 1.$$

(This should be read as $(4x^2 \otimes y) + (5x \otimes y) + \dots$; addition comes after $\otimes$.) Of course there should be no distinction between writing $4x^2 \otimes y$ and $x^2 \otimes 4y$ or even $2x^2 \otimes 2y$. While we want to keep the x and y separate, the scalars should be free to float around.

¹Seriously, $\otimes$ looks *terrifying* to non-mathematicians, and even to many math undergraduates.

Of course, there's no need to do everything in terms of just the monomials. We are free to write

$$(x + 1) \otimes (y + 1).$$

If you like, you can expand this as

$$x \otimes y + 1 \otimes y + x \otimes 1 + 1 \otimes 1.$$

Same thing. The point is that we can take any two of our polynomials and artificially “tensor” them together.

The definition of the tensor product does exactly this, and nothing else. I'll only define this for vector spaces for simplicity. The definition for modules over a commutative ring R is exactly the same.

Definition 7.1.1. Let V and W be vector spaces over the same field k . The **tensor product** $V \otimes_k W$ is the abelian group generated by elements of the form $v \otimes w$, subject to relations

$$\begin{aligned} (v_1 + v_2) \otimes w &= v_1 \otimes w + v_2 \otimes w \\ v \otimes (w_1 + w_2) &= v \otimes w_1 + v \otimes w_2 \\ (c \cdot v) \otimes w &= v \otimes (c \cdot w). \end{aligned}$$

As a vector space, its action is given by $c \cdot (v \otimes w) = (c \cdot v) \otimes w = v \otimes (c \cdot w)$.

Here's another way to phrase the same idea. We define a **pure tensor** as an element of the form $v \otimes w$ for $v \in V$ and $w \in W$. But we let the $\otimes$ wall be “permeable” in the sense that

$$(c \cdot v) \otimes w = v \otimes (c \cdot w) = c \cdot (v \otimes w)$$

and we let multiplication and addition distribute as we expect. Finally, the elements of $V \otimes W$ are not the pure tensors, but in fact *sums* of these pure tensors!

As the example we gave suggested, the basis of $V \otimes W$ is literally the “product” of the bases of V and W . In particular, this fulfills our desire that $\dim(V \otimes W) = \dim V \cdot \dim W$.

Proposition 7.1.2 (Basis of $V \otimes W$)

Let V and W be finite-dimensional k -vector spaces. If $e_1, \dots, e_m$ is a basis of V and $f_1, \dots, f_n$ is a basis of W , then the basis of $V \otimes_k W$ is precisely $e_i \otimes f_j$, where $i = 1, \dots, m$ and $j = 1, \dots, n$.

Proof. Omitted; it's easy at least to see that this basis is spanning. □

Example 7.1.3 (Explicit computation)

Let V have basis e_1, e_2 and W have basis f_1, f_2 . Let $v = 3e_1 + 4e_2 \in V$ and $w = 5f_1 + 6f_2 \in W$. Let's write $v \otimes w$ in this basis for $V \otimes_k W$:

$$\begin{aligned} v \otimes w &= (3e_1 + 4e_2) \otimes (5f_1 + 6f_2) \\ &= (3e_1) \otimes (5f_1) + (4e_2) \otimes (5f_1) + (3e_1) \otimes (6f_2) + (4e_2) \otimes (6f_2) \\ &= 15(e_1 \otimes f_1) + 20(e_2 \otimes f_1) + 18(e_1 \otimes f_2) + 24(e_2 \otimes f_2). \end{aligned}$$

So you can see why tensor products are a nice “product” to consider if we're really interested in $V \times W$ in a way that's more intimate than just a direct sum.

Abuse of Notation 7.1.4. We'll almost always abbreviate $\otimes_k$ to just $\otimes$.

Remark 7.1.5. Observe that to define a linear map $V \otimes W \rightarrow X$, I only have to say what happens to each pure tensor $v \otimes w$, since the pure tensors *generate* $V \otimes W$. But again, keep in mind that $V \otimes W$ **consists of sums of these pure tensors!** In other words, $V \otimes W$ is generated by pure tensors.

Remark 7.1.6. Much like the Cartesian product $A \times B$ of sets, you can tensor together any two vector spaces V and W over the same field k ; the relationship between V and W is completely irrelevant. One can think of the $\otimes$ as a “wall” through which one can pass scalars in k , but otherwise keeps the elements of V and W separated. Thus, $\otimes$ is **content-agnostic**.

This also means that even if V and W have some relation to each other, the tensor product doesn't remember this. So for example $v \otimes 1 \neq 1 \otimes v$, just like $(g, 1_G) \neq (1_G, g)$ in the group $G \times G$.

§7.2 Dual module

Prototypical example for this section: Rotate a column matrix by 90 degrees.

Consider the following vector space:

Example 7.2.1 (Functions from $\mathbb{R}^3 \rightarrow \mathbb{R}$)

The set of real functions $f(x, y, z)$ is an infinite-dimensional real vector space. Indeed, we can add two functions to get $f + g$, and we can think of functions like $2f$.

This is a terrifyingly large vector space, but you can do some reasonable reductions. For example, you can restrict your attention to just the *linear maps* from $\mathbb{R}^3$ to $\mathbb{R}$.

That's exactly what we're about to do. This definition might seem strange at first, but bear with me.

Definition 7.2.2. Let V be a k -vector space. Then $V^\vee$, the **dual space** of V , is defined as the vector space whose elements are *linear maps from V to k* .

The addition and multiplication are pointwise: it's the same notation we use when we write $cf + g$ to mean $c \cdot f(x) + g(x)$. The dual space itself is less easy to think about.

Let's try to find a basis for $V^\vee$. First, here is a very concrete interpretation of the vector space. Suppose for example $V = \mathbb{R}^3$. We can think of elements of V as column matrices, like

$$v = \begin{bmatrix} 2 \\ 5 \\ 9 \end{bmatrix} \in V.$$

Then a linear map $f : V \rightarrow k$ can be interpreted as a *row matrix*:

$$f = [3 \quad 4 \quad 5] \in V^\vee.$$

Then

$$f(v) = [3 \quad 4 \quad 5] \begin{bmatrix} 2 \\ 5 \\ 9 \end{bmatrix} = 71.$$

More precisely: **to specify a linear map $V \rightarrow k$, I only have to tell you where each basis element of V goes.** In the above example, f sends e_1 to 3, e_2 to 4, and e_3 to 5. So f sends

$$2e_1 + 5e_2 + 9e_3 \mapsto 2 \cdot 3 + 5 \cdot 4 + 9 \cdot 5 = 71.$$

Let's make all this precise.

Proposition 7.2.3 (The dual basis for $V^\vee$)

Let V be a finite-dimensional vector space with basis $e_1, \dots, e_n$. For each i consider the function $e_i^\vee : V \rightarrow k$ by

$$e_i^\vee(e_j) = \begin{cases} 1 & i = j \\ 0 & i \neq j. \end{cases}$$

In more humane terms, $e_i^\vee(v)$ gives the coefficient of e_i in v .

Then $e_1^\vee, e_2^\vee, \dots, e_n^\vee$ is a basis of $V^\vee$.

I want to emphasize that this choice of basis for $V^\vee$ is absolutely not canonical, and depends very heavily on which basis of V is originally selected. It is not “God-given”, in other words.

Example 7.2.4 (Explicit example of element in $V^\vee$)

In this notation, $f = 3e_1^\vee + 4e_2^\vee + 5e_3^\vee$. Do you see why the “sum” notation works as expected here? Indeed

$$\begin{aligned} f(e_1) &= (3e_1^\vee + 4e_2^\vee + 5e_3^\vee)(e_1) \\ &= 3e_1^\vee(e_1) + 4e_2^\vee(e_1) + 5e_3^\vee(e_1) \\ &= 3 \cdot 1 + 4 \cdot 0 + 5 \cdot 0 = 3. \end{aligned}$$

That's exactly what we wanted.

Also, you might be inclined to point out that $V \cong V^\vee$ at this point, since there's an obvious isomorphism $e_i \mapsto e_i^\vee$. You might call it “rotating the column matrix by 90° ”. I want to point out again that

- This isomorphism is absolutely not canonical, and depends very much on which basis you choose. In other words, if I pick a different basis, then the isomorphism will be completely different.

It *is* true that V and $V^\vee$ are isomorphic, but the fact is that any two k -vector spaces of the same dimension are isomorphic, so the fact that $V \cong V^\vee$ is not especially impressive.

- If V is infinite-dimensional it is no longer necessarily true that $V \cong V^\vee$.

§7.3 The trace

The point of introducing both the tensor product and the dual module is that together we can use them to get an algebraic way of thinking about “matrices”. Here's the intuition. If V is three-dimensional and W is five-dimensional, then we can think of the maps

$V \rightarrow W$ as a 5×3 array of numbers. We want to think of these maps as a vector space: (since one can add or scale matrices). So it had better be a vector space with dimension 15, but just saying “ $k^{\oplus 15}$ ” is not really that satisfying (what is the basis?).

To do better, we consider the tensor product

$$V^\vee \otimes W$$

which somehow is a product of maps out of V and the target space W . We claim that this is in fact the space we want: i.e. **there is a natural bijection between elements of $V^\vee \otimes W$ and linear maps from V to W** .

First, how do we interpret an element of $V^\vee \otimes W$ as a map $V \rightarrow W$? For concreteness, suppose V has a basis e_1, e_2, e_3 , and W has a basis f_1, f_2, f_3, f_4, f_5 . Consider an element of $V^\vee \otimes W$, say

$$e_1^\vee \otimes (f_2 + 2f_4) + 4e_2^\vee \otimes f_5.$$

We want to interpret this element as a function $V \rightarrow W$: so given a $v \in V$, we want to output an element of W . There’s really only one way to do this: feed in $v \in V$ into the $V^\vee$ guys on the left. That is, take the map

$$v \mapsto e_1^\vee(v) \cdot (f_2 + 2f_4) + 4e_2^\vee(v) \cdot f_5 \in W.$$

So, there’s a natural way to interpret any element $\xi_1 \otimes w_1 + \cdots + \xi_m \otimes w_m \in V^\vee \otimes W$ as a linear map $V \rightarrow W$. The claim is that in fact, every linear map $V \rightarrow W$ has a unique such interpretation.

First, for notational convenience,

Definition 7.3.1. Let $\text{Hom}(V, W)$ be the set of linear maps from V to W (which one can interpret as matrices which send V to W), viewed as a vector space over k . Let $\text{Mat}(V) = \text{Hom}(V, V)$. (Hom and Mat stand for homomorphism and matrices.)

In this notation, $V^\vee \stackrel{\text{def}}{=} \text{Hom}(V, k)$.

We can now write down something that’s more true generally.

Theorem 7.3.2 ($V^\vee \otimes W \iff \text{linear maps } V \rightarrow W$)

Let V and W be finite-dimensional vector spaces. We described a map

$$\Psi : V^\vee \otimes W \rightarrow \text{Hom}(V, W)$$

by sending $\xi_1 \otimes w_1 + \cdots + \xi_m \otimes w_m$ to the linear map

$$v \mapsto \xi_1(v)w_1 + \cdots + \xi_m(v)w_m.$$

Then Ψ is an isomorphism of vector spaces, i.e. every linear map $V \rightarrow W$ can be uniquely represented as an element of $V^\vee \otimes W$ in this way.

The beauty is that we don’t actually have to refer to the basis at all to state the theorem.

Proof. This looks intimidating, but it’s actually not difficult. We proceed in two steps:

1. First, we check that Ψ is *surjective*; every linear map has at least one representation of this form. To see this, take any $T : V \rightarrow W$. Suppose V has basis e_1, e_2, e_3 and that $T(e_1) = w_1$, $T(e_2) = w_2$ and $T(e_3) = w_3$. Then the element

$$e_1^\vee \otimes w_1 + e_2^\vee \otimes w_2 + e_3^\vee \otimes w_3$$

works, as it is contrived to agree with T on the basis elements e_i .

2. So it suffices to check now that $\dim V^\vee \otimes W = \dim \operatorname{Hom}(V, W)$. Certainly, $V^\vee \otimes W$ has dimension $\dim V \cdot \dim W$. But by viewing $\operatorname{Hom}(V, W)$ as $\dim V \cdot \dim W$ matrices, we see that it too has dimension $\dim V \cdot \dim W$. $\square$

So there is a **natural isomorphism** $V^\vee \otimes W \cong \operatorname{Hom}(V, W)$. While we did use a basis liberally in the *proof that it works*, this doesn't change the fact that the isomorphism is “God-given”, depending only on the spirit of V and W itself and not which basis we choose to express the vector spaces in.

The above is perhaps a bit dense, so here is a concrete example.

Example 7.3.3 (Explicit example)

Let $V = \mathbb{R}^2$ and take a basis e_1, e_2 of V . Then define $T : V \rightarrow V$ by

$$T = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}.$$

Then the isomorphism sends T to

$$T \mapsto e_1^\vee \otimes e_1 + 2e_2^\vee \otimes e_1 + 3e_1^\vee \otimes e_2 + 4e_2^\vee \otimes e_2 \in V^\vee \otimes V.$$

The beauty of the basis-free definition is that even if we change the basis, the above expression will look completely different, but the actual element in $V^\vee \otimes V$ doesn't change.

We are now ready to give the definition of a trace. Recall that a square matrix T can be thought of as a map $T : V \rightarrow V$. According to the above theorem there is a map

$$\operatorname{Mat}(V) \cong V^\vee \otimes V$$

so every map $V \rightarrow V$ can be thought of as an element of $V^\vee \otimes V$. But we can also define an *evaluation map* $\operatorname{ev} : V^\vee \otimes V \rightarrow k$ by “collapsing” each pure tensor: $f \otimes v \mapsto f(v)$. So this gives us a composed map

$$\operatorname{Mat}(V) \xrightarrow{\cong} V^\vee \otimes V \xrightarrow{\operatorname{ev}} k.$$

This result is called the **trace** of a matrix T .

Example 7.3.4 (Example of a trace)

Continuing the previous example,

$$\operatorname{Tr} T = e_1^\vee(e_1) + 2e_2^\vee(e_1) + 3e_1^\vee(e_2) + 4e_2^\vee(e_2) = 1 + 0 + 0 + 4 = 5.$$

And that is why the trace is the sum of the diagonal entries.

§7.4 (Optional) Two more digressions on dual spaces

- Suppose $T : V \rightarrow W$ is a linear map. Then we actually get a map $T^\vee : W^\vee \rightarrow V^\vee$, called the **dual map**, as follows: $T^\vee(f) = f \circ T$. In the concrete interpretation, if T is a matrix, then T multiplies column vectors in V on the left to get column vectors in W . Then $T^\vee$ is the same matrix, but it multiplies row vectors in $W^\vee$

on the right to get row vectors in V . Here is an explicit example: if $V = \mathbb{R}^3$ and $W = \mathbb{R}^2$ then

$$\begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} = \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} \text{ and } \begin{bmatrix} w'_1 & w'_2 \end{bmatrix} \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix} = \begin{bmatrix} v'_1 & v'_2 & v'_3 \end{bmatrix}.$$

But once again, the definition $T^\vee(f) = f \circ T$ has the nice property that it doesn't depend on the choice of a basis. That shows that $T^\vee$ really is well defined, depending only on $T : V \rightarrow W$. And in particular, this definition works even if V or W is infinite-dimensional.

- Why is the dual so named? We'll show that the dual of the dual (i.e. $(V^\vee)^\vee$) is somehow exactly the same as V . It's of course true that they are isomorphic, but somehow that's not very impressive (any vector spaces of the same dimension are isomorphic). But we will show that there is a very natural isomorphism between them; this will be more meaningful.

Let V be a finite-dimensional vector space. For every $v \in V$, we define the map $\text{ev}_v : V^\vee \rightarrow k$ called "evaluation at v ": it takes a function $f \in V^\vee$, and since $f : V \rightarrow k$, we can return the value $f(v)$. Succinctly, $\text{ev}_v(f) = f(v)$.

Since $\text{ev}_v : V^\vee \rightarrow k$, we see that $\text{ev}_v \in (V^\vee)^\vee$. In fact, you can show that the map $v \mapsto \text{ev}_v$ is an isomorphism from V to $(V^\vee)^\vee$ for dimension reasons. So unlike the situation $V \cong V^\vee$, there is a *canonical* isomorphism

$$V \cong (V^\vee)^\vee$$

that doesn't depend at all on what V looks like, or on an arbitrary choice of basis. In other words, the structures of V and $(V^\vee)^\vee$ really are the same at heart.

§7.5 Wedge product

Prototypical example for this section: $\Lambda^2(\mathbb{R}^2)$.

We're now going to define something called the wedge product. It will look a ton like the tensor product, but we'll have one extra relation.

For simplicity, I'll first define the wedge product $\Lambda^2(V)$. But we will later replace 2 with any n .

Definition 7.5.1. Let V be a k -vector space. The 2-wedge product $\Lambda^2(V)$ is the abelian group generated by elements of the form $v \wedge w$ (where $v, w \in V$), subject to relations

$$\begin{aligned} (v_1 + v_2) \wedge w &= v_1 \wedge w + v_2 \wedge w \\ v \wedge (w_1 + w_2) &= v \wedge w_1 + v \wedge w_2 \\ (c \cdot v) \wedge w &= v \wedge (c \cdot w) \\ v \wedge v &= 0. \end{aligned}$$

As a vector space, its action is given by $c \cdot (v \wedge w) = (c \cdot v) \wedge w = v \wedge (c \cdot w)$.

Exercise 7.5.2. Use these properties to show that $v \wedge w = -(w \wedge v)$ for any $v, w \in V$.

This looks almost exactly the same as the definition for a tensor product, with two subtle differences. The first is that we only have V now, rather than V and W . Secondly, there is a new *mysterious* relation $v \wedge v = 0$, which implies that

$$v \wedge w = -(w \wedge v).$$

What's that doing there? It seems kind of weird.

I'll give you a hint.

Example 7.5.3 (Wedge product explicit computation)

Let $V = \mathbb{R}^2$, and let $v = ae_1 + be_2$, $w = ce_1 + de_2$. Now let's compute $v \wedge w$ in $\Lambda^2(V)$.

$$\begin{aligned} v \wedge w &= (ae_1 + be_2) \wedge (ce_1 + de_2) \\ &= ac(e_1 \wedge e_1) + bd(e_2 \wedge e_2) + ad(e_1 \wedge e_2) + bc(e_2 \wedge e_1) \\ &= ad(e_1 \wedge e_2) + bc(e_2 \wedge e_1) \\ &= (ad - bc)(e_1 \wedge e_2). \end{aligned}$$

What is $ad - bc$? You might already recognize it:

- You might know that the area of the parallelogram formed by v and w is $ad - bc$.
- You might recognize it as the determinant of $\begin{bmatrix} a & c \\ b & d \end{bmatrix}$. In fact, you might even know that the determinant is meant to interpret hypervolumes.

This is absolutely no coincidence. The wedge product is designed to interpret signed areas. That is, $v \wedge w$ is meant to interpret the area of the parallelogram formed by v and w . You can see why the condition $(cv) \wedge w = v \wedge (cw)$ would make sense now. And now of course you know why $v \wedge v$ ought to be zero: it's an area zero parallelogram!

The **miracle of wedge products** is that the only additional condition we need to add to the tensor product axioms is that $v \wedge w = -(w \wedge v)$. Then suddenly, the wedge will do all our work of interpreting volumes for us.

Question 7.5.4. Let V be a real finite-dimensional vector space. Convince yourself that if $e_1, \dots, e_n$ is a basis of V , then a basis of $\Lambda^2(V)$ is $e_i \wedge e_j$ where $i < j$. Hence $\Lambda^2(V)$ has dimension $\binom{n}{2}$.

Now I have the courage to define a multi-dimensional wedge product. It's just the same thing with more wedges.

Definition 7.5.5. Let V be a vector space and m a positive integer. The space $\Lambda^m(V)$ is generated by wedges of the form

$$v_1 \wedge v_2 \wedge \dots \wedge v_m$$

subject to relations

$$\begin{aligned} \dots \wedge (v_1 + v_2) \wedge \dots &= (\dots \wedge v_1 \wedge \dots) + (\dots \wedge v_2 \wedge \dots) \\ \dots \wedge (cv_1) \wedge v_2 \wedge \dots &= \dots \wedge v_1 \wedge (cv_2) \wedge \dots \\ \dots \wedge v \wedge v \wedge \dots &= 0. \end{aligned}$$

As a vector space

$$c \cdot (v_1 \wedge v_2 \wedge \dots \wedge v_m) = (cv_1) \wedge v_2 \wedge \dots \wedge v_m = v_1 \wedge (cv_2) \wedge \dots \wedge v_m = \dots$$

This definition is pretty wordy, but in English the three conditions say

- We should be able to add products like before,
- You can put constants onto any of the m components (as is directly pointed out in the “vector space” action), and
- Switching any two adjacent wedges negates the whole wedge.

It’s a pretty standard generalization of $\Lambda^2(V)$. You can convince yourself that any element of the form

$$\cdots \wedge v \wedge \cdots \wedge v \wedge \cdots$$

should still be zero.

Just like $e_1 \wedge e_2$ was a basis earlier, we can find the basis for general m and n .

Proposition 7.5.6 (Basis of the wedge product)

Let V be a vector space with basis $e_1, \dots, e_n$. A basis for $\Lambda^m(V)$ consists of the elements

$$e_{i_1} \wedge e_{i_2} \wedge \cdots \wedge e_{i_m}$$

where

$$1 \leq i_1 < i_2 < \cdots < i_m \leq n.$$

Hence $\Lambda^m(V)$ has dimension $\binom{n}{m}$.

Sketch of Proof. We knew earlier that $e_{i_1} \otimes \cdots \otimes e_{i_m}$ was a basis for the tensor product. Here we have the additional property that (a) if two basis elements re-appear then the whole thing becomes zero, and (b) we can shuffle around elements, and so we arbitrarily decide to put the basis elements in increasing order. $\square$

§7.6 The determinant

Prototypical example for this section: $(ae_1 + be_2) \wedge (ce_1 + de_2) = (ad - bc)e_1 \wedge e_2$.

Now we’re ready to define the determinant. Suppose $T : V \rightarrow V$ is a square matrix. We claim that the map $\Lambda^m(V) \rightarrow \Lambda^m(V)$ given on wedges by

$$v_1 \wedge v_2 \wedge \cdots \wedge v_m \mapsto T(v_1) \wedge T(v_2) \wedge \cdots \wedge T(v_m)$$

is a linear map. You can check this yourself if you like. We call that map $\Lambda^m(T)$.

Example 7.6.1 (Example of $\Lambda^m(T)$)

In $V = \mathbb{R}^4$ with standard basis $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3, \mathbf{e}_4$, let $T(\mathbf{e}_1) = \mathbf{e}_2$, $T(\mathbf{e}_2) = 2\mathbf{e}_3$, $T(\mathbf{e}_3) = \mathbf{e}_3$ and $T(\mathbf{e}_4) = 2\mathbf{e}_2 + \mathbf{e}_3$. Then, for example, $\Lambda^2(T)$ sends

$$\begin{aligned} \mathbf{e}_1 \wedge \mathbf{e}_2 + \mathbf{e}_3 \wedge \mathbf{e}_4 &\mapsto T(\mathbf{e}_1) \wedge T(\mathbf{e}_2) + T(\mathbf{e}_3) \wedge T(\mathbf{e}_4) \\ &= \mathbf{e}_2 \wedge 2\mathbf{e}_3 + \mathbf{e}_3 \wedge (2\mathbf{e}_2 + \mathbf{e}_3) \\ &= 2(\mathbf{e}_2 \wedge \mathbf{e}_3 + \mathbf{e}_3 \wedge \mathbf{e}_2) \\ &= 0. \end{aligned}$$

Now here's something interesting. Suppose V has dimension n , and let $m = n$. Then $\Lambda^n(V)$ has dimension $\binom{n}{n} = 1$ — it's a one dimensional space! Hence $\Lambda^n(V) \cong k$.

So $\Lambda^n(T)$ can be thought of as a linear map from k to k . But we know that *a linear map from k to k is just multiplication by a constant*. Hence $\Lambda^n(T)$ is multiplication by some constant.

Definition 7.6.2. Let $T : V \rightarrow V$, where V is an n -dimensional vector space. Then $\Lambda^n(T)$ is multiplication by a constant c ; we define the **determinant** of T as $c = \det T$.

Example 7.6.3 (The determinant of a 2×2 matrix)

Let $V = \mathbb{R}^2$ again with basis $\mathbf{e}_1$ and $\mathbf{e}_2$. Let

$$T = \begin{bmatrix} a & c \\ b & d \end{bmatrix}.$$

In other words, $T(\mathbf{e}_1) = a\mathbf{e}_1 + b\mathbf{e}_2$ and $T(\mathbf{e}_2) = c\mathbf{e}_1 + d\mathbf{e}_2$.

Now let's consider $\Lambda^2(V)$. It has a basis $\mathbf{e}_1 \wedge \mathbf{e}_2$. Now $\Lambda^2(T)$ sends it to

$$\mathbf{e}_1 \wedge \mathbf{e}_2 \xrightarrow{\Lambda^2(T)} T(\mathbf{e}_1) \wedge T(\mathbf{e}_2) = (a\mathbf{e}_1 + b\mathbf{e}_2) \wedge (c\mathbf{e}_1 + d\mathbf{e}_2) = (ad - bc)(\mathbf{e}_1 \wedge \mathbf{e}_2).$$

So $\Lambda^2(T) : \Lambda^2(V) \rightarrow \Lambda^2(V)$ is multiplication by $\det T = ad - bc$, because it sent $\mathbf{e}_1 \wedge \mathbf{e}_2$ to $(ad - bc)(\mathbf{e}_1 \wedge \mathbf{e}_2)$.

And that is the definition of a determinant. Once again, since we defined it in terms of $\Lambda^n(T)$, this definition is totally independent of the choice of basis. In other words, the determinant can be defined based on $T : V \rightarrow V$ alone without any reference to matrices.

Question 7.6.4. Why does $\Lambda^n(S \circ T) = \Lambda^n(S) \circ \Lambda^n(T)$? (Remember that these are one-dimensional maps.)

In this way, we also get

$$\det(S \circ T) = \det(S) \det(T)$$

for free. The general nasty formula for a determinant in terms of the matrix also follows from our work, and is just a generalization of the work we did for $n = 2$. Simply write out

$$(a_{11}e_1 + a_{21}e_2 + \dots) \wedge \dots \wedge (a_{1n}e_1 + a_{2n}e_2 + \dots + a_{nn}e_n)$$

and do a full expansion.

Exercise 7.6.5. Convince yourself this gives the right answer. (For example, expand this for $n = 3$.)

§7.7 Problems to think about

Problem 7A. Show that for any real numbers x_{ij} (here $1 \leq i, j \leq n$) we have

$$\det \begin{bmatrix} x_{11} & x_{12} & \dots & x_{1n} \\ x_{21} & x_{22} & \dots & x_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ x_{n1} & x_{n2} & \dots & x_{nn} \end{bmatrix} = \det \begin{bmatrix} x_{11} + x_{12} & x_{12} & \dots & x_{1n} \\ x_{21} + x_{22} & x_{22} & \dots & x_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ x_{n1} + x_{n2} & x_{n2} & \dots & x_{nn} \end{bmatrix}.$$



Problem 7B[†] (Traces kind of commute). Let $T : V \rightarrow W$ and $S : W \rightarrow V$ be linear maps between finite-dimensional vector spaces V and W . Show that

$$\text{Tr}(T \circ S) = \text{Tr}(S \circ T).$$

Problem 7C[†] (Product of traces). Let $T : V \rightarrow V$ and $S : W \rightarrow W$ be linear maps of finite-dimensional vector spaces V and W . Consider $T \otimes S : V \otimes W \rightarrow V \otimes W$. Prove that

$$\text{Tr}(T \otimes S) = \text{Tr}(T) \text{Tr}(S).$$



Problem 7D (Sweden 2010, edited). A herd of 1000 cows of nonzero weight is given. Prove that we can remove one cow such that the remaining 999 cows cannot be split into two halves of equal weights.



Problem 7E. Let V be a vector space over K and $T : V \rightarrow V$. Show that

$$\det(a \cdot \text{id}_V - T) = \sum_{n=0}^{\dim V} a^{\dim V - n} \cdot (-1)^n \text{Tr}(\Lambda^n(T))$$

where the trace is taking by viewing $\Lambda^n(T) : \Lambda^n(V) \rightarrow \Lambda^n(V)$.

8 Spectral theory

This chapter will develop the theory of eigenvalues and eigenvectors, the so-called “Jordan canonical form”, and from it the characteristic polynomial. In what follows, V is a finite-dimensional vector space over a ground field k .

§8.1 Why you should care

We know that a square matrix T is really just a linear map from V to V . What’s the simplest type of linear map? It would just be multiplication by some scalar λ , which would have associated matrix

$$\begin{bmatrix} \lambda & 0 & \dots & 0 \\ 0 & \lambda & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \lambda \end{bmatrix}.$$

That’s perhaps *too* simple, though. If we had a basis $e_1, \dots, e_n$ then another very “simple” operation would just be scaling each basis element by λ_i , i.e. a **diagonal** matrix of the form

$$T = \begin{bmatrix} \lambda_1 & 0 & \dots & 0 \\ 0 & \lambda_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \lambda_n \end{bmatrix}.$$

These maps are more general. Indeed, you can, for example, compute T^{100} in a heartbeat: the map sends $e_1 \rightarrow \lambda_1^{100} e_1$. Try doing that with an arbitrary $n \times n$ matrix.

Of course, most linear maps are probably not that nice. Or are they? Let’s consider a map $T : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ by $e_1 \mapsto 2e_1$ and $e_2 \mapsto e_1 + 3e_2$. This doesn’t look anywhere as nice until we realize we can rewrite it as

$$\begin{aligned} e_1 &\mapsto 2e_1 \\ e_1 + e_2 &\mapsto 3(e_1 + e_2). \end{aligned}$$

So our completely random-looking map, under a suitable change of basis, looks like the very nice maps we described before! That’s the motivation for this chapter.

§8.2 Eigenvectors and eigenvalues

Let k be a field and V a vector space over it. In the above example, we saw that there were two very nice vectors, e_1 and $e_1 + e_2$, for which V did something very simple. Naturally, these vectors have a name.

Definition 8.2.1. Let $T : V \rightarrow V$ and $v \in V$, $v \neq 0$. We say that v is an **eigenvector** if $T(v) = \lambda v$ for some $\lambda \in k$ (possibly zero, but remember $v \neq 0$). The value λ is called an **eigenvalue** of T . (Of course, no mention to a basis anywhere.)

Unfortunately, it’s not exactly true that eigenvalues always exist.

Example 8.2.2 (Eigenvalues need not exist)

Let $V = \mathbb{R}^2$ and let T be the map which rotates a vector by 90° around the origin. Then $T(v)$ is not a multiple of v for any $v \in V$, other than the trivial $v = 0$.

However, it is true if we replace k with an algebraically closed field ¹

Theorem 8.2.3 (Eigenvalues always exist over algebraically closed fields)

Suppose k is also algebraically closed field. Then if $T : V \rightarrow V$ is a linear map, there exists an eigenvalue $\lambda \in k$.

Proof. (From [Ax97]) The idea behind this proof is to consider “polynomials” in T . For example, $2T^2 - 4T + 5$ would be shorthand for $2T(T(v)) - 4T(v) + 5v$. In this way we can consider “polynomials” $P(T)$; this lets us tie in the “algebraically closed” condition. These polynomials behave nicely:

Question 8.2.4. Show that $P(T) + Q(T) = (P + Q)(T)$ and $P(T) \circ Q(T) = (P \cdot Q)(T)$.

Let $n = \dim V < \infty$ and fix any nonzero vector $v \in V$, and consider vectors $v, T(v), \dots, T^n(v)$. They must be linearly dependent for dimension reasons; thus there is a nonzero polynomial P such that $P(T)$ is zero when applied to v . WLOG suppose P is a monic polynomial, and thus $P(z) = (z - r_1) \dots (z - r_m)$ say. Then we get

$$0 = (T - r_1 \text{id}) \circ (T - r_2 \text{id}) \circ \dots \circ (T - r_m \text{id})(v)$$

which means at least one of $T - r_i \text{id}$ is not injective, done. $\square$

So in general we like to consider algebraically closed fields. This is not a big loss: any real matrix can be interpreted as a complex matrix whose entries just happen to be real, for example.

§8.3 The Jordan form

So that you know exactly where I’m going, here’s the main theorem.

¹A field is **algebraically closed** if all its polynomials have roots, the archetypal example being $\mathbb{C}$.

Theorem 8.3.1 (Jordan canonical form)

Let $T : V \rightarrow V$ be a linear map over an algebraically closed field k . Then we can choose a basis of V such that the matrix T has the form

$$\begin{bmatrix} \lambda_1 & 1 & & & & \\ 0 & \lambda_1 & & & & \\ & & \lambda_2 & & & \\ & & & \lambda_3 & 1 & 0 \\ & & & 0 & \lambda_3 & 1 \\ & & & 0 & 0 & \lambda_3 \\ & & & & & \ddots \\ & & & & & & \lambda_m & 1 \\ & & & & & & 0 & \lambda_m \end{bmatrix}$$

Each of the blocks is called a **Jordan block**. (The blank entries are all zero.)

What does this mean? Basically, it means our dream is almost true. What happens is that V can get broken down as a direct sum

$$V = J_1 \oplus J_2 \oplus \cdots \oplus J_m$$

and T acts on each of these subspaces independently. These subspaces correspond to the blocks in the matrix above. In the simplest case, $\dim J_i = 1$, so J_i has a basis element e for which $T(e) = \lambda_i e$; in other words, we just have a simple eigenvalue. But on occasion, the situation is not quite so simple, and we have a block of size greater than 1; this leads to 1's just above the diagonals. I'll explain later how to interpret the 1's. For now, you should note that even if $\dim J_i \geq 2$, we still have a basis element which is an eigenvector with element λ_i .

Example 8.3.2

Let $T : k^6 \rightarrow k^6$ and suppose T is given by the matrix

$$T = \begin{bmatrix} 3 & 0 & 0 & 0 & 0 & 0 \\ 0 & 7 & 0 & 0 & 0 & 0 \\ 0 & 0 & 2 & 1 & 0 & 0 \\ 0 & 0 & 0 & 2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 5 & 0 \\ 0 & 0 & 0 & 0 & 0 & 3 \end{bmatrix}.$$

Reading the matrix, we can compute all the eigenvectors and eigenvalues: they are

$$T(e_1) = 3e_1$$

$$T(e_2) = 7e_2$$

$$T(e_3) = 2e_3$$

$$T(e_5) = 5e_5$$

$$T(e_6) = 3e_6.$$

Here's the idea behind the proof. We would like to be able to break down V directly into components as above, but some of the λ_i 's from the different Jordan blocks could

coincide, and this turns out to prove a technical difficulty for detecting them. So instead, we're going to break down V first into subspaces based on the values of λ 's; these are called the **generalized eigenspaces**. In other words, the generalized eigenspace of a given λ is just all the Jordan blocks which have eigenvalue λ . Only after that will we break the generalized eigenspace into the individual Jordan blocks.

The sections below record this proof in the other order: the next part deals with breaking generalized eigenspaces into Jordan blocks, and the part after that is the one where we break down V into generalized eigenspaces.

§8.4 Nilpotent maps

Bear with me for a moment. First, define:

Definition 8.4.1. A map $T : V \rightarrow V$ is **nilpotent** if T^m is the zero map for some integer m . (Here T^m means “ T applied m times”.)

What's an example of a nilpotent map?

Example 8.4.2 (The “Descending staircase”)

Let $V = k^3$ have basis e_1, e_2, e_3 . Then the map T which sends

$$e_3 \mapsto e_2 \mapsto e_1 \mapsto 0$$

is nilpotent, since $T(e_1) = T^2(e_2) = T^3(e_3) = 0$, and hence $T^3(v) = 0$ for all $v \in V$.

That's a pretty nice example. As another example, we can have multiple such staircases.

Example 8.4.3 (Double staircase)

Let $V = k^{\oplus 5}$ have basis e_1, e_2, e_3, e_4, e_5 . Then the map

$$e_3 \mapsto e_2 \mapsto e_1 \mapsto 0 \text{ and } e_5 \mapsto e_4 \mapsto 0$$

is nilpotent.

However this isn't really that different from the previous example; it's just the same idea repeated multiple times. And in fact we now claim that *all* nilpotent maps have essentially that form.

Theorem 8.4.4 (Nilpotent Jordan)

Let $T : V \rightarrow V$ be a nilpotent map. Then we can write $V = \bigoplus_{i=1}^m V_i$ where each V_i has a basis of the form $v_i, T(v_i), \dots, T^{\dim V_i - 1}(v_i)$ for some $v_i \in V_i$. Hence every nilpotent map can be viewed as independent staircases.

Each chain $v_i, T(v_i), T(T(v_i)), \dots$ is just one staircase. The proof is given later, but first let me point out where this is going.

Here's the punch line. Let's take the double staircase again. Expressing it as a matrix

gives, say

$$S = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

Then we can compute

$$S + \lambda \text{id} = \begin{bmatrix} \lambda & 1 & 0 & 0 & 0 \\ 0 & \lambda & 1 & 0 & 0 \\ 0 & 0 & \lambda & 0 & 0 \\ 0 & 0 & 0 & \lambda & 1 \\ 0 & 0 & 0 & 0 & \lambda \end{bmatrix}.$$

It's a bunch of Jordan blocks! This gives us a plan to proceed: we need to break V into a bunch of subspaces such that $T - \lambda \text{id}$ is nilpotent over each subspace. Then Nilpotent Jordan will finish the job.

§8.5 Reducing to the nilpotent case

Definition 8.5.1. Let $T : V \rightarrow V$. A subspace $W \subseteq V$ is called **T -invariant** if $T(w) \in W$ for any $w \in W$. In this way, T can be thought of as a map $W \rightarrow W$.

In this way, the Jordan form is a decomposition of V into invariant subspaces.

Now I'm going to be cheap, and define:

Definition 8.5.2. A map $T : V \rightarrow V$ is called **indecomposable** if it's impossible to write $V = W_1 \oplus W_2$ where both W_1 and W_2 are T -invariant.

Picture of a *decomposable* map:

$$\left[\begin{array}{c|ccc} W_1 & 0 & 0 & 0 \\ & 0 & 0 & 0 \\ \hline 0 & 0 & & \\ 0 & 0 & & \\ 0 & 0 & & \end{array} \right] \begin{array}{c} \\ \\ W_2 \\ \\ \end{array}$$

As you might expect, we can break a space apart into “indecomposable” parts.

Proposition 8.5.3

Given any map $T : V \rightarrow V$, we can write

$$V = V_1 \oplus V_2 \oplus \cdots \oplus V_m$$

where each V_i is T -invariant, and for any i the map $T : V_i \rightarrow V_i$ is indecomposable.

Proof. Same as the proof that every integer is the product of primes. $\square$

Incredibly, with just that we're almost done! Consider a decomposition as above, so that $T : V_1 \rightarrow V_1$ is an indecomposable map. Then T has an eigenvalue λ_1 , so let $S = T - \lambda_1 \text{id}$; hence $\ker S \neq \{0\}$.

Question 8.5.4. Show that V_1 is also S -invariant, so we can consider $S : V_1 \rightarrow V_1$.

By Problem 6D*, we have

$$V_1 = \ker S^N \oplus \operatorname{im} S^N$$

for some N . But we assumed T was indecomposable, so this can only happen if $\operatorname{im} S^N = \{0\}$ and $\ker S^N = V$ (since $\ker S^N$ contains our eigenvector). Hence S is nilpotent, so it's a collection of staircases. In fact, since T is indecomposable, there is only one staircase. Hence V_1 is a Jordan block, as desired.

§8.6 (Optional) Proof of nilpotent Jordan

The proof is just induction on $\dim V$. Assume $\dim V \geq 1$, and let $W = T(V)$ be the image of V . Since T is nilpotent, we must have $W \subsetneq V$. Moreover, if $W = \{0\}$ (i.e. T is the zero map) then we're already done. So assume $\{0\} \subsetneq W \subsetneq V$.

By the inductive hypothesis, we can select a good basis of W :

$$\mathcal{B}' = \left\{ \begin{aligned} &T(v_1), T(T(v_1)), T(T(T(v_1))), \dots \\ &T(v_2), T(T(v_2)), T(T(T(v_2))), \dots \\ &\dots, \\ &T(v_\ell), T(T(v_\ell)), T(T(T(v_\ell))), \dots \end{aligned} \right\}$$

for some $T(v_i) \in W$ (here we have taken advantage of the fact that each element of W is itself of the form $T(v)$ for some v).

Also, note that there are exactly ℓ elements of $\mathcal{B}'$ which are in $\ker T$ (namely the last element of each of the ℓ staircases). We can thus complete it to a basis $v_{\ell+1}, \dots, v_m$ (where $m = \dim \ker T$). (In other words, the last element of each staircase plus the $m - \ell$ new ones are a basis for $\ker T$.)

Now consider

$$\mathcal{B} = \left\{ \begin{aligned} &v_1, T(v_1), T(T(v_1)), T(T(T(v_1))), \dots \\ &v_2, T(v_2), T(T(v_2)), T(T(T(v_2))), \dots \\ &\dots, \\ &v_\ell, T(v_\ell), T(T(v_\ell)), T(T(T(v_\ell))), \dots \\ &v_{\ell+1}, v_{\ell+2}, \dots, v_m \end{aligned} \right\}.$$

Question 8.6.1. Check that there are exactly $\ell + \dim W + (\dim \ker T - \ell) = \dim V$ elements.

Exercise 8.6.2. Show that all the elements are linearly independent. (Assume for contradiction there is some linear dependence, then take T of both sides.)

Hence $\mathcal{B}$ is a basis of the desired form.

§8.7 Characteristic polynomials, and Cayley-Hamilton

Take a map $T : V \rightarrow V$, where V is n -dimensional, and suppose its eigenvalues are $a_1, a_2, \dots, a_n$ (with repetition). Then the **characteristic polynomial** is given by

$$p_T(X) = (X - a_1)(X - a_2) \dots (X - a_n).$$

Note that if we've written T in Jordan form, i.e.

$$T = \begin{bmatrix} a_1 & * & 0 & \dots & 0 \\ 0 & a_2 & * & \dots & 0 \\ 0 & 0 & a_3 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & a_n \end{bmatrix}$$

(here each $*$ is either 0 or 1), then we can “hack” together the “immoral” definition

$$p_T(X) \stackrel{\text{def}}{=} \det \begin{bmatrix} X - a_1 & * & 0 & \dots & 0 \\ 0 & X - a_2 & * & \dots & 0 \\ 0 & 0 & X - a_3 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & X - a_n \end{bmatrix} = \det (X \cdot \text{id}_n - T).$$

The latter definition is what you'll see in most linear algebra books because it lets you define the characteristic polynomial without having to say anything, and more importantly it lets you construct the characteristic polynomial even if you use the evil “array of numbers” way of teaching linear algebra. I'll admit it does have the merit that it means that given any matrix, it's easy to compute the characteristic polynomial and hence compute the eigenvalues. But what a *terrible* way to define the characteristic polynomial.

In particular, it obscures the following theorem, which is actually a completely triviality.

Theorem 8.7.1 (Cayley-Hamilton theorem)

For any $T : V \rightarrow V$, the map $p_T(T)$ is the zero map.

Here, by $p_T(T)$ we mean that if

$$p_T(X) = X^n + c_{n-1}X^{n-1} + \dots + c_0$$

then

$$p_T(T) = T^n + c_{n-1}T^{n-1} + \dots + c_1T + c_0I$$

is the zero map, where T^k denotes T applied k times.

Example 8.7.2 (Example of Cayley-Hamilton, immorally)

Suppose $T = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$. Using the immoral definition of characteristic polynomial, we find that $p_T(X) = (X - 1)(X - 4) - (-2)(-3) = X^2 - 5X - 2$. Indeed, you can verify that

$$T^2 - 5T - 2 = \begin{bmatrix} 7 & 10 \\ 15 & 22 \end{bmatrix} - 5 \cdot \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} - 2 \cdot \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}.$$

If you are only told the immoral definition of p_T and the evil view that matrices are arrays of numbers, then this looks like a complete miracle. Indeed, just look at the terrible proofs on Wikipedia.

But if you use the abstract viewpoint of T as a map, then it is a total tautology:

Proof of Cayley-Hamilton. Suppose we write V in Jordan normal form as

$$V = J_1 \oplus \cdots \oplus J_m$$

where J_i has eigenvalue λ_i and dimension d_i . By definition,

$$p_T(X) = (X - \lambda_1)^{d_1} (X - \lambda_2)^{d_2} \cdots (X - \lambda_m)^{d_m}.$$

By definition, $(X - \lambda_1)^{d_1}$ is the zero map on J_1 . So T is zero on J_1 . Similarly it's zero on each of the other J_i 's — end of story. $\square$

Also, a quick fun fact: all of this actually gives us another way to show the trace is independent from the actual matrix, as follows.

Exercise 8.7.3. Let $T : V \rightarrow V$, and suppose $p_T(X) = X^n + c_{n-1}X^{n-1} + \cdots + c_0$. Show that in any basis, the sum of the diagonal entries of T is $-c_{n-1}$. Moreover, $\det T = (-1)^n c_0$ is the product of the eigenvalues.

The fact that $\det T$ is the *product* of the eigenvalues should actually not be surprising in retrospect. Suppose $T : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ has a basis of eigenvectors e_1 and e_2 with eigenvalues 4 and 5; hence $T(e_1) = 4e_1$ and $T(e_2) = 5e_2$, and so of course we would say that T stretches areas by a factor of 20!

§8.8 (Digression) Tensoring up

Earlier I said that we'll generally just work over algebraically closed fields, because, for example, a real matrix can be interpreted as a complex matrix whose entries just happen to be real numbers. I just want to briefly tell you, in terms of tensor products, exactly what we have done.

Let's take the space $V = \mathbb{R}^3$, with basis e_1, e_2, e_3 . Thus objects in V are of the form

$$r_1 e_1 + r_2 e_2 + r_3 e_3$$

where r_1, r_2, r_3 are real numbers. We want to consider essentially the same vector space, but with complex coefficients z_i rather than real coefficients e_i .

So here's what we do: view $\mathbb{C}$ as a $\mathbb{R}$ -vector space (with basis $\{1, i\}$, say) and consider the **complexification**

$$V_{\mathbb{C}} \stackrel{\text{def}}{=} \mathbb{C} \otimes_{\mathbb{R}} V.$$

Then you can check that our elements are actually of the form

$$z_1 \otimes e_1 + z_2 \otimes e_2 + z_3 \otimes e_3.$$

(Here, the tensor product is over $\mathbb{R}$, so we have $z \otimes re_i = (zr) \otimes e_i$ for $r \in \mathbb{R}$.) Then $V_{\mathbb{C}}$ can be thought as a vector space over $\mathbb{C}$, since you can multiply on the left by complex numbers. In this way, the tensor product lets us “fuse on” complex coefficients even though technically we shouldn't be able to do so.

If $T : V \rightarrow W$ is a map, then $T_{\mathbb{C}} : V_{\mathbb{C}} \rightarrow W_{\mathbb{C}}$ is just the map $z \otimes v \mapsto z \otimes T(v)$. You'll see this written sometimes as $T_{\mathbb{C}} = \text{id} \otimes T$.

§8.9 Problems to think about

Problem 8A. For a matrix X , we define the exponential map by

$$\exp(X) = 1 + X + \frac{X^2}{2!} + \frac{X^3}{3!} + \dots$$

(take it for granted that this converges to a matrix). Prove that

$$\det(\exp(X)) = e^{\operatorname{Tr} X}.$$



Problem 8B (Putnam 1988). Let V be an n -dimensional vector space. Let $T : V \rightarrow V$ be a linear map and suppose there exists $n + 1$ eigenvectors, any n of which are linearly independent. Does it follow that T is a scalar multiple of the identity?



Problem 8C (Putnam 2015). Define S to be the set of real matrices $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ such that a, b, c, d form an arithmetic progression in that order. Find all $M \in S$ such that for some integer $k > 1$, $M^k \in S$.

DRAFT (Evan Chen)
Updated February 15, 2017

9 Inner product spaces

It will often turn out that our vector spaces which look more like $\mathbb{R}^n$ not only have the notion of addition, but also a notion of *orthogonality* and the notion of *distance*. All this is achieved by endowing the vector space with a so-called **inner form**, which you likely already know as the “dot product” for $\mathbb{R}^n$. Indeed, in $\mathbb{R}^n$ you already know that

- $v \cdot w = 0$ if and only if v and w are perpendicular, and
- $|v|^2 = v \cdot v$.

The purpose is to quickly set up this structure in full generality. Some highlights of the chapter:

- We’ll see that the high school “dot product” formulation is actually very natural: it falls out from the two axioms we listed above. If you ever wondered why $\sum a_i b_i$ behaves as nicely as it does, now you’ll know.
- We show how the inner form can be used to make V into a *metric space*, giving it more geometric structure.
- We’ll identify $V \cong V^\vee$ in a way that wasn’t possible before, and as a corollary deduce the nice result that symmetric matrices with real entries always have real eigenvalues.

Throughout this chapter, *all vector spaces are over $\mathbb{C}$ or $\mathbb{R}$* , unless otherwise specified. We’ll generally prefer working over $\mathbb{C}$ instead of $\mathbb{R}$ since, like we saw in the chapter on spectral theory, $\mathbb{C}$ is algebraically closed in a nice way that $\mathbb{R}$ is not. Every real matrix can be thought of as a matrix with complex entries anyways.

§9.1 Defining the norm

Prototypical example for this section: Dot product in $\mathbb{R}^n$.

First, let’s define the inner form for real spaces. Rather than the notation $v \cdot w$ it is most customary to use $\langle v, w \rangle$ for general vector spaces.

Definition 9.1.1. Let V be a real vector space. A **real inner form**¹ is a function

$$\langle \bullet, \bullet \rangle : V \times V \rightarrow \mathbb{R}$$

which satisfies the following properties:

- The form is **symmetric**: for any $v, w \in V$ we have

$$\langle v, w \rangle = \langle w, v \rangle.$$

Of course, one would expect this property from a product.

¹Other names include “inner product”, “dot product”, “positive definite symmetric bilinear form”, ...

- The form is **linear in the first argument**, meaning that $\langle cx, y \rangle = c \langle x, y \rangle$ and $\langle x + y, z \rangle = \langle x, z \rangle + \langle y, z \rangle$. This is often summarized by the single equation

$$\langle cx + y, z \rangle = c \langle x, z \rangle + \langle y, z \rangle.$$

Of course by symmetry it is linear in the second argument too. This should be thought of as “distributive over addition”.

- The form is **positive definite**, meaning $\langle v, v \rangle \geq 0$ is a nonnegative real number, and equality takes place only if $v = 0_V$.

Example 9.1.2 ($\mathbb{R}^n$)

As we already know, one can define the inner form on $\mathbb{R}^n$ as follows. Let $\mathbf{e}_1 = (1, 0, \dots, 0)$, $\mathbf{e}_2 = (0, 1, \dots, 0)$, $\dots$, $\mathbf{e}_n = (0, \dots, 0, 1)$ be the standard basis. Then we let

$$\langle a_1 \mathbf{e}_1 + \dots + a_n \mathbf{e}_n, b_1 \mathbf{e}_1 + \dots + b_n \mathbf{e}_n \rangle \stackrel{\text{def}}{=} a_1 b_1 + \dots + a_n b_n.$$

It’s easy to see this is bilinear (symmetric and linear in both arguments). To see it is positive definite, note that if $a_i = b_i$ then the dot product is $a_1^2 + \dots + a_n^2$, which is zero exactly when all a_i are zero.

The definition for a complex product space is similar, but has a minor difference: rather than symmetry we instead have *conjugate symmetry* meaning $\langle v, w \rangle = \overline{\langle w, v \rangle}$. Thus, while we still have linearity in the first argument, we actually have a different linearity for the second argument. To be explicit:

Definition 9.1.3. Let V be a complex vector space. A **complex inner form** $\langle \bullet, \bullet \rangle : V \times V \rightarrow \mathbb{C}$ is a function which satisfies the following properties:

- The form has **conjugate symmetry**, which means that for any $v, w \in V$ we have

$$\langle v, w \rangle = \overline{\langle w, v \rangle}.$$

- The form is **linear in the first argument**, so again we have

$$\langle cx + y, z \rangle = c \langle x, z \rangle + \langle y, z \rangle.$$

However it no longer follows that it is linear in the second argument, see below!

- The form is **positive definite**, meaning $\langle v, v \rangle$ is a nonnegative real number, and equals zero exactly when $v = 0_V$.

Question 9.1.4. Show that in a complex inner form, we instead have

$$\langle x, cy + z \rangle = \bar{c} \langle x, y \rangle + \langle x, z \rangle.$$

Example 9.1.5 ($\mathbb{C}^n$)

The dot product in $\mathbb{C}^n$ is defined as follows: let $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n$ be the standard basis. For complex numbers w_i, z_i we set

$$\langle w_1 \mathbf{e}_1 + \dots + w_n \mathbf{e}_n, z_1 \mathbf{e}_1 + \dots + z_n \mathbf{e}_n \rangle \stackrel{\text{def}}{=} w_1 \bar{z}_1 + \dots + w_n \bar{z}_n.$$

Question 9.1.6. Check that the above is in fact a complex inner form.

Definition 9.1.7. A real or complex vector space with an inner form is called an **inner product space**.

Remark 9.1.8. Note that we can apply the axioms for a complex inner form to those for a real inner form verbatim; in $\mathbb{R}$, conjugation is the identity. Thus for real vector spaces, for example, “conjugate symmetry” is just “symmetry”. Consequently, when proving results about inner product spaces it usually suffices to do just the complex case.

The above example explains why we have to require the complex inner form to satisfy conjugate symmetry rather than just symmetry. If we had tried to define the dot product as $\sum w_i z_i$, then we would have lost the condition of being positive definite, because there is no guarantee that $\langle v, v \rangle = \sum z_i^2$ will even be a real number at all. On the other hand, with conjugate symmetry we actually enforce $\langle v, v \rangle = \overline{\langle v, v \rangle}$, i.e. $\langle v, v \rangle \in \mathbb{R}$ for every v .

Now that we have a dot product, we can talk both about the norm and orthogonality.

§9.2 Norms

Prototypical example for this section: $\mathbb{R}^n$ becomes its usual Euclidean space with the vector norm.

The inner form equips our vector space with a notion of distance, which we call the norm.

Definition 9.2.1. Let V be an inner product space. The **norm** of $v \in V$ is defined by

$$\|v\| = \sqrt{\langle v, v \rangle}.$$

This definition makes sense because we assumed our form to be positive definite.

Example 9.2.2 ($\mathbb{R}^n$ and $\mathbb{C}^n$ are normed vector spaces)

When $V = \mathbb{R}^n$ or $V = \mathbb{C}^n$ with the standard dot product norm, then the norm of v corresponds to the absolute value that we are used to.

Our goal now is to prove that

With the metric $d(v, w) = \|v - w\|$, V becomes a metric space.

To do this we have to establish the triangle inequality. We use the following lemma as a stepping stone: Let’s now prove something we all know and love:

Lemma 9.2.3 (Cauchy-Schwarz)

Let V be an inner product space. For any $v, w \in V$ we have

$$|\langle v, w \rangle| \leq \|v\| \|w\|$$

with equality if and only if v and w are linearly dependent.

Proof. The key to the proof is to think about the equality case: we'll use the inequality $\langle cv - w, cv - w \rangle \geq 0$. Deferring the choice of c until later, we compute

$$\begin{aligned} 0 &\leq \langle cv - w, cv - w \rangle \\ &= \langle cv, cv \rangle - \langle cv, w \rangle - \langle w, cv \rangle + \langle w, w \rangle \\ &= |c|^2 \langle v, v \rangle - c \langle v, w \rangle - \bar{c} \langle w, v \rangle + \langle w, w \rangle \\ &= |c|^2 \|v\|^2 + \|w\|^2 - c \langle v, w \rangle - \overline{c \langle v, w \rangle} \\ 2 \operatorname{Re} [c \langle v, w \rangle] &\leq |c|^2 \|v\|^2 + \|w\|^2 \end{aligned}$$

At this point, a good choice of c is

$$c = \frac{\|w\|}{\|v\|} \cdot \frac{|\langle v, w \rangle|}{\langle v, w \rangle}.$$

since then

$$\begin{aligned} c \langle v, w \rangle &= \frac{\|w\|}{\|v\|} |\langle v, w \rangle| \in \mathbb{R} \\ |c| &= \frac{\|w\|}{\|v\|} \end{aligned}$$

whence the inequality becomes

$$\begin{aligned} 2 \frac{\|w\|}{\|v\|} |\langle v, w \rangle| &\leq 2 \|w\|^2 \\ |\langle v, w \rangle| &\leq \|v\| \|w\|. \end{aligned}$$

□

Thus:

Theorem 9.2.4 (Triangle inequality)

We always have

$$\|v\| + \|w\| \geq \|v + w\|$$

with equality if and only if v and w are linearly dependent.

Exercise 9.2.5. Prove this by squaring both sides, and applying Cauchy-Schwarz.

In this way, our vector space now has a topological structure of a metric space.

§9.3 Orthogonality

Prototypical example for this section: Still $\mathbb{R}^n$!

Our next goal is to give the geometric notion of “perpendicular”. The definition is easy enough:

Definition 9.3.1. Two nonzero vectors v and w in an inner product space are **orthogonal** if $\langle v, w \rangle = 0$.

As we expect from our geometric intuition in $\mathbb{R}^n$, this implies independence:

Lemma 9.3.2 (Orthogonal vectors are independent)

Any set of nonzero pairwise orthogonal vectors is linearly independent.

Proof. Consider a dependence

$$a_1 v_1 + \cdots + a_n v_n = 0$$

for a_i in $\mathbb{R}$ or $\mathbb{C}$. Then

$$0_V = \left\langle v_1, \sum a_i v_i \right\rangle = a_1 \|v_1\|^2.$$

Hence $a_1 = 0$, since we assumed $\|v_1\| \neq 0$. Similarly $a_2 = \cdots = a_n = 0$. $\square$

In light of this, we can now consider a stronger condition on our bases:

Definition 9.3.3. An **orthonormal** basis of a finite-dimensional inner product space V is a basis $e_1, \dots, e_n$ such that $\|e_i\| = 1$ for every i and $\langle e_i, e_j \rangle = 0$ for any $i \neq j$.

Example 9.3.4 ($\mathbb{R}^n$ and $\mathbb{C}^n$ have standard bases)

In $\mathbb{R}^n$ and $\mathbb{C}^n$ equipped with the standard dot product, the standard basis $\mathbf{e}_1, \dots, \mathbf{e}_n$ is also orthonormal.

This is no loss of generality:

Theorem 9.3.5 (Gram-Schmidt)

Let V be a finite-dimensional inner product space. Then it has an orthonormal basis.

Sketch of Proof. One constructs the orthonormal basis explicitly from any basis $e_1, \dots, e_n$ of V . Define $\text{proj}_u(v) = \frac{\langle v, u \rangle}{\langle u, u \rangle} u$. Then recursively define

$$\begin{aligned} u_1 &= e_1 \\ u_2 &= e_2 - \text{proj}_{u_1}(e_2) \\ u_3 &= e_3 - \text{proj}_{u_1}(e_3) - \text{proj}_{u_2}(e_3) \\ &\vdots \\ u_n &= e_n - \text{proj}_{u_1}(e_n) - \cdots - \text{proj}_{u_{n-1}}(e_n). \end{aligned}$$

One can show the u_i are pairwise orthogonal and not zero. $\square$

Thus, we can generally assume our bases are orthonormal.

Worth remarking:

Example 9.3.6 (The dot product is the “only” inner form)

Let V be a finite-dimensional inner product space, and consider *any* orthonormal basis $e_1, \dots, e_n$. Then we have that

$$\langle a_1 e_1 + \cdots + a_n e_n, b_1 e_1 + \cdots + b_n e_n \rangle = \sum_{i,j=1}^n a_i \overline{b_j} \langle e_i, e_j \rangle = \sum_{i=1}^n a_i \overline{b_i}$$

owing to the fact that the $\{e_i\}$ are orthonormal.

So the dot product is actually a *highly natural* inner form to consider. It arises just out of the geometric consideration that in an orthonormal basis we ought to have norm 1 and pairwise orthogonality; this is enough to obtain the dot product by extending linearly.

§9.4 Identifying with the dual space

Earlier I complained that there was no natural isomorphism $V \cong V^\vee$. But in fact, given an inner form we can actually make such an identification: that is we can naturally associate every linear map $\xi : V \rightarrow k$ with a vector $v \in V$.

To see how we might do this, suppose $V = \mathbb{R}^3$ for now with standard basis $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$. How might we use the standard product to represent a map from $V \rightarrow \mathbb{R}$? For example, take $\xi \in V^\vee$ by $\xi(\mathbf{e}_1) = 3, \xi(\mathbf{e}_2) = 4$ and $\xi(\mathbf{e}_3) = 5$. Actually, I claim that

$$\xi(v) = \langle 3\mathbf{e}_1 + 4\mathbf{e}_2 + 5\mathbf{e}_3, v \rangle$$

for every v . And this is obvious, once you do the computation.

More generally:

Theorem 9.4.1 (V is isomorphic to $V^\vee$)

Let V be a finite-dimensional inner product space and $V^\vee$ its dual. Then the map $V \rightarrow V^\vee$ by

$$v \mapsto \langle v, - \rangle \in V^\vee$$

is an isomorphism.

Proof. It suffices to show that the map is injective and surjective.

- Injective: suppose $\langle v_1, w \rangle = \langle v_2, w \rangle$ for every vector w . This means $\langle v_1 - v_2, w \rangle = 0$ for every vector $w \in V$. This can only happen if $v_1 - v_2 = 0$; for example, take $w = v_1 - v_2$ and use positive definiteness.
- Surjective: take an orthonormal basis, and mimic the argument we gave earlier.

Actually, since we already know $\dim V = \dim V^\vee$ we only had to prove one of the above. As a matter of personal taste, I find the proof of injectivity more elegant, and the proof of surjectivity more enlightening, so I included both. $\square$

Thus

Once V is given an inner form, V and $V^\vee$ are canonically isomorphic.

§9.5 The transpose of a matrix

Let V be a finite-dimensional inner product space and let $T : V \rightarrow V$. The **adjoint** (soon **conjugate transpose**) of T , denoted $T^\dagger$, is defined as follows: for every vector $w \in V$, we let $T^\dagger(w)$ be the unique vector with

$$\langle v, T^\dagger(w) \rangle = \langle T(v), w \rangle$$

for every $v \in V$. This is well-defined, because $v \mapsto \langle T(v), w \rangle$ is some function in $V^\vee$, and hence by the isomorphism we described earlier it can be written uniquely in the form $\langle v, x \rangle$ for some $x \in V$; we then let $T^\dagger(w) = x$.

By symmetry, of course, we also have $\langle T^\dagger(v), w \rangle = \langle v, T(w) \rangle$.

Of course, the name conjugate transpose suggests something else:

Theorem 9.5.1 (Adjoint is conjugate transpose)

Fix an orthonormal basis of a finite-dimensional inner product V . If we write T as a matrix in this basis, then $T^\dagger$ is the conjugate transpose.

Proof. One-line version: take v and w to be basis elements, and this falls right out.

Full proof: let

$$T = \begin{bmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{n1} & \cdots & a_{nn} \end{bmatrix}$$

in this basis $e_1, \dots, e_n$. Then, letting $w = e_i$ and $v = e_j$ we deduce that

$$\langle e_i, T^\dagger(e_j) \rangle = \langle T(e_i), e_j \rangle = a_{ji} \implies \langle T^\dagger(e_j), e_i \rangle = \overline{a_{ji}}$$

for any i , which is enough to deduce the result. $\square$

In this way we can talk about the transpose of a matrix in a meaningful way.

§9.6 Spectral theory of normal maps

The world would be a very beautiful place if it turned out that we could pick a basis of eigenvectors that was also orthonormal. This is of course far too much to hope for; even without the orthonormal condition, we saw in the previous chapter that we could still have 1's off the diagonal. However, it turns out that there is a good characterization anyways.

Definition 9.6.1. We say T is **normal** if $TT^\dagger = T^\dagger T$.

We say a complex T is **self-adjoint** or **Hermitian** if $T = T^\dagger$; i.e. as a matrix in any orthonormal basis, T is its own conjugate transpose. For real T we say “self-adjoint”, “Hermitian” or **symmetric**.

Theorem 9.6.2 (Normal $\iff$ diagonalizable in inner form)

Let V be a finite-dimensional complex inner product space. A linear map $T : V \rightarrow V$ is normal if and only if one can pick an orthonormal basis of eigenvectors.

Proof. This is long, and maybe should be omitted on a first reading. If T has an orthonormal basis of eigenvectors, this result is immediate.

Now assume T is normal. We first prove T is diagonalizable; this is the hard part.

Claim 9.6.3. If T is normal, then $\ker T = \ker T^r = \ker T^\dagger$ for $r \geq 1$. (Here T^r is T applied r times.)

Proof of Claim. Let $S = T^\dagger \circ T$, which is self-adjoint. We first note that S is Hermitian and $\ker S = \ker T$. To see it's Hermitian, note $\langle Sv, w \rangle = \langle Tv, Tw \rangle = \langle v, Sw \rangle$. Taking $v = w$ also implies $\ker S \subseteq \ker T$ (and hence equality since obviously $\ker T \subseteq \ker S$).

First, since we have $\langle S^r(v), S^{r-2}(v) \rangle = \langle S^{r-1}(v), S^{r-1}(v) \rangle$, an induction shows that $\ker S = \ker S^r$ for $r \geq 1$. Now, since T is normal, we have $S^r = (T^\dagger)^r \circ T^r$, and thus we have the inclusion

$$\ker T \subseteq \ker T^r \subseteq \ker S^r = \ker S = \ker T$$

where the last equality follows from the first claim. Thus in fact $\ker T = \ker T^r$.

Finally, to show equality with $\ker T^\dagger$ we

$$\begin{aligned}\langle Tv, Tv \rangle &= \langle v, T^\dagger Tv \rangle \\ &= \langle v, TT^\dagger v \rangle \\ &= \langle T^\dagger v, T^\dagger v \rangle.\end{aligned}$$

■

Now consider the given T , and any λ .

Question 9.6.4. Show that $(T - \lambda \text{id})^\dagger = T^\dagger - \bar{\lambda} \text{id}$. Thus if T is normal, so is $T - \lambda \text{id}$.

In particular, for any eigenvalue λ of T , we find that $\ker(T - \lambda \text{id}) = \ker(T - \lambda \text{id})^r$. This implies that all the Jordan blocks of T have size 1; i.e. that T is in fact diagonalizable. Finally, we conclude that the eigenvectors of T and $T^\dagger$ match, and the eigenvalues are complex conjugates.

So, diagonalize T . We just need to show that if v and w are eigenvectors of T with distinct eigenvalues, then they are orthogonal. (We can use Gram-Schmidt on any eigenvalue that appears multiple times.) To do this, suppose $T(v) = \lambda v$ and $T(w) = \mu w$ (thus $T^\dagger(w) = \bar{\mu}w$). Then

$$\lambda \langle v, w \rangle = \langle \lambda v, w \rangle = \langle Tv, w \rangle = \langle v, T^\dagger(w) \rangle = \langle v, \bar{\mu}w \rangle = \bar{\mu} \langle v, w \rangle.$$

Since $\lambda \neq \bar{\mu}$, we conclude $\langle v, w \rangle = 0$. □

This means that not only can we write

$$T = \begin{bmatrix} \lambda_1 & \dots & \dots & 0 \\ 0 & \lambda_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \lambda_n \end{bmatrix}$$

but moreover that the basis associated with this matrix happens to be orthonormal vectors.

As a corollary:

Theorem 9.6.5 (Hermitian matrices have real eigenvalues)

A Hermitian matrix T is diagonalizable, and all its eigenvalues are real.

Proof. Obviously Hermitian $\implies$ normal, so write it in the orthonormal basis of eigenvectors. To see that the eigenvalues are real, note that $T = T^\dagger$ means $\lambda_i = \bar{\lambda}_i$ for every i . □

§9.7 Problems to think about

Problem 9A (Pythagorean theorem). Show that if $\langle v, w \rangle = 0$ in an inner product space, then $\|v\|^2 + \|w\|^2 = \|v + w\|^2$.



Problem 9B (Taiwan IMO camp). In a town there are n people and k clubs. Each club has an odd number of members, and any two clubs have an even number of common members. Prove that $k \leq n$.

Problem 9C* (Inner product structure of tensors). Let V and W be finite-dimensional inner product spaces over k , where k is either $\mathbb{R}$ or $\mathbb{C}$.

- (a) Find a canonical way to make $V \otimes_k W$ into an inner product space too.
- (b) Let $e_1, \dots, e_n$ be an orthonormal basis of V and $f_1, \dots, f_m$ be an orthonormal basis of W . What's an orthonormal basis of $V \otimes W$?



Problem 9D (Putnam 2014). Let n be a positive integer. What is the largest k for which there exist $n \times n$ matrices $M_1, \dots, M_k$ and $N_1, \dots, N_k$ with real entries such that for all i and j , the matrix product $M_i N_j$ has a zero entry somewhere on its diagonal if and only if $i \neq j$?

DRAFT (Evan Chen)
Updated February 15, 2017

III

Groups, Rings, and More

10	Group actions overkill AIME problems	122
10.1	Definition of a group action	122
10.2	Stabilizers and orbits	123
10.3	Burnside's lemma	124
10.4	Conjugation of elements	125
10.5	Problems to think about	127
11	Find all groups	128
11.1	Sylow theorems	128
11.2	(Optional) Proving Sylow's theorem	129
11.3	(Optional) Simple groups and Jordan-Hölder	131
11.4	Problems to think about	132
12	Rings and ideals	134
12.1	Number theory motivation	134
12.2	Definition and examples of rings	134
12.3	Integral domains and fields	136
12.4	Ideals	137
12.5	Generating ideals	139
12.6	Principal ideal domains and Noetherian rings	140
12.7	Prime ideals	142
12.8	Maximal ideals	143
12.9	Problems	144
13	The PID structure theorem	145
13.1	Finitely generated abelian groups	145
13.2	Some ring theory prerequisites	146
13.3	The structure theorem	147
13.4	Reduction to maps of free R -modules	148
13.5	Smith normal form	149
13.6	Problems to think about	151

10 Group actions overkill AIME problems

Consider this problem from the 1996 AIME:

(AIME 1996) Two of the squares of a 7×7 checkerboard are painted yellow, and the rest are painted green. Two color schemes are equivalent if one can be obtained from the other by applying a rotation in the plane of the board. How many inequivalent color schemes are possible?

What's happening here? Let X be the set of the $\binom{49}{2}$ possible colorings of the board. What's the natural interpretation of "rotation"? Answer: the group $\mathbb{Z}_4 = \langle r \mid r^4 = 1 \rangle$ somehow "acts" on this set X by sending one state $x \in X$ to another state $r \cdot x$, which is just x rotated by 90° . Intuitively we're just saying that two configurations are the same if they can be reached from one another by this "action".

We can make all of this precise using the idea of a group action.

§10.1 Definition of a group action

Prototypical example for this section: The AIME problem.

Definition 10.1.1. Let X be a set and G a group. A **group action** is a binary operation $\cdot : G \times X \rightarrow X$ which lets a $g \in G$ send an $x \in X$ to $g \cdot x$. It satisfies the axioms

- $(g_1 g_2) \cdot x = g_1 \cdot (g_2 \cdot x)$ for any $g_1, g_2 \in G$ for all $x \in X$.
- $1_G \cdot x = x$ for any $x \in X$.

Here are some more examples of group actions.

Example 10.1.2 (Examples of group actions)

Let $G = (G, \star)$ be a group. Here are some examples of group actions.

- The group $\mathbb{Z}_4$ can act on the set of ways to color a 7×7 board either yellow or green.
- The group $\mathbb{Z}_4 = \langle r \mid r^4 = 1 \rangle$ acts on the xy -plane $\mathbb{R}^2$ as follows: $r \cdot (x, y) = (y, -x)$. In other words, it's a rotation by 90° .
- The dihedral group D_{2n} acts on the set of ways to color the vertices of an n -gon.
- The group S_n acts on $X = \{1, 2, \dots, n\}$ by applying the permutation σ : $\sigma \cdot x \stackrel{\text{def}}{=} \sigma(x)$.
- The group G can act on itself (i.e. $X = G$) by left multiplication: put $g \cdot g' \stackrel{\text{def}}{=} g \star g'$.

§10.2 Stabilizers and orbits


Prototypical example for this section: Again the AIME problem.

Given a group action G on X , we can define an equivalence relation $\sim$ on X as follows: $x \sim y$ if $x = g \cdot y$ for some $g \in G$. For example, in the AIME problem, $\sim$ means “one can be obtained from the other by a rotation”.

Question 10.2.1. Why is this an equivalence relation?

In that case, the AIME problem wants the number of equivalence classes under $\sim$. So let’s give these equivalence classes a name: **orbits**. We usually denote orbits by $\mathcal{O}$.

As usual, orbits carve out X into equivalence classes.



It turns out that a very closely related concept is:

Definition 10.2.2. The **stabilizer** of a point $x \in X$, denoted $\text{Stab}_G(x)$, is the set of $g \in G$ which fix x ; in other words

$$\text{Stab}_G(x) \stackrel{\text{def}}{=} \{g \in G \mid g \cdot x = x\}.$$

Example 10.2.3

Consider the AIME problem again, with X the possible set of states (again $G = \mathbb{Z}_4$). Let x be the configuration where two opposite corners are colored yellow. Evidently 1_G fixes x , but so does the 180° rotation r^2 . But r and r^3 do not preserve x , so $\text{Stab}_G(x) = \{1, r^2\} \cong \mathbb{Z}_2$.

Question 10.2.4. Why is $\text{Stab}_G(x)$ a subgroup of G ?

Once we realize the stabilizer is a group, this leads us to what I privately call the “fundamental theorem of how big an orbit is”.


Theorem 10.2.5 (Orbit-stabilizer theorem)

Let $\mathcal{O}$ be an orbit, and pick any $x \in \mathcal{O}$. Let $S = \text{Stab}_G(x)$ be a subgroup of G . There is a natural bijection between $\mathcal{O}$ and left cosets. In particular,

$$|\mathcal{O}| |S| = |G|.$$

In particular, the stabilizers of each $x \in \mathcal{O}$ have the same size.

Proof. The point is that every coset gS just specifies an element of $\mathcal{O}$, namely $g \cdot x$. The fact that S is a stabilizer implies that it is irrelevant which representative we pick.



Since the $|\mathcal{O}|$ cosets partition G , each of size $|S|$, we obtain the second result. $\square$

§10.3 Burnside's lemma

Now for the crux of this chapter: a way to count the number of orbits.

Theorem 10.3.1 (Burnside's lemma)

Let G act on a set X . The number of orbits of the action is equal to

$$\frac{1}{|G|} \sum_{g \in G} |\text{FixPt } g|$$

where $\text{FixPt } g$ is the set of points $x \in X$ such that $g \cdot x = x$.

The proof is deferred as a bonus problem, since it has a very olympiad-flavored solution. As usual, this lemma was not actually proven by Burnside; Cauchy got there first, and thus it is sometimes called *the lemma that is not Burnside's*. Example application:

Example 10.3.2 (AIME 1996)

Two of the squares of a 7×7 checkerboard are painted yellow, and the rest are painted green. Two color schemes are equivalent if one can be obtained from the other by applying a rotation in the plane of the board. How many inequivalent color schemes are possible?

We know that $G = \mathbb{Z}_4$ acts on the set X of $\binom{49}{2}$ possible coloring schemes. Now we can compute $\text{FixPt } g$ explicitly for each $g \in \mathbb{Z}_4$.

- If $g = 1_G$, then every coloring is fixed, for a count of $\binom{49}{2} = 1176$.
- If $g = r^2$ there are exactly 24 coloring schemes fixed by g : this occurs when the two squares are reflections across the center, which means they are preserved under a 180° rotation.
- If $g = r$ or $g = r^3$, then there are no fixed coloring schemes.

As $|G| = 4$, the average is

$$\frac{1176 + 24 + 0 + 0}{4} = 300.$$

Exercise 10.3.3 (MathCounts Chapter Target Round). A circular spinner has seven sections of equal size, each of which is colored either red or blue. Two colorings are considered the same if one can be rotated to yield the other. In how many ways can the spinner be colored? (Answer: 20)

Consult [Ma13b] for some more examples of “hands-on” applications.

§10.4 Conjugation of elements

Prototypical example for this section: In S_n , conjugacy classes are “cycle types”.

A particularly common type of action is the so-called **conjugation**. We let G act on itself as follows:

$$g : h \mapsto ghg^{-1}.$$

You might think this definition is a little artificial. Who cares about the element ghg^{-1} ? Let me try to convince you this definition is not so unnatural.

Example 10.4.1 (Conjugacy in S_n)

Let $G = S_5$, and fix a $\pi \in S_5$. Here’s the question: is $\pi\sigma\pi^{-1}$ related to σ ? To illustrate this, I’ll write out a completely random example of a permutation $\sigma \in S_5$.

$$\begin{array}{rclcl} & 1 & \mapsto & 3 & \pi(1) & \mapsto & \pi(3) \\ & 2 & \mapsto & 1 & \pi(2) & \mapsto & \pi(1) \\ \text{If } \sigma = & 3 & \mapsto & 5 & \text{then } \pi\sigma\pi^{-1} = & \pi(3) & \mapsto & \pi(5) \\ & 4 & \mapsto & 2 & & \pi(4) & \mapsto & \pi(2) \\ & 5 & \mapsto & 4 & & \pi(5) & \mapsto & \pi(4) \end{array}$$

Thus our fixed π doesn’t really change the structure of σ at all: it just “renames” each of the elements 1, 2, 3, 4, 5 to $\pi(1)$, $\pi(2)$, $\pi(3)$, $\pi(4)$, $\pi(5)$.

But wait, you say. That’s just a very particular type of group behaving nicely under conjugation. Why does this mean anything more generally? All I have to say is: remember Cayley’s theorem! (This was Problem 1E†.)

In any case, we may now define:

Definition 10.4.2. The **conjugacy classes** of a group G are the orbits of G under the conjugacy action.

Let’s see what the conjugacy classes of S_n are, for example.

Example 10.4.3 (Conjugacy classes of S_n correspond to cycle types)

Intuitively, the discussion above says that two elements of S_n should be conjugate if they have the same “shape”, regardless of what the elements are named. The right way to make the notion of “shape” rigorous is cycle notation. For example, consider the permutation

$$\sigma_1 = (1\ 3\ 5)(2\ 4)$$

in cycle notation, meaning $1 \mapsto 3 \mapsto 5 \mapsto 1$ and $2 \mapsto 4 \mapsto 2$. It is conjugate to the permutation

$$\sigma_2 = (1\ 2\ 3)(4\ 5)$$

or any other way of relabeling the elements. So, we could think of σ as having conjugacy class

$$(-\ -\ -)(-\ -).$$

More generally, you can show that two elements of S_n are conjugate if and only if they have the same “shape” under cycle decomposition.

Question 10.4.4. Show that the number of conjugacy classes of S_n equals the number of partitions of n .



Figure 10.1: An element in a conjugacy class of size one. (Source: [Ge].)

Question 10.4.5. Figure out why the claim in Figure 10.1 is true.

As long as I’ve put the above picture, I may as well also define:

Definition 10.4.6. Let G be a group. The **center** of G , denoted $Z(G)$, is the set of elements $x \in G$ such that $xg = gx$ for every $g \in G$. More succinctly,

$$Z(G) \stackrel{\text{def}}{=} \{x \in G \mid gx = xg \ \forall g \in G\}.$$

You can check this is indeed a subgroup of G .

Question 10.4.7. Why is $Z(G)$ normal in G ?

Question 10.4.8. What are the conjugacy classes of elements in the center?

A trivial result that gets used enough that I should explicitly call it out:

Corollary 10.4.9 (Conjugacy in abelian groups is trivial)

If G is abelian, then the conjugacy classes all have size one.

§10.5 Problems to think about

Problem 10A (PUMaC 2009 C8). Taotao wants to buy a bracelet consisting of seven beads, each of which is orange, white or black. (The bracelet can be rotated and reflected in space.) Find the number of possible bracelets.

Problem 10B. Show that two elements in the same conjugacy class have the same order.



Problem 10C. Prove Burnside's lemma.

Problem 10D (The “class equation”). Let G be a finite group. We define the **centralizer** $C_G(g) = \{x \in G \mid xg = gx\}$ for each $g \in G$. Show that

$$|G| = |Z(G)| + \sum_{s \in S} \frac{|G|}{|C_G(s)|}$$

where $S \subseteq G$ is defined as follows: for each conjugacy class $C \subseteq G$ with $|C| > 1$, we pick a representative of C and add it to S .



Problem 10E[†] (Classical). Assume G is a finite group and p is the smallest prime dividing its order. Let H be a subgroup of G with $|G|/|H| = p$. Show that H is normal in G .

11 Find all groups

The following problem will hopefully never be proposed at the IMO.

Let n be a positive integer and let $S = \{1, \dots, n\}$. Find all functions $f : S \times S \rightarrow S$ such that

- (a) $f(x, 1) = f(1, x) = x$ for all $x \in S$.
- (b) $f(f(x, y), z) = f(x, f(y, z))$ for all $x, y, z \in S$.
- (c) For every $x \in S$ there exists a $y \in S$ such that $f(x, y) = f(y, x) = 1$.

Nonetheless, it's remarkable how much progress we've made on this "problem". In this chapter I'll try to talk about some things we have accomplished.

§11.1 Sylow theorems

Here we present the famous Sylow theorems, some of the most general results we have about finite groups.

Theorem 11.1.1 (The Sylow theorems)

Let G be a group of order $p^n m$, where $\gcd(p, m) = 1$ and p is a prime. A **Sylow p -subgroup** is a subgroup of order p^n . Let n_p be the number of Sylow p -subgroups of G . Then

- (a) $n_p \equiv 1 \pmod{p}$. In particular, $n_p \neq 0$ and a Sylow p -subgroup exists.
- (b) n_p divides m .
- (c) Any two Sylow p -subgroups are conjugate subgroups (hence isomorphic).

Sylow's theorem is really huge for classifying groups; in particular, the conditions $n_p \equiv 1 \pmod{p}$ and $n_p \mid m$ can often pin down the value of n_p to just a few values. Here are some results which follow from the Sylow theorems.

- A Sylow p -subgroup is normal if and only if $n_p = 1$.
- Any group G of order pq , where $p < q$ are primes, must have $n_q = 1$, since $n_q \equiv 1 \pmod{q}$ yet $n_q \mid p$. Thus G has a normal subgroup of order q .
- Since any abelian group has all subgroups normal, it follows that any abelian group has exactly one Sylow p -subgroup for every p dividing its order.
- If $p \neq q$, the intersection of a Sylow p -subgroup and a Sylow q -subgroup is just $\{1_G\}$. That's because the intersection of any two subgroups is also a subgroup, and Lagrange's theorem tells us that its order must divide both a power of p and a power of q ; this can only happen if the subgroup is trivial.

Here's an example of another "practical" application.

Proposition 11.1.2 (Triple product of primes)

If $|G| = pqr$ is the product of distinct primes, then G must have a normal Sylow subgroup.

Proof. WLOG, assume $p < q < r$. Notice that $n_p \equiv 1 \pmod{p}$, $n_p | qr$ and cyclically, and assume for contradiction that $n_p, n_q, n_r > 1$.

Since $n_r | pq$, we have $n_r = pq$ since n_r divides neither p nor q as $n_r \geq 1 + r > p, q$. Also, $n_p \geq 1 + p$ and $n_q \geq 1 + q$. So we must have at least $1 + p$ Sylow p -subgroups, at least $1 + q$ Sylow q -subgroups, and at least pq Sylow r -subgroups.

But these groups are pretty exclusive.

Question 11.1.3. Take the $n_p + n_q + n_r$ Sylow subgroups and consider two of them, say H_1 and H_2 . Show that $|H_1 \cap H_2| = 1$ as follows: check that $H_1 \cap H_2$ is a subgroup of both H_1 and H_2 , and then use Lagrange's theorem.

We claim that there are too many elements now. Indeed, if we count the non-identity elements contributed by these subgroups, we get

$$n_p(p-1) + n_q(q-1) + n_r(r-1) \geq (1+p)(p-1) + (1+q)(q-1) + pq(r-1) > pqr$$

which is more elements than $|G|$ has! $\square$

§11.2 (Optional) Proving Sylow's theorem

The proof of Sylow's theorem is somewhat involved, and in fact many proofs exist. I'll present one below here. It makes extensive use of group actions, so I want to recall a few facts first. If G acts on X , then

- The orbits of the action form a partition of X .
- if $\mathcal{O}$ is any orbit, then the orbit-Stabilizer theorem says that

$$|\mathcal{O}| = |G| / |\text{Stab}_G(x)|$$

for any $x \in \mathcal{O}$.

- In particular: suppose in the above that G is a **p -group**, meaning $|G| = p^t$ for some t . Then either $|\mathcal{O}| = 1$ or p divides $|\mathcal{O}|$. In the case $\mathcal{O} = \{x\}$, then by definition, x is a **fixed point** of every element of G : we have $g \cdot x = x$ for every x .

Note that when I say x is a fixed point, I mean it is fixed by **every** element of the group, i.e. the orbit really has size one. Hence that's a really strong condition.

Definitions

Prototypical example for this section: Conjugacy in S_n .

I've defined conjugacy of elements previously, but I now need to define it for groups:

Definition 11.2.1. Let G be a group, and let X denote the set of subgroups of G . Then **conjugation** is the action of G on X that sends

$$H \mapsto gHg^{-1} = \{ghg^{-1} \mid h \in H\}.$$

If H and K are subgroups of G such that $H = gKg^{-1}$ for some $g \in G$ (in other words, they are in the same orbit under this action), then we say they are **conjugate** subgroups.

Because we somehow don't think of conjugate elements as "that different" (for example, in permutation groups), the following shouldn't be surprising:

Question 11.2.2. Show that for any subgroup H of a group G , the map $H \rightarrow gHg^{-1}$ by $h \mapsto ghg^{-1}$ is in fact an isomorphism. This implies that any two conjugate subgroups are isomorphic.

Definition 11.2.3. For any subgroup H of G the **normalizer** of H is defined as

$$N_G(H) \stackrel{\text{def}}{=} \{g \in G \mid gHg^{-1} = H\}.$$

In other words, it is the stabilizer of H under the conjugation action.

We are now ready to present the proof.

Step 1: Prove that a Sylow p -subgroup exists

What follows is something like the probabilistic method. By considering the set X of ALL subsets of size p^n at once, we can exploit the "deep number theoretic fact" that

$$|X| = \binom{p^nm}{p^n} \not\equiv 0 \pmod{p}.$$

(It's not actually deep: use Lucas' theorem.)

Here is the proof.

- Let G act on X by $g \cdot X \stackrel{\text{def}}{=} \{gx \mid x \in X\}$.
- Take an orbit $\mathcal{O}$ with size not divisible by p . (This is possible because of our deep number theoretic fact. Since $|X|$ is nonzero mod p and the orbits partition X , the claimed orbit must exist.)
- let $S \in \mathcal{O}$, $H = \text{Stab}_G(S)$. then p^n divides $|H|$, by the orbit-Stabilizer theorem.
- Consider a second action: let H act on S by $h \cdot s \stackrel{\text{def}}{=} hs$ (we know $hs \in S$ since $H = \text{Stab}_G(S)$).
- Observe that $\text{Stab}_H(s) = \{1_H\}$. Then all orbits of the second action must have size $|H|$. Thus $|H|$ divides $|S| = p^n$.
- This implies $|H| = p^n$, and we're done.

Step 2: Any two Sylow p -subgroups are conjugate

Let P be a Sylow p -subgroup (which exists by the previous step). We now prove that for any p -group Q , $Q \subseteq gPg^{-1}$. Note that if Q is also a Sylow p -subgroup, then $Q = gPg^{-1}$ for size reasons; this implies that any two Sylow subgroups are indeed conjugate.

Let Q act on the set of left cosets of P by left multiplication. Note that

- Q is a p -group, so any orbit has size divisible by p unless it's 1.
- But the number of left cosets is m , which isn't divisible by p .

Hence some coset gP is a fixed point for every q , meaning $qgP = gP$ for all q . Equivalently, $qg \in gP$ for all $q \in Q$, so $Q \subseteq gPg^{-1}$ as desired.

Step 3: Showing $n_p \equiv 1 \pmod{p}$

Let $\mathcal{S}$ denote the set of all the Sylow p -subgroups. Let $P \in \mathcal{S}$ be arbitrary.

Question 11.2.4. Why does $|\mathcal{S}|$ equal n_p ? (In other words, are you awake?)

Now we can proceed with the proof. Let P act on $\mathcal{S}$ by conjugation. Then:

- Because P is a p -group, $n_p \pmod{p}$ is the number of fixed points of this action. Now we claim P is the only fixed point of this action.
- Let Q be any other fixed point, meaning $xQx^{-1} = Q$ for any $x \in P$.
- Define the normalizer $N_G(Q) = \{g \in G \mid gQg^{-1} = Q\}$. It contains both P and Q .
- Now for the crazy part: apply Step 2 to $N_G(Q)$. Since P and Q are Sylow p -subgroups of it, they must be conjugate.
- Hence $P = Q$, as desired.

Step 4: n_p divides m

Since $n_p \equiv 1 \pmod{p}$, it suffices to show n_p divides $|G|$. Let G act on the set of all Sylow p -groups by conjugation. Step 2 says this action has only one orbit, so the Orbit-Stabilizer theorem implies n_p divides $|G|$.

§11.3 (Optional) Simple groups and Jordan-Hölder

Prototypical example for this section: Decomposition of $\mathbb{Z}_{12}$ is $1 \trianglelefteq \mathbb{Z}_2 \trianglelefteq \mathbb{Z}_4 \trianglelefteq \mathbb{Z}_{12}$.

Just like every integer breaks down as the product of primes, we can try to break every group down as a product of “basic” groups. Armed with our idea of quotient groups, the right notion is this.

Definition 11.3.1. A **simple group** is a group with no normal subgroups other than itself and the trivial group.

Question 11.3.2. For which n is $\mathbb{Z}_n$ simple? (Hint: remember that $\mathbb{Z}_n$ is abelian.)

Then we can try to define what it means to “break down a group”.

Definition 11.3.3. A **composition series** of a group G is a sequence of subgroups $H_0, H_1, \dots, H_n$ such that

$$\{1\} = H_0 \trianglelefteq H_1 \trianglelefteq H_2 \trianglelefteq \dots \trianglelefteq H_n \trianglelefteq G$$

of maximal length (i.e. n is as large as possible, but all H_i are of course distinct). The **composition factors** are the groups $H_1/H_0, H_2/H_1, \dots, H_n/H_{n-1}$.

You can show that the “maximality” condition implies that the composition factors are all simple groups.

Let’s say two composition series are equivalent if they have the same composition factors (up to permutation); in particular they have the same length. Then it turns out that the following theorem *is* true.

Theorem 11.3.4 (Jordan-Hölder)

Every finite group G admits a unique composition series up to equivalence.

Example 11.3.5 (Fundamental theorem of arithmetic when $n = 12$)

Let's consider the group $\mathbb{Z}_{12}$. It's not hard to check that the possible composition series are

$$\{1\} \trianglelefteq \mathbb{Z}_2 \trianglelefteq \mathbb{Z}_4 \trianglelefteq \mathbb{Z}_{12} \text{ with factors } \mathbb{Z}_2, \mathbb{Z}_2, \mathbb{Z}_3$$

$$\{1\} \trianglelefteq \mathbb{Z}_2 \trianglelefteq \mathbb{Z}_6 \trianglelefteq \mathbb{Z}_{12} \text{ with factors } \mathbb{Z}_2, \mathbb{Z}_3, \mathbb{Z}_2$$

$$\{1\} \trianglelefteq \mathbb{Z}_3 \trianglelefteq \mathbb{Z}_6 \trianglelefteq \mathbb{Z}_{12} \text{ with factors } \mathbb{Z}_3, \mathbb{Z}_2, \mathbb{Z}_2.$$

These correspond to the factorization $12 = 2^2 \cdot 3$.

This suggests that classifying all finite simple groups would be great progress, since every finite group is somehow a “product” of simple groups; the only issue is that there are multiple ways of building a group from constituents.

Amazingly, we actually *have* a full list of simple groups, but the list is really bizarre. Every finite simple group falls in one of the following categories:

- $\mathbb{Z}_p$ for p a prime,
- For $n \geq 5$, the subgroup of S_n consisting of “even” permutations.
- A simple group of Lie type (which I won't explain), and
- Twenty-six “sporadic” groups which do not fit into any nice family.

The two largest of the sporadic groups have cute names. The **baby monster group** has order

$$2^{41} \cdot 3^{13} \cdot 5^6 \cdot 7^2 \cdot 11 \cdot 13 \cdot 17 \cdot 19 \cdot 23 \cdot 31 \cdot 47 \approx 4 \cdot 10^{33}$$

and the **monster group** (also “**friendly giant**”) has order

$$2^{46} \cdot 3^{20} \cdot 5^9 \cdot 7^6 \cdot 11^2 \cdot 13^3 \cdot 17 \cdot 19 \cdot 23 \cdot 29 \cdot 31 \cdot 41 \cdot 47 \cdot 59 \cdot 71 \approx 8 \cdot 10^{53}.$$

It contains twenty of the sporadic groups as subquotients (including itself), and these twenty groups are called the “**happy family**”.

Math is weird.

Question 11.3.6. Show that “finite simple group of order 2” is redundant in the sense that any group of order 2 is both finite and simple.

§11.4 Problems to think about

Problem 11A* (Cauchy's theorem). Let G be a group and let p be a prime dividing $|G|$. Prove¹ that G has an element of order p .

Problem 11B. Let G be a finite simple group. Show that $|G| \neq 56$.



Problem 11C (Engel's PSS?). Consider the set of all words consisting of the letters a and b . Given such a word, we can change the word either by inserting a word of the form www , where w is a word, anywhere in the given word, or by deleting such a sequence from the word. Can we turn the word ab into the word ba ?



Problem 11D. Let p be a prime. Show that the only simple group with order p^n (for some positive integer n) is the cyclic group $\mathbb{Z}_p$.

DRAFT (Evan Chen)
Updated February 15, 2017

¹ Cauchy's theorem can be proved without the Sylow theorems, and in fact can often be used to give alternate proofs of Sylow.

12 Rings and ideals

This is a typical chapter on commutative algebra.

§12.1 Number theory motivation

Commutative algebra is closely tied to algebraic geometry: lots of the ideas in commutative algebra have nice “geometric” interpretations that motivate the definitions. Unfortunately, algebraic geometry also has commutative algebra as a prerequisite, leading to a chicken-and-egg problem: to appreciate commutative algebra fully you need to know algebraic geometry, but to learn algebraic geometry you need to know commutative algebra.

To try and patch this as best I could, I’ll try to flavor and motivate the examples with olympiad number theory. This way, your intuition from all those shortlist N6’s you had to do can hopefully carry over to make sense of the examples here. Basically, we’ll try to generalize properties of the ring $\mathbb{Z}$ to any abelian structure in which we can also multiply. That’s why, for example, you can talk about “irreducible polynomials in $\mathbb{Q}[x]$ ” in the same way you can talk about “primes in $\mathbb{Z}$ ”, or about “factoring polynomials modulo p ” in the same way we can talk “unique factorization in $\mathbb{Z}$ ”. Even if you only care about $\mathbb{Z}$ (say, you’re a number theorist), this has a lot of value: I assure you that trying to solve $x^n + y^n = z^n$ (for $n > 2$) requires going into a ring other than $\mathbb{Z}$!

For all the sections that follow, keep $\mathbb{Z}$ in mind as your prototype.

§12.2 Definition and examples of rings

Prototypical example for this section: $\mathbb{Z}$ all the way! Also $R[x]$ and various fields.

Well, I guess I’ll define a ring¹.

Definition 12.2.1. A **ring** is a triple $(R, +, \times)$, the two operations usually called addition and multiplication, such that

- (i) $(R, +)$ is an abelian group, with identity 0_R , or just 0 .
- (ii) $\times$ is an associative, binary operation on R with some identity, written 1_R or just 1 .
- (iii) Multiplication distributes over addition.

The ring R is **commutative** if $\times$ is commutative.

Abuse of Notation 12.2.2. As usual, we will abbreviate $(R, +, \times)$ to just R .

Abuse of Notation 12.2.3. For simplicity, assume all rings are commutative for the rest of this chapter. We’ll run into some noncommutative rings eventually, but for such rings we won’t need the full theory of this chapter anyways.

These definitions are just here for completeness. The examples are much more important.

¹Or, according to some authors, a “ring with identity”. You see, some authors don’t require rings to have multiplicative identity. But I’ve never seen a ring without identity that I care about. So, just be warned that some authors have different definitions of ring. For us, “ring” always means “ring with 1”.

Example 12.2.4 (Typical rings and fields)

- (a) The sets $\mathbb{Z}$, $\mathbb{Q}$, $\mathbb{R}$ and $\mathbb{C}$ are all rings with the usual addition and multiplication.
- (b) The integers modulo n are also a ring with the usual addition and multiplication. We also denote it by $\mathbb{Z}_n$.
- (c) The **trivial ring** is the ring R with one element $0_R = 1_R$.

Since I've defined this structure, I may as well state the obligatory facts about it.

Fact 12.2.5. For any ring R and $r \in R$, $r \cdot 0_R = 0_R$. Moreover, $r \cdot (-1_R) = -r$.

Here are some more examples of rings.

Example 12.2.6 (Polynomial ring)

Given any ring R , the ring $R[x]$ is defined as the set of polynomials with coefficients in R :

$$R[x] = \{a_n x^n + a_{n-1} x^{n-1} + \cdots + a_0 \mid a_0, \dots, a_n \in R\}.$$

Addition and multiplication are done exactly in the way you would expect.

Remark 12.2.7 (Digression). Happily, polynomial division also does what we expect: if $p(x) \in R[x]$ and $p(a) = 0$, then $(x - a)q(x) = p(x)$ for some polynomial q . Proof: just do polynomial long division. With that, note the caveat that

$$x^2 - 1 \equiv (x - 1)(x + 1) \pmod{8}$$

has *four* roots 1, 3, 5, 7 in $\mathbb{Z}_8$.

The problem is that $2 \cdot 4 = 0$ even though 2 and 4 are not zero; we call 2 and 4 *zero divisors* for that reason. In an *integral domain* (a ring without zero divisors), this pathology goes away, and just about everything you know about polynomials carries over. (I'll say this all again next section.)

Example 12.2.8 (Multi-variable polynomial ring)

We can consider polynomials in n variables with coefficients in R , denoted $R[x_1, \dots, x_n]$. (We can even adjoin infinitely many x 's if we like!)

Example 12.2.9 (Gaussian integers are a ring)

The **Gaussian integers** are the set of complex numbers with integer real and imaginary parts, that is

$$\mathbb{Z}[i] = \{a + bi \mid a, b \in \mathbb{Z}\}.$$

Example 12.2.10 (Product ring)

Given two rings R and S the **product ring**, denoted $R \times S$, is defined as ordered pairs (r, s) with both operations done component-wise. For example, the Chinese remainder theorem says that $\mathbb{Z}_{15} \cong \mathbb{Z}_3 \times \mathbb{Z}_5$.

Remark 12.2.11. Equivalently, we can define $R \times S$ as the abelian group $R \oplus S$, and endow it with the multiplication where $r \cdot s = 0$ for $r \in R$, $s \in S$.

§12.3 Integral domains and fields

Prototypical example for this section: $\mathbb{Z}$ is an integral domain; $\mathbb{Q}$ is a field.

As you already know, if the multiplication is invertible, then we call the ring a field. To be explicit, let me write the relevant definitions.

Definition 12.3.1. A **unit** of a ring R is an element $u \in R$ which is invertible: for some $x \in R$ we have $ux = 1_R$.

Example 12.3.2 (Examples of units)

- (a) The units of $\mathbb{Z}$ are ± 1 , because these are the only things which “divide 1” (which is the reason for the name “unit”).
- (b) On the other hand, in $\mathbb{Q}$ everything is a unit (except 0). For example, $\frac{3}{5}$ is a unit since $\frac{3}{5} \cdot \frac{5}{3} = 1$.
- (c) The Gaussian integers $\mathbb{Z}[i]$ have four units: ± 1 and $\pm i$.

Definition 12.3.3. A nontrivial (commutative) ring is a **field** when all its nonzero elements are units.

Colloquially, we say that

A field is a structure where you can add, subtract, multiply, and divide.

Remark 12.3.4. You might say at this point that “fields are nicer than rings”, but as you’ll see in this chapter, the conditions for being a field are somehow “too strong”. To give an example of what I mean: if you try to think about the concept of “divisibility” in $\mathbb{Z}$, you’ve stepped into the vast and bizarre realm of number theory. Try to do the same thing in $\mathbb{Q}$ and you get nothing: any nonzero a “divides” any nonzero b because $b = a \cdot \frac{b}{a}$.

I know at least one person who instead thinks of this as an argument for why people shouldn’t care about number theory (studying chaos rather than order).

Now it would be nice if we could still conclude the zero product property: if $ab = 0$ then either $a = 0$ or $b = 0$. If our ring is a field, this is true: if $b \neq 0$, then we can multiply by b^{-1} to get $a = 0$. But many other rings we consider like $\mathbb{Z}$ and $\mathbb{Z}[x]$ also have this property, despite not being full-fledged fields.

Not for all rings though: in $\mathbb{Z}_{15}$,

$$3 \cdot 5 \equiv 0 \pmod{15}.$$

If $a, b \neq 0$ but $ab = 0$ then we say a and b are **zero divisors** of the ring R . So we give a name to such rings.

Definition 12.3.5. A nontrivial (commutative) ring with no zero divisors is called an **integral domain**.

Example 12.3.6 (Integral domains and fields)

- (a) $\mathbb{Q}, \mathbb{R}, \mathbb{C}$ are fields, since the notion $\frac{1}{c}$ makes sense in them.
- (b) $\mathbb{R}[x]$ is not a field, since there is no polynomial $P(x)$ with $xP(x) = 1$. However, $\mathbb{R}[x]$ is an integral domain, because if $P(x)Q(x) \equiv 0$ then one of P or Q is zero.
- (c) $\mathbb{Z}[x]$ is also an example of an integral domain. In fact, $R[x]$ is an integral domain for any integral domain R (why?).
- (d) $\mathbb{Z}_n$ is a field exactly when n is prime. When n is not prime, it is a ring but not an integral domain.

§12.4 Ideals

Prototypical example for this section: $5\mathbb{Z}$ is an ideal of $\mathbb{Z}$.

This section is going to go briskly – it’s the obvious generalization of all the stuff we did with quotient groups.²

First, we define a homomorphism and isomorphism.

Definition 12.4.1. Let $R = (R, +_R, \times_R)$ and $S = (S, +_S, \times_S)$ be rings. A **ring homomorphism** is a map $\phi : R \rightarrow S$ such that

- (i) $\phi(x +_R y) = \phi(x) +_S \phi(y)$ for each $x, y \in R$.
- (ii) $\phi(x \times_R y) = \phi(x) \times_S \phi(y)$ for each $x, y \in R$.
- (iii) $\phi(1_R) = 1_S$.

If ϕ is a bijection then ϕ is an **isomorphism** and we say that rings R and S are **isomorphic**.

Just what you would expect. The only surprise is that we also demand $\phi(1_R)$ to go to 1_S . This condition is not extraneous: consider the map $\mathbb{Z} \rightarrow \mathbb{Z}$ called “multiply by zero”.

Now, just like we were able to mod out by groups, we’d also like to define quotient rings. So once again,

Definition 12.4.2. The **kernel** of a ring homomorphism $\phi : R \rightarrow S$, denoted $\ker \phi$, is the set of $r \in R$ such that $\phi(r) = 0$.

In group theory, we were able to characterize the “normal” subgroups by a few obviously necessary conditions (namely, $gHg^{-1} = H$). We can do the same thing for rings, and it’s in fact easier because our operations are commutative.

First, note two obvious facts:

- If $\phi(x) = \phi(y) = 0$, then $\phi(x + y) = 0$ as well. So $\ker \phi$ should be closed under addition.
- If $\phi(x) = 0$, then for any $r \in R$ we have $\phi(rx) = \phi(r)\phi(x) = 0$ too. So for $x \in \ker \phi$ and any $r \in R$, we have $rx \in \ker \phi$.

²I once found an abstract algebra textbook which teaches rings before groups. At the time I didn’t understand why, but now I think I get it – modding out by things in commutative rings is far more natural, and you can start talking about all the various flavors of rings and fields. You also have (in my opinion) more vivid first examples for rings than for groups. I actually sympathize a lot with this approach – maybe I’ll convert Napkin to follow it one day.

A (nonempty) subset $I \subseteq R$ is called an ideal if it satisfies these properties. That is,

Definition 12.4.3. A nonempty subset $I \subseteq R$ is an **ideal** if it is closed under addition, and for each $x \in I$, $rx \in I$ for all $r \in R$.

Note that in the second condition, r need not be in I ! So this is stronger than just saying R is closed under multiplication.

Remark 12.4.4. If R is not commutative, we also need the condition $xr \in I$. That is, the ideal is *two-sided*: it absorbs multiplication from both the left and the right. But since most of our rings are commutative we needn't worry with this distinction.

Example 12.4.5 (Prototypical example of an ideal)

Consider the set $I = 5\mathbb{Z} = \{\dots, -10, -5, 0, 5, 10, \dots\}$ as an ideal in $\mathbb{Z}$. We indeed see I is the kernel of the “take mod 5” homomorphism:

$$\mathbb{Z} \rightarrow \mathbb{Z}/5\mathbb{Z}.$$

It's clearly closed under addition, but it absorbs multiplication from *all* elements of $\mathbb{Z}$: given $15 \in I$, $999 \in \mathbb{Z}$, we get $15 \cdot 999 \in I$.

Question 12.4.6. Let R be a ring and I an ideal. Convince yourself that if I contains any units, then $I = R$. (As an example, consider any ideal in the Gaussian integers which contains i .)

Question 12.4.7. Verify that in $\mathbb{R}$ (or actually any nontrivial field), the only two ideals are $(0) = \{0\}$ and $(1) = \mathbb{R}$.

Now we claim that these conditions are sufficient. More explicitly,

Theorem 12.4.8 (Ring analog of normal subgroups)

Let R be a ring and $I \subseteq R$. Then I is the kernel of some homomorphism if and only if it's an ideal.

Proof. It's quite similar to the proof for the normal subgroup thing, and you might try it yourself as an exercise.

Obviously the conditions are necessary. To see they're sufficient, we *define* a ring by “cosets”

$$S = \{r + I \mid r \in R\}.$$

These are the equivalence where we say $r_1 \sim r_2$ if $r_1 - r_2 \in I$ (think of this as taking “mod I ”). To see that these form a ring, we just have to check that the addition and multiplication we put on them is well-defined. Specifically, we want to check that if $r_1 \sim s_1$ and $r_2 \sim s_2$, then $r_1 + r_2 \sim s_1 + s_2$ and $r_1 r_2 \sim s_1 s_2$. We actually already did the first part – just think of R and S as abelian groups, forgetting for the moment that we can multiply. The multiplication is more interesting.

Exercise 12.4.9 (Recommended). Show that if $r_1 \sim s_1$ and $r_2 \sim s_2$, then $r_1 r_2 \sim s_1 s_2$. You will need to use the fact that I absorbs multiplication from *any* elements of R , not just those in I .

Anyways, since this addition and multiplication is well-defined there is now a surjective homomorphism $R \rightarrow S$ with kernel exactly I . $\square$

Definition 12.4.10. Given an ideal I , we define as above the **quotient ring**

$$R/I \stackrel{\text{def}}{=} \{r + I \mid r \in R\}.$$

It's the ring of these equivalence classes.

Example 12.4.11 ($\mathbb{Z}/5\mathbb{Z}$)

The integers modulo 5 formed by “modding out additively by 5” are the $\mathbb{Z}_5$ we have already met.

But here's an important point: just as we don't actually think of $\mathbb{Z}/5\mathbb{Z}$ as consisting of $k + 5\mathbb{Z}$ for $k = 0, \dots, 4$, we also don't really want to think about R/I as elements $r + I$. The better way to think about it is

R/I is the result when we declare that elements of I are all zero; that is, we “mod out by elements of I ”.

For example, modding out by $5\mathbb{Z}$ means that we consider all elements in $\mathbb{Z}$ divisible by 5 to be zero. This gives you the usual modular arithmetic!

§12.5 Generating ideals

Let's give you some practice with ideals.

Exercise 12.5.1. Show that the only ideals of $\mathbb{Z}$ are precisely those sets of the form $n\mathbb{Z}$, where n is an integer.

Thus, while ideals of fields are not terribly interesting, ideals of $\mathbb{Z}$ look eerily like elements of $\mathbb{Z}$. Let's make this more precise.

Definition 12.5.2. Let R be a ring. The **ideal generated** by a set of elements $x_1, \dots, x_n \in R$ is denoted by $I = (x_1, x_2, \dots, x_n)$ and given by

$$I = \{r_1 x_1 + \dots + r_n x_n \mid r_i \in R\}.$$

One can think of this as “the smallest ideal containing all the x_i ”.

The “additive structure” of ideals is emphasized if I say

An ideal is an R -module. The ideal $(x_1, \dots, x_n)$ is the submodule spanned by $x_1, \dots, x_n$.

In particular, if $I = (x)$ then I consists of exactly the “multiples of x ”, i.e. numbers of the form rx for $r \in R$.

Remark 12.5.3. We can also apply this definition to infinite generating sets, as long as only finitely many of the r_i are not zero (since infinite sums don't make sense in general).

Example 12.5.4 (Examples of generated ideals)

- (a) As $(n) = n\mathbb{Z}$ for all $n \in \mathbb{Z}$, every ideal in $\mathbb{Z}$ is of the form (n) .
- (b) In $\mathbb{Z}[i]$, we have $(5) = \{5a + 5bi \mid a, b \in \mathbb{Z}\}$.
- (c) In $\mathbb{Z}[x]$, the ideal (x) consists of polynomials with zero constant terms.
- (d) In $\mathbb{Z}[x, y]$, the ideal (x, y) again consists of polynomials with zero constant terms.
- (e) In $\mathbb{Z}[x]$, the ideal $(x, 5)$ consists of polynomials whose constant term is divisible by 5.

Question 12.5.5. Please check that the set $I = \{r_1x_1 + \cdots + r_nx_n \mid r_i \in R\}$ is indeed always an ideal (closed under addition, and absorbs multiplication).

Now suppose $I = (x_1, \dots, x_n)$. What does R/I look like? According to what I said at the end of the last section, it's what happens when we “mod out” by each of the elements x_i . For example...

Example 12.5.6 (Modding out by generated ideals)

- (a) Let $R = \mathbb{Z}$ and $I = (5)$. Then R/I is literally $\mathbb{Z}/5\mathbb{Z}$, or the “integers modulo 5”: it is the result of declaring $5 = 0$.
- (b) Let $R = \mathbb{Z}[x]$ and $I = (x)$. Then R/I means we send x to zero; hence $R/I \cong \mathbb{Z}$ as given any polynomial $p(x) \in R$, we simply get its constant term.
- (c) Let $R = \mathbb{Z}[x]$ again and now let $I = (x - 3)$. Then R/I should be thought of as the quotient when $x - 3 \equiv 0$, that is, $x \equiv 3$. So given a polynomial $p(x)$ its image after we mod out should be thought of as $p(3)$. Again $R/I \cong \mathbb{Z}$, but in a different way.
- (d) Finally, let $I = (x - 3, 5)$. Then R/I not only sends x to three, but also 5 to zero. So given $p \in R$, we get $p(3) \pmod{5}$. Then $R/I \cong \mathbb{Z}/5\mathbb{Z}$.

By the way, given an ideal I of a ring R , it's totally legit to write

$$x \equiv y \pmod{I}$$

to mean that $x - y \in I$. Everything you learned about modular arithmetic carries over.

§12.6 Principal ideal domains and Noetherian rings

Prototypical example for this section: $\mathbb{Z}$ is a PID, $\mathbb{Z}[x]$ is not but is at least Noetherian. $\mathbb{Z}[x_1, x_2, \dots]$ is not Noetherian.

What happens if we put multiple generators in an ideal, like $(10, 15) \subseteq \mathbb{Z}$? Well, we have by definition that $(10, 15)$ is given as a set by

$$(10, 15) \stackrel{\text{def}}{=} \{10x + 15y \mid x, y \in \mathbb{Z}\}.$$

If you're good at Number Theory you'll instantly recognize that this as just $5\mathbb{Z} = (5)$. Surprise! In $\mathbb{Z}$, the ideal (a, b) is exactly $\gcd(a, b)\mathbb{Z}$. And that's exactly the reason you often see the GCD of two numbers denoted (a, b) .

We call such an ideal (one generated by a single element) a **principal ideal**. So, in $\mathbb{Z}$, every ideal is principal. But the same is not true in more general rings.

Example 12.6.1 (A non-principal ideal)

In $\mathbb{Z}[x]$, $I = (x, 2015)$ is *not* a principal ideal. For if $I = (p)$ for some polynomial $p \in I$ then p divides x and 2015. This can only occur if $p = \pm 1$, but then I contains a unit.

A ring with the property that all its ideals are principal is called a **principal ideal domain**, which is abbreviated PID. We like PID's because they effectively let us take the "greatest common factor" in a similar way as the GCD in $\mathbb{Z}$.

If it's too much to ask that an ideal is generated by *one* element, perhaps we can at least ask that our ideals are generated by *finitely many* elements. Unfortunately, in certain weird rings this is also not the case.

Example 12.6.2 (Non-Noetherian ring)

Consider the ring $R = \mathbb{Z}[x_1, x_2, x_3, \dots]$ which has *infinitely* many free variables. Then the ideal $I = (x_1, x_2, \dots) \subseteq R$ cannot be written with a finite generating set.

Nonetheless, most "sane" rings we work in *do* have the property that their ideals are finitely generated. We now name such rings and give two equivalent definitions:

Proposition 12.6.3 (The equivalent definitions of a Noetherian ring)

For a ring R , the following are equivalent:

- (a) Every ideal I of R is finitely generated (i.e. can be written with a finite generating set).
- (b) There does *not* exist an infinite ascending chain of ideals

$$I_1 \subsetneq I_2 \subsetneq I_3 \subsetneq \dots$$

The absence of such chains is often called the **ascending chain condition**.

Such rings are called **Noetherian**.

Example 12.6.4 (Non-Noetherian ring breaks ACC)

In the ring $R = \mathbb{Z}[x_1, x_2, x_3, \dots]$ we have an infinite ascending chain

$$(x_1) \subsetneq (x_1, x_2) \subsetneq (x_1, x_2, x_3) \subsetneq \dots$$

From the example, you can kind of see why the proposition is true: from an infinitely generated ideal you can extract an ascending chain by throwing elements in one at a time. I'll leave the proof to you if you want to do it.³

³On the other hand, every undergraduate class in this topic I've seen makes you do it as homework. Admittedly I haven't gone to that many such classes.

Question 12.6.5. Why are fields Noetherian? Why are PID's (such as $\mathbb{Z}$) Noetherian?

This leaves the question: is our prototypical non-example of a PID, $\mathbb{Z}[x]$, a Noetherian ring? The answer is a glorious yes, according to the celebrated Hilbert basis theorem.

Theorem 12.6.6 (Hilbert basis theorem)

Given a Noetherian ring R , the ring $R[x]$ is also Noetherian. Thus by induction, $R[x_1, x_2, \dots, x_n]$ is Noetherian for any integer n .

The proof of this theorem is really olympiad flavored, so I couldn't possibly spoil it – I've left it as a problem at the end of this chapter.

Noetherian rings really shine in algebraic geometry, and it's a bit hard for me to motivate them right now, other than to just say “almost all rings you'll ever care about are Noetherian”. Please bear with me!

§12.7 Prime ideals

We know that every integer can be factored (up to sign) as a unique product of primes; for example $15 = 3 \cdot 5$ and $-10 = -2 \cdot 5$. You might remember the proof involves the so-called Bézout's lemma, which essentially says that $(a, b) = (\gcd(a, b))$; in other words we've carefully used the fact that $\mathbb{Z}$ is a PID.

It turns out that for general rings, the situation is not as nice as factoring elements because most rings are not PID's. The classic example of something going wrong is

$$6 = 2 \cdot 3 = (1 - \sqrt{5})(1 + \sqrt{5})$$

in $\mathbb{Z}[\sqrt{-5}]$. Nonetheless, we can sidestep the issue and talk about factoring *ideals*: somehow the example $10 = 2 \cdot 5$ should be $(10) = (2) \cdot (5)$, which says “every multiple of 10 is the product of a multiple of 2 and a multiple of 5”. I'd have to tell you then how to multiply two ideals, which I do in the chapter on unique factorization.

Let's at least figure out what primes are. In $\mathbb{Z}$, we have that $p \neq 1$ is prime if whenever $p \mid xy$, either $p \mid x$ or $p \mid y$. We port over this definition to our world of ideals.

Definition 12.7.1. An ideal I is **proper** if $I \neq R$, or equivalently, I contains no units.

Definition 12.7.2. A proper ideal $I \subsetneq R$ is a **prime ideal** if whenever $xy \in I$, either $x \in I$ or $y \in I$.

The condition that I is proper is analogous to the fact that we don't consider 1 to be a prime number.

Example 12.7.3 (Examples of prime ideals)

- (a) The ideal (7) of $\mathbb{Z}$ is prime.
- (b) The ideal (x) of $\mathbb{Z}[x]$ is prime.
- (c) The ideal $(5) = 5\mathbb{Z} + 5i\mathbb{Z}$ of $\mathbb{Z}[i]$ is *not* prime, since the elements $3 + i$ and $3 - i$ have product $10 \in (5)$, yet neither is itself in (5) .

Remark 12.7.4. Ideals have the nice property that they get rid of “sign issues”. For example, in $\mathbb{Z}$, do we consider -3 to be a prime? When phrased with ideals, this annoyance goes away: $(-3) = (3)$.

More generally, for a ring R , talking about ideals lets us ignore multiplication by a unit. (Note that -1 is a unit in $\mathbb{Z}$.)

Exercise 12.7.5. What do you call a ring R for which the zero ideal (0) is prime?

We also have:

Theorem 12.7.6 (Prime ideal $\iff$ quotient is integral domain)

An ideal I is prime if and only if R/I is an integral domain.

Exercise 12.7.7. (Mandatory) Convince yourself the theorem is true. (A possible start is to consider $R = \mathbb{Z}$ and $I = (15)$.)

I now must regrettably inform you that unique factorization is still not true even with the notion of a “prime” ideal (though again I haven’t told you how to multiply two ideals yet). But it will become true with some additional assumptions, namely that we have a Dedekind domain.

§12.8 Maximal ideals

Here’s another flavor of an ideal.

Definition 12.8.1. A proper ideal I of a ring R is **maximal** if no proper ideal J contains it.

Example 12.8.2 (Examples of maximal ideals)

- (a) The ideal $I = (7)$ of $\mathbb{Z}$ is maximal, because if an ideal J contains 7 and an element n not in I it must contain $\gcd(7, n) = 1$, and hence $J = \mathbb{Z}$.
- (b) The ideal (x) is *not* maximal in $\mathbb{Z}[x]$, because it’s contained in $(x, 5)$ (among others).
- (c) On the other hand, $(x, 5)$ is indeed maximal in $\mathbb{Z}[x]$, as we will see in a moment.

Exercise 12.8.3. What do you call a ring R for which the zero ideal (0) is maximal?

There’s an analogous theorem to the one for prime ideals.

Theorem 12.8.4

An ideal I is maximal if and only if R/I is a field.

Proof. See if you can convince yourself that there is a one-to-one correspondence between

- (i) Ideals J with $I \subseteq J \subseteq R$, and
- (ii) Ideals of R/I .

You may want to start by taking $R = \mathbb{Z}$ and $I = (15)$. In any case, the theorem follows from this. $\square$

Question 12.8.5. Show that maximal ideals are prime. (The proof is very short.)

In practice, because modding out by generated ideals is pretty convenient, this is a very efficient way to check whether an ideal is maximal.

Example 12.8.6 (Modding out in $\mathbb{Z}[x]$)

- (a) This instantly implies that $(x, 5)$ is a maximal ideal in $\mathbb{Z}[x]$, because if we mod out by x and 5 in $\mathbb{Z}[x]$, we just get $\mathbb{Z}/5\mathbb{Z}$, which is a field.
- (b) On the other hand, modding out by just x gives $\mathbb{Z}$, which is an integral domain but not a field; that's why (x) is prime but not maximal.

As we say, any maximal ideal is prime. But now note that $\mathbb{Z}$ has the special property that all of its prime ideals are also maximal. It's with this condition and a few other minor conditions that you get a so-called *Dedekind domain* where prime factorization of ideals *does* work. More on that later.

§12.9 Problems

Not olympiad problems, but again the spirit is very close to what you might see in an olympiad.

Problem 12A*. Prove that a finite integral domain is in fact a field!

Problem 12B* (Krull's theorem). Let R be a ring and J an ideal.

- (a) Prove that if R is Noetherian, then J is contained in a maximal ideal I .
- (b) Use Zorn's lemma (Chapter 57) to prove the result even if R isn't Noetherian.

Problem 12C*. Prove that any nonzero prime ideal of $\mathbb{Z}[\sqrt{2}]$ is also a maximal ideal.



Problem 12D. Prove the Hilbert basis theorem.



Problem 12E. Let $A \subseteq B \subseteq C$ be rings. Suppose C is a finitely generated A -module. Does it follow that B is a finitely generated A -module?

13 The PID structure theorem

The main point of this chapter is to discuss a classification theorem for finitely generated abelian groups. This won't take long to do, and if you like you can read just the first section and then move on.

However, since I'm here, I will go ahead and state the result as a special case of the much more general *structure theorem*. Its corollaries include

- All finite-dimensional vector spaces are $k^{\oplus n}$.
- The classification theorem for finitely generated abelian groups,
- The Jordan decomposition of a matrix from before,
- Another canonical form for a matrix: “Frobenius normal form”.

§13.1 Finitely generated abelian groups

Remark 13.1.1. We talk about abelian groups in what follows, but really the morally correct way to think about these structures is as $\mathbb{Z}$ -modules.

Definition 13.1.2. An abelian group $G = (G, +)$ is **finitely generated** if it is finitely generated as a $\mathbb{Z}$ -module. (That is, there exists a finite collection $b_1, \dots, b_m \in G$, such that every $x \in G$ can be written in the form $c_1 b_1 + \dots + c_m b_m$ for some $c_1, \dots, c_m \in \mathbb{Z}$.)

Example 13.1.3 (Examples of finitely generated abelian groups)

- (a) $\mathbb{Z}$ is finitely generated (by 1).
- (b) $\mathbb{Z}_n$ is finitely generated (by 1).
- (c) $\mathbb{Z}^{\oplus 2}$ is finitely generated (by two elements $(1, 0)$ and $(0, 1)$).
- (d) $\mathbb{Z}^{\oplus 3} \oplus \mathbb{Z}_9 \oplus \mathbb{Z}_{2016}$ is finitely generated by five elements.
- (e) $\mathbb{Z}_3 \oplus \mathbb{Z}_5$ is finitely generated by two elements.

Exercise 13.1.4. In fact $\mathbb{Z}_3 \oplus \mathbb{Z}_5$ is generated by *one* element. What is it?

You might notice that these examples are not very diverse. That's because they are actually the only examples:

Theorem 13.1.5 (Fundamental theorem of finitely generated abelian groups)

Let G be a finitely generated abelian group. Then there exists an integer r , prime powers $q_1, \dots, q_m$ (not necessarily distinct) such that

$$G \cong \mathbb{Z}^{\oplus r} \oplus \mathbb{Z}_{q_1} \oplus \mathbb{Z}_{q_2} \oplus \dots \oplus \mathbb{Z}_{q_m}.$$

This decomposition is unique up to permutation of the $\mathbb{Z}_{q_i}$.

Definition 13.1.6. The **rank** of a finitely generated abelian group G is the integer r above.

Now, we could prove this theorem, but it is more interesting to go for the gold and state and prove the entire structure theorem.

§13.2 Some ring theory prerequisites

Prototypical example for this section: $R = \mathbb{Z}$.

Before I can state the main theorem, I need to define a few terms for UFD's, which behave much like $\mathbb{Z}$:

Our intuition from the case $R = \mathbb{Z}$ basically carries over verbatim.

We don't even need to deal with prime ideals and can factor elements instead.

Definition 13.2.1. If R is a UFD, then $p \in R$ is a **prime element** if (p) is a prime ideal and $p \neq 0$. For UFD's this is equivalent to: if $p = xy$ then either x or y is a unit.

So for example in $\mathbb{Z}$ the set of prime elements is $\{\pm 2, \pm 3, \pm 5, \dots\}$. Now, since R is a UFD, every element r factors into a product of prime elements

$$r = up_1^{e_1} p_2^{e_2} \cdots p_m^{e_m}$$

Definition 13.2.2. We say r **divides** s if $s = r'r$ for some $r' \in R$. This is written $r \mid s$.

Example 13.2.3 (Divisibility in $\mathbb{Z}$)

The number 0 is divisible by every element of $\mathbb{Z}$. All other divisibility as expected.

Question 13.2.4. Show that $r \mid s$ if and only if the exponent of each prime in r is less than or equal to the corresponding exponent in s .

Now, the case of interest is the even stronger case when R is a PID:

Proposition 13.2.5 (PID's are Noetherian UFD's)

If R is a PID, then it is Noetherian and also a UFD.

Proof. The fact that R is Noetherian is obvious. For R to be a UFD we essentially repeat the proof for $\mathbb{Z}$, using the fact that (a, b) is principal in order to extract $\gcd(a, b)$. $\square$

In this case, we have a Chinese remainder theorem for elements.

Theorem 13.2.6 (Chinese remainder theorem for rings)

Let m and n be relatively prime elements, meaning $(m) + (n) = (1)$. Then

$$R/(mn) \cong R/m \times R/n.$$

Here the ring product is as defined in Example 12.2.10.

Proof. This is the same as the proof of the usual Chinese remainder theorem. First, since $(m, n) = (1)$ we have $am + bn = 1$ for some a and b . Then we have a map

$$R/m \times R/n \rightarrow R/(mn) \quad \text{by} \quad (r, s) \mapsto r \cdot bn + s \cdot am.$$

One can check that this map is well-defined and an isomorphism of rings. (Diligent readers invited to do so.) $\square$

Finally, we need to introduce the concept of a Noetherian R -module.

Definition 13.2.7. An R -module M is **Noetherian** if it satisfies one of the two equivalent conditions:

- Its submodules obey the ascending chain condition: there is no infinite sequence of modules $M_1 \subsetneq M_2 \subsetneq \dots$.
- All submodules of M (including M itself) are finitely generated.

This generalizes the notion of a Noetherian ring: a Noetherian ring R is one for which R is Noetherian as an R -module.

Question 13.2.8. Check these two conditions are equivalent. (Copy the proof for rings.)

§13.3 The structure theorem

Our structure theorem takes two forms:

Theorem 13.3.1 (Structure theorem, invariant form)

Let R be a PID and let M be any finitely generated R -module. Then

$$M \cong \bigoplus_{i=1}^m R/s_i$$

for some s_i satisfying $s_1 \mid s_2 \mid \dots \mid s_m$.

Corollary 13.3.2 (Structure theorem, primary form)

Let R be a PID and let M be any finitely generated R -module. Then

$$M \cong R^{\oplus r} \oplus R/(q_1) \oplus R/(q_2) \oplus \dots \oplus R/(q_m)$$

where $q_i = p_i^{e_i}$ for some prime element p_i and integer $e_i \geq 1$.

Proof of corollary. Factor each s_i into prime factors (since R is a UFD), then use the Chinese remainder theorem. $\square$

Remark 13.3.3. In both theorems the decomposition is unique up to permutations of the summands; good to know, but I won't prove this.

§13.4 Reduction to maps of free R -modules

Definition 13.4.1. A **free R -module** is a module of the form $R^{\oplus n}$ (or more generally, $\bigoplus_I R$ for some indexing set I , just to allow an infinite basis).

The proof of the structure theorem proceeds in two main steps. First, we reduce the problem to a *linear algebra* problem involving free R -modules $R^{\oplus d}$. Once that's done, we just have to play with matrices; this is done in the next section.

Suppose M is finitely generated by d elements. Then there is a surjective map of R -modules

$$R^{\oplus d} \twoheadrightarrow M$$

whose image on the basis of $R^{\oplus d}$ are the generators of M . Let K denote the kernel.

We claim that K is finitely generated as well. To this end we prove that

Lemma 13.4.2 (Direct sum of Noetherian modules is Noetherian)

Let M and N be two Noetherian R -modules. Then the direct sum $M \oplus N$ is also a Noetherian R -module.

Proof. It suffices to show that if $L \subseteq M \oplus N$, then L is finitely generated. One guess is that $L = P \oplus Q$, where P and Q are the projections of L onto M and N . Unfortunately this is false (take $M = N = \mathbb{Z}$ and $L = \{(n, n) \mid n \in \mathbb{Z}\}$) so we will have to be more careful.

Consider the submodules

$$\begin{aligned} A &= \{x \in M \mid (x, 0) \in L\} \subseteq M \\ B &= \{y \in N \mid \exists x \in M : (x, y) \in L\} \subseteq N. \end{aligned}$$

(Note the asymmetry for A and B : the proof doesn't work otherwise.) Then A is finitely generated by $a_1, \dots, a_k$, and B is finitely generated by $b_1, \dots, b_\ell$. Let $x_i = (a_i, 0)$ and let $y_i = (*, b_i)$ be elements of L (where the $*$'s are arbitrary things we don't care about). Then x_i and y_i together generate L . $\square$

Question 13.4.3. Deduce that for R a PID, $R^{\oplus d}$ is Noetherian.

Hence $K \subseteq R^{\oplus d}$ is finitely generated as claimed. So we can find another surjective map $R^{\oplus f} \twoheadrightarrow K$. Consequently, we have a composition

$$\begin{array}{ccccc} & & K & & \\ & \nearrow & \hookrightarrow & \searrow & \\ R^{\oplus f} & \xrightarrow{T} & R^{\oplus d} & \longrightarrow & M \end{array}$$

Observe that M is the *cokernel* of the linear map T , i.e. we have that

$$M \cong R^{\oplus d} / \text{im}(T).$$

So it suffices to understand the map T well.

§13.5 Smith normal form

The idea is now that we have reduced our problem to studying linear maps $T : R^{\oplus m} \rightarrow R^{\oplus n}$, which can be thought of as a generic matrix

$$T = \begin{bmatrix} a_{11} & \cdots & a_{1m} \\ \vdots & \ddots & \vdots \\ a_{n1} & \cdots & a_{nm} \end{bmatrix}$$

for a basis $e_1, \dots, e_m$ of $R^{\oplus m}$ and $f_1, \dots, f_n$ of N .

Of course, as you might expect it ought to be possible to change the given basis of T such that T has a nicer matrix form. We already saw this in *Jordan form*, where we had a map $T : V \rightarrow V$ and changed the basis so that T was “almost diagonal”. This time, we have *two* sets of bases we can change, so we would hope to get a diagonal basis, or even better.

Before proceeding let’s think about how we might edit the matrix: what operations are permitted? Here are some examples:

- Swapping rows and columns, which just corresponds to re-ordering the basis.
- Adding a multiple of a column to another column. For example, if we add 3 times the first column to the second column, this is equivalent to replacing the basis

$$(e_1, e_2, e_3, \dots, e_m) \mapsto (e_1, e_2 + 3e_1, e_3, \dots, e_m).$$

- Adding a multiple of a row to another row. One can see that adding 3 times the first row to the second row is equivalent to replacing the basis

$$(f_1, f_2, f_3, \dots, f_n) \mapsto (f_1 - 3f_2, f_2, f_3, \dots, f_n).$$

More generally,

If A is an invertible $n \times n$ matrix we can replace T with AT .

This corresponds to replacing

$$(f_1, \dots, f_n) \mapsto (A(f_1), \dots, A(f_n))$$

(the “invertible” condition just guarantees the latter is a basis). Of course similarly we can replace X with XB where B is an invertible $m \times m$ matrix; this corresponds to

$$(e_1, \dots, e_m) \mapsto (B^{-1}(e_1), \dots, B^{-1}(e_m))$$

Armed with this knowledge, we can now approach:

Theorem 13.5.1 (Smith normal form)

Let R be a PID. Let $M = R^{\oplus m}$ and $N = R^{\oplus n}$ be free R -modules and let $T : M \rightarrow N$ be a linear map. Set $k = \min(m, n)$.

Then we can select a pair of new bases for M and N such that T has only diagonal entries $s_1, s_2, \dots, s_k$ and $s_1 \mid s_2 \mid \cdots \mid s_k$.

So if $m > n$, the matrix should take the form

$$\begin{bmatrix} s_1 & 0 & 0 & 0 & \dots & 0 \\ 0 & s_2 & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \dots & \vdots \\ 0 & 0 & 0 & s_n & \dots & 0 \end{bmatrix}.$$

and similarly when $m \leq n$.

Question 13.5.2. Show that Smith normal form implies the structure theorem.

Remark 13.5.3. Note that this is not a generalization of Jordan form.

- In Jordan form we consider maps $T : V \rightarrow V$; note that the source and target space are the *same*, and we are considering one basis for the space V .
- In Smith form the maps $T : M \rightarrow N$ are between *different* modules, and we pick *two* sets of bases (one for M and one for N).

Example 13.5.4 (Example of Smith normal form)

To give a flavor of the idea of the proof, let's work through a concrete example with the $\mathbb{Z}$ -matrix

$$\begin{bmatrix} 18 & 38 & 48 \\ 14 & 30 & 38 \end{bmatrix}.$$

The GCD of all the entries is 2, and so motivated by this, we perform the **Euclidean algorithm on the left column**: subtract the second row from the first row, then three times the first row from the second:

$$\begin{bmatrix} 18 & 38 & 48 \\ 14 & 30 & 38 \end{bmatrix} \mapsto \begin{bmatrix} 4 & 8 & 10 \\ 14 & 30 & 38 \end{bmatrix} \mapsto \begin{bmatrix} 4 & 8 & 10 \\ 2 & 6 & 2 \end{bmatrix}.$$

Now that the GCD of 2 is present, we move it to the upper-left by switching the two rows, and then kill off all the entries in the same row/column; since 2 was the GCD all along, we isolate 2 completely:

$$\begin{bmatrix} 4 & 8 & 10 \\ 2 & 6 & 2 \end{bmatrix} \mapsto \begin{bmatrix} 2 & 6 & 2 \\ 4 & 8 & 10 \end{bmatrix} \mapsto \begin{bmatrix} 2 & 6 & 2 \\ 0 & -4 & 6 \end{bmatrix} \mapsto \begin{bmatrix} 2 & 0 & 0 \\ 0 & -4 & 6 \end{bmatrix}.$$

This reduces the problem to a 1×2 matrix. So we just apply the Euclidean algorithm again there:

$$\begin{bmatrix} 2 & 0 & 0 \\ 0 & -4 & 6 \end{bmatrix} \mapsto \begin{bmatrix} 2 & 0 & 0 \\ 0 & -4 & 2 \end{bmatrix} \mapsto \begin{bmatrix} 2 & 0 & 0 \\ 0 & 0 & 2 \end{bmatrix} \mapsto \begin{bmatrix} 2 & 0 & 0 \\ 0 & 2 & 0 \end{bmatrix}.$$

Now all we have to do is generalize this proof to work with any PID. It's intuitively clear how to do this: the PID condition more or less lets you perform a Euclidean algorithm.

Proof of Smith normal form. Begin with a generic matrix

$$T = \begin{bmatrix} a_{11} & \dots & a_{1m} \\ \vdots & \ddots & \vdots \\ a_{n1} & \dots & a_{nm} \end{bmatrix}$$

We want to show, by a series of operations (gradually changing the given basis) that we can rearrange the matrix into Smith normal form.

Define $\gcd(x, y)$ to be any generator of the principal ideal (x, y) .

Claim 13.5.5 (“Euclidean algorithm”). If a and b are entries in the same row or column, we can change bases to replace a with $\gcd(a, b)$ and b with something else.

Proof. We do just the case of columns. By hypothesis, $\gcd(a, b) = xa + yb$ for some $x, y \in R$. We must have $(x, y) = (1)$ now (we’re in a UFD). So there are u and v such that $xu + yv = 1$. Then

$$\begin{bmatrix} x & y \\ -v & u \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} \gcd(a, b) \\ \text{something} \end{bmatrix}$$

and the first matrix is invertible (check this!), as desired. $\blacksquare$

Let $s_1 = (a_{ij})_{i,j}$ be the GCD of all entries. Now by repeatedly applying this algorithm, we can cause s to appear in the upper left hand corner. Then, we use it to kill off all the entries in the first row and the first column, thus arriving at a matrix

$$\begin{bmatrix} s_1 & 0 & 0 & \dots & 0 \\ 0 & a'_{22} & a'_{23} & \dots & a'_{2n} \\ 0 & a'_{32} & a'_{33} & \dots & a'_{3n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & a'_{m2} & a'_{m3} & \dots & a'_{mn} \end{bmatrix}.$$

Now we repeat the same procedure with this lower-right $(m-1) \times (n-1)$ matrix, and so on. This gives the Smith normal form. $\square$

With the Smith normal form, we have in the original situation that

$$M \cong R^{\oplus d} / \text{im } T$$

and applying the theorem to T completes the proof of the structure theorem.

§13.6 Problems to think about

Now, we can apply our structure theorem!

Problem 13A[†] (Finite-dimensional vector spaces are all isomorphic). A vector space V over a field k has a finite spanning set of vectors. Show that $V \cong k^{\oplus n}$ for some n .

Problem 13B[†] (Frobenius normal form). Let $T : V \rightarrow V$ where V is a finite-dimensional vector space over an arbitrary field k (not necessarily algebraically closed). Show that one can write T as a block-diagonal matrix whose blocks are all of the form

$$\begin{bmatrix} 0 & 0 & 0 & \dots & 0 & * \\ 1 & 0 & 0 & \dots & 0 & * \\ 0 & 1 & 0 & \dots & 0 & * \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & 1 & * \end{bmatrix}.$$

(View V as a $k[x]$ -module with action $x \cdot v = T(v)$.)

Problem 13C[†] (Jordan normal form). Let $T : V \rightarrow V$ where V is a finite-dimensional vector space over an arbitrary field k which is algebraically closed. Prove that T can be written in Jordan form.

IV

Complex Analysis

14 Holomorphic functions	153
14.1 The nicest functions on earth	153
14.2 Complex differentiation	155
14.3 Contour integrals	156
14.4 Cauchy-Goursat theorem	158
14.5 Cauchy's integral theorem	159
14.6 Holomorphic functions are analytic	161
14.7 Problems to think about	163
15 Meromorphic functions	164
15.1 The second nicest functions on earth	164
15.2 Meromorphic functions	164
15.3 Winding numbers and the residue theorem	167
15.4 Argument principle	169
15.5 Philosophy: why are holomorphic functions so nice?	170
15.6 Problems to think about	170
16 Holomorphic square roots and logarithms	171
16.1 Motivation: square root of a complex number	171
16.2 Square roots of holomorphic functions	173
16.3 Covering projections	174
16.4 Complex logarithms	174
16.5 Some special cases	175
16.6 Problems to think about	176

14 Holomorphic functions

Throughout this chapter, we denote by U an open subset of the complex plane, and by Ω an open subset which is also simply connected. The main references for this chapter were [Ya12; Ba10].

§14.1 The nicest functions on earth

In high school you were told how to differentiate and integrate real-valued functions. In this chapter on complex analysis, we'll extend it to differentiation and integration of complex-valued functions.

Big deal, you say. Calculus was boring enough. Why do I care about complex calculus?

Perhaps it's easiest to motivate things if I compare real analysis to complex analysis. In real analysis, your input lives inside the real line $\mathbb{R}$. This line is not terribly discerning – you can construct a lot of unfortunate functions. Here are some examples.

Example 14.1.1 (Optional: evil real functions)

You can skim over these very quickly: they're just here to make a point.

(a) The **Devil's Staircase** (or Cantor function) is a continuous function $H : [0, 1] \rightarrow [0, 1]$ which has derivative zero “almost everywhere”, yet $H(0) = 0$ and $H(1) = 1$.

(b) The **Weierstraß function**

$$x \mapsto \sum_{n=0}^{\infty} \left(\frac{1}{2}\right)^n \cos(2015^n \pi x)$$

is continuous *everywhere* but differentiable *nowhere*.

(c) The function

$$x \mapsto \begin{cases} x^{100} & x \geq 0 \\ -x^{100} & x < 0 \end{cases}$$

has the first 99 derivatives but not the 100th one.

(d) If a function has all derivatives (we call these **smooth** functions), then it has a Taylor series. But for real functions that Taylor series might still be wrong. The function

$$x \mapsto \begin{cases} e^{-1/x} & x > 0 \\ 0 & x \leq 0 \end{cases}$$

has derivatives at every point. But if you expand the Taylor series at $x = 0$, you get $0 + 0x + 0x^2 + \dots$, which is wrong for *any* $x > 0$ (even $x = 0.0001$).

Let's even put aside the pathology. If I tell you the value of a real smooth function on the interval $[-1, 1]$, that still doesn't tell you anything about the function as a whole. It could be literally anything, because it's somehow possible to “fuse together” smooth functions.




Figure 14.1: The Weierstraß Function (image from [Ee]).

So what about complex functions? If you just consider them as functions $\mathbb{R}^2 \rightarrow \mathbb{R}^2$, you now have the interesting property that you can integrate along things that are not just line segments: you can write integrals across curves in the plane. But $\mathbb{C}$ has something more: it is a *field*, so you can *multiply* and *divide* two complex numbers.

So we restrict our attention to differentiable functions called *holomorphic functions*. It turns out that the multiplication on $\mathbb{C}$ makes all the difference. The primary theme in what follows is that holomorphic functions are *really, really nice*, and that knowing tiny amounts of data about the function can determine all its values.

The two main highlights of this chapter, from which all other results are more or less corollaries:

- Contour integrals of loops are always zero.
- A holomorphic function is essentially given by its Taylor series; in particular, single-differentiable implies infinitely differentiable. Thus, holomorphic functions behave quite like polynomials.

Some of the resulting corollaries:

- It'll turn out that knowing just the values of a holomorphic function on the boundary of the unit circle will tell you the values in its interior.
- Knowing just the values of the function at $1, \frac{1}{2}, \frac{1}{3}, \dots$ are enough to determine the whole function!
- Bounded holomorphic functions $\mathbb{C} \rightarrow \mathbb{C}$ must be constant
- And more...

As [Pu02] writes: “Complex analysis is the good twin and real analysis is the evil one: beautiful formulas and elegant theorems seem to blossom spontaneously in the complex domain, while toil and pathology rule the reals”.

§14.2 Complex differentiation

Prototypical example for this section: Polynomials are holomorphic; $\bar{z}$ is not.

Let $f : U \rightarrow \mathbb{C}$ be a complex function. Then for some $z_0 \in U$, we define the **derivative** at z_0 to be

$$\lim_{h \rightarrow 0} \frac{f(z_0 + h) - f(z_0)}{h}.$$

Note that this limit may not exist; when it does we say f is **differentiable** at z_0 .

What do I mean by a “complex” limit $h \rightarrow 0$? It’s what you might expect: for every $\varepsilon > 0$ there should be a $\delta > 0$ such that

$$0 < |h| < \delta \implies \left| \frac{f(z_0 + h) - f(z_0)}{h} - L \right| < \varepsilon.$$

If you like topology, you are encouraged to think of this in terms of neighborhoods in the complex plane. (This is why we require U to be open: it makes it possible to take δ -neighborhoods in it.)


But note that having a complex derivative is actually much stronger than a real function having a derivative. In the real line, h can only approach zero from below and above, and for the limit to exist we need the “left limit” to equal the “right limit”. But the complex numbers form a *plane*: h can approach zero from many directions, and we need all the limits to be equal.

Example 14.2.1 (Important: conjugation is *not* holomorphic)

Let $f(z) = \bar{z}$ be complex conjugation, $f : \mathbb{C} \rightarrow \mathbb{C}$. This function, despite its simple nature, is not holomorphic! Indeed, at $z = 0$ we have,

$$\frac{f(h) - f(0)}{h} = \frac{\bar{h}}{h}.$$

This does not have a limit as $h \rightarrow 0$, because depending on “which direction” we approach zero from we have different values.



If a function $f : U \rightarrow \mathbb{C}$ is complex differentiable at all the points in its domain it is called **holomorphic**. In the special case of a holomorphic function with domain $U = \mathbb{C}$,

we call the function **entire**.¹

Example 14.2.2 (Examples of holomorphic functions)

In all the examples below, the derivative of the function is the same as in their real analogues (e.g. the derivative of e^z is e^z).

- (a) Any polynomial $z \mapsto z^n + c_{n-1}z^{n-1} + \cdots + c_0$ is holomorphic.
- (b) The complex exponential $\exp : x + yi \mapsto e^x(\cos y + i \sin y)$ can be shown to be holomorphic.
- (c) $\sin$ and $\cos$ are holomorphic when extended to the complex plane by $\cos z = \frac{e^{iz} + e^{-iz}}{2}$ and $\sin z = \frac{e^{iz} - e^{-iz}}{2i}$.
- (d) As usual, the sum, product, chain rules and so on apply, and hence **sums, products, nonzero quotients, and compositions of holomorphic functions are also holomorphic**.

You are welcome to try and prove these results, but I won't bother to do so.

§14.3 Contour integrals

Prototypical example for this section: $\oint_{\gamma} z^m dz$ around the unit circle.

In the real line we knew how to integrate a function across a line segment $[a, b]$: essentially, we'd "follow along" the line segment adding up the values of f we see to get some area. Unlike in the real line, in the complex plane we have the power to integrate over arbitrary paths: for example, we might compute an integral around a unit circle. A contour integral lets us formalize this.

First of all, if $f : \mathbb{R} \rightarrow \mathbb{C}$ and $f(t) = u(t) + iv(t)$ for $u, v \in \mathbb{R}$, we can define an integral $\int_a^b$ by just adding the real and imaginary parts:

$$\int_a^b f(t) dt = \left(\int_a^b u(t) dt \right) + i \left(\int_a^b v(t) dt \right).$$

Now let $\alpha : [a, b] \rightarrow \mathbb{C}$ be a path, thought of as a complex differentiable² function. Such a path is called a **contour**, and we define its **contour integral** by

$$\oint_{\alpha} f(z) dz = \int_a^b f(\alpha(t)) \cdot \alpha'(t) dt.$$

You can almost think of this as a u -substitution (which is where the α' comes from). In particular, it turns out this integral does not depend on how α is "parametrized": a circle given by

$$[0, 2\pi] \rightarrow \mathbb{C} : t \mapsto e^{it}$$

and another circle given by

$$[0, 1] \rightarrow \mathbb{C} : t \mapsto e^{2\pi it}$$

¹Sorry, I know the word "holomorphic" sounds so much cooler. I'll try to do things more generally for that sole reason.

²This isn't entirely correct here: you want the path α to be continuous and mostly differentiable, but you allow a finite number of points to have "sharp bends"; in other words, you can consider paths which are combinations of n smooth pieces. But for this we also require that α has "bounded length".

and yet another circle given by

$$[0, 1] \rightarrow \mathbb{C} : t \mapsto e^{2\pi it^5}$$

will all give the same contour integral, because the paths they represent have the same geometric description: “run around the unit circle once”.

In what follows I try to use α for general contours and γ in the special case of loops. Let’s see an example of a contour integral.

Theorem 14.3.1

Take $\gamma : [0, 2\pi] \rightarrow \mathbb{C}$ to be the unit circle specified by

$$t \mapsto e^{it}.$$

Then for any integer m , we have


$$\oint_{\gamma} z^m dz = \begin{cases} 2\pi i & m = -1 \\ 0 & \text{otherwise} \end{cases}$$

Proof. The derivative of e^{it} is ie^{it} . So, by definition the answer is the value of

$$\begin{aligned} \int_0^{2\pi} (e^{it})^m \cdot (ie^{it}) dt &= \int_0^{2\pi} i(e^{it})^{1+m} dt \\ &= i \int_0^{2\pi} \cos[(1+m)t] + i \sin[(1+m)t] dt \\ &= - \int_0^{2\pi} \sin[(1+m)t] dt + i \int_0^{2\pi} \cos[(1+m)t] dt. \end{aligned}$$

This is now an elementary calculus question. One can see that this equals $2\pi i$ if $m = -1$ and otherwise the integrals vanish. $\square$

Let me try to explain why this intuitively ought to be true for $m = 0$. In that case we just have $\oint_{\gamma} 1 dz$. So as the integral walks around the unit circle, it “sums up” all the tangent vectors at every point (that’s the direction it’s walking in), multiplied by 1. And given the nice symmetry of the circle, it should come as no surprise that everything cancels out. The theorem says that even if we multiply by z^m for $m \neq -1$, we get the same cancellation.



Definition 14.3.2. Given $\alpha : [0, 1] \rightarrow \mathbb{C}$, we denote by $\bar{\alpha}$ the “backwards” contour $\bar{\alpha}(t) = \alpha(1 - t)$.


Question 14.3.3. What's the relation between $\oint_{\alpha} f dz$ and $\oint_{\bar{\alpha}} f dz$? Prove it.

This might seem a little boring. Things will get really cool really soon, I promise.

§14.4 Cauchy-Goursat theorem


Prototypical example for this section: $\oint_{\gamma} z^m dz = 0$ for $m \geq 0$. But if $m < 0$, Cauchy's theorem does not apply.

Let $\Omega \subseteq \mathbb{C}$ be simply connected (for example, $\Omega = \mathbb{C}$), and consider two paths α, β with the same start and end points.



What's the relation between $\oint_{\alpha} f(z) dz$ and $\oint_{\beta} f(z) dz$? You might expect there to be some relation between them, considering that the space Ω is simply connected. But you probably wouldn't expect there to be *much* of a relation.

As a concrete example, let $\Psi : \mathbb{C} \rightarrow \mathbb{C}$ be the function $z \mapsto z - \operatorname{Re}[z]$ (for example, $\Psi(2015 + 3i) = 3i$). Let's consider two paths from -1 to 1 . Thus β just walking along the real axis, and α which follows an upper semicircle.



Obviously $\oint_{\beta} \Psi(z) dz = 0$. But heaven knows what $\oint_{\alpha} \Psi(z) dz$ is supposed to equal. We can compute it now just out of non-laziness. If you like, you are welcome to compute it yourself (it's a little annoying but not hard). If I myself didn't mess up, it is

$$\oint_{\alpha} \Psi(z) dz = - \oint_{\bar{\alpha}} \Psi(z) dz = - \int_0^{\pi} (i \sin(t)) \cdot i e^{it} dt = \frac{1}{2} \pi i$$

which in particular is not zero.

But somehow Ψ is not a really natural function. It's not respecting any of the nice, continuous structure of $\mathbb{C}$ since it just rudely lops off the real part of its inputs. More precisely,

Question 14.4.1. Show that $\Psi(z) = z - \operatorname{Re}[z]$ is not holomorphic. (Hint: $\bar{z}$ is not holomorphic.)

Now here's a miracle: for holomorphic functions, the two integrals are *always equal*. Equivalently, (by considering α followed by $\bar{\beta}$) contour integrals of loops are always zero. This is the celebrated Cauchy-Goursat theorem (also called the Cauchy integral theorem, but later we'll have a "Cauchy Integral Formula" so blah).

Theorem 14.4.2 (Cauchy-Goursat theorem)

Let γ be a loop, and $f : \Omega \rightarrow \mathbb{C}$ a holomorphic function where Ω is open in $\mathbb{C}$ and simply connected. Then

$$\oint_{\gamma} f(z) dz = 0.$$

Remark 14.4.3 (Sanity check). This might look surprising considering that we saw $\oint_{\gamma} z^{-1} dz = 2\pi i$ earlier. The subtlety is that z^{-1} is not even defined at $z = 0$. On the other hand, the function $\mathbb{C} \setminus \{0\} \rightarrow \mathbb{C}$ by $z \mapsto \frac{1}{z}$ is holomorphic! The defect now is that $\Omega = \mathbb{C} \setminus \{0\}$ is not simply connected. So the theorem passes our sanity checks, albeit just barely.

The typical proof of Cauchy's Theorem assumes additionally that the partial derivatives of f are continuous and then applies the so-called Green's theorem. But it was Goursat who successfully proved the fully general theorem we've stated above, which assumed only that f was holomorphic. I'll only outline the proof, and very briefly. You can show that if $f : \Omega \rightarrow \mathbb{C}$ has an antiderivative $F : \Omega \rightarrow \mathbb{C}$ which is also holomorphic, and moreover Ω is simply connected, then you get a "fundamental theorem of calculus", a la

$$\int_{\alpha} f(z) dz = F(\alpha(b)) - F(\alpha(a))$$

where $\alpha : [a, b] \rightarrow \mathbb{C}$ is some path. So to prove Cauchy-Goursat, you just have to show this antiderivative F exists. Goursat works very hard to prove the result in the special case that γ is a triangle, and hence by induction for any polygon. Once he has the result for a rectangle, he uses this special case to construct the function F explicitly. Goursat then shows that F is holomorphic, completing the proof.

Anyways, the theorem implies that $\oint_{\gamma} z^m dz = 0$ when $m \geq 0$. So much for all our hard work earlier. But so far we've only played with circles. This theorem holds for *any* contour which is a loop. So what else can we do?

§14.5 Cauchy's integral theorem

We now present a stunning application of Cauchy-Goursat, a "representation theorem": essentially, it says that values of f inside a disk are determined by just the values on the boundary! In fact, we even write down the exact formula. As [Ya12] says, "any time a certain type of function satisfies some sort of representation theorem, it is likely that many more deep theorems will follow." Let's pull back the curtain:

Theorem 14.5.1 (Cauchy's integral formula)


Let $\gamma : [0, 2\pi] \rightarrow \mathbb{C}$ be a circle in the plane given by $t \mapsto Re^{it}$, which bounds a disk D . Suppose $f : U \rightarrow \mathbb{C}$ is holomorphic such that U contains the circle and its interior. Then for any point a in the interior of D , we have

$$f(a) = \frac{1}{2\pi i} \oint_{\gamma} \frac{f(z)}{z-a} dz.$$


Note that we don't require U to be simply connected, but the reason is pretty silly: we're only going to ever integrate f over D , which is an open disk, and hence the disk is simply connected anyways.

The presence of $2\pi i$, which you saw earlier in the form $\oint_{\text{circle}} z^{-1} dz$, is no accident. In fact, that's the central result we're going to use to prove the result.

Proof. There are several proofs out there, but I want to give the one that really draws out the power of Cauchy's theorem. Here's the picture we have: there's a point a sitting inside a circle γ , and we want to get our hands on the value $f(a)$.



We're going to do a trick: construct a **keyhole contour** $\Gamma_{\delta,\varepsilon}$ which has an outer circle γ , plus an inner circle $\overline{\gamma}_{\varepsilon}$, which is a circle centered at a with radius ε , running clockwise (so that γ_{ε} runs counterclockwise). The "width" of the corridor is δ . See picture:



Hence $\Gamma_{\delta,\varepsilon}$ consists of four smooth curves.

Question 14.5.2. Draw a *simply connected* open set Ω which contains the entire $\Gamma_{\delta,\varepsilon}$ but does not contain the point a .

Hence, the function $\frac{f(z)}{z-a}$ manages to be holomorphic on all of Ω . Thus Cauchy's theorem applies and tells us that

$$0 = \oint_{\Gamma_{\delta,\varepsilon}} \frac{f(z)}{z-a} dz.$$

As we let $\delta \rightarrow 0$, the two walls of the keyhole will cancel each other (because f is continuous, and the walls run in opposite directions). So taking the limit as $\delta \rightarrow 0$, we

are left with just γ and γ_ε , which (taking again orientation into account) gives

$$\oint_{\gamma} \frac{f(z)}{z-a} dz = - \oint_{\gamma_\varepsilon} \frac{f(z)}{z-a} dz = \oint_{\gamma_\varepsilon} \frac{f(z)}{z-a} dz.$$

Thus **we've managed to replace γ with a much smaller circle γ_ε centered around a** , and the rest is just algebra.

To compute the last quantity, write

$$\begin{aligned} \oint_{\gamma_\varepsilon} \frac{f(z)}{z-a} dz &= \oint_{\gamma_\varepsilon} \frac{f(z) - f(a)}{z-a} dz + f(a) \cdot \oint_{\gamma_\varepsilon} \frac{1}{z-a} dz \\ &= \oint_{\gamma_\varepsilon} \frac{f(z) - f(a)}{z-a} dz + 2\pi i f(a). \end{aligned}$$

where we've used Theorem 14.3.1. Thus, all we have to do is show that

$$\oint_{\gamma_\varepsilon} \frac{f(z) - f(a)}{z-a} dz = 0.$$

For this we can basically use the weakest bound possible, the so-called *ML lemma* which I'll just cite without proof: it just says "bound the function everywhere by its maximum".

Lemma 14.5.3 (*ML estimation lemma*)

Let f be a holomorphic function and α a path. Suppose $M = \max_{z \text{ on } \alpha} |f(z)|$, and let L be the length of α . Then

$$\left| \oint_{\alpha} f(z) dz \right| \leq ML.$$

(This is straightforward to prove if you know the definition of length: $L = \int_a^b |\alpha'(t)| dt$, where $\alpha : [a, b] \rightarrow \mathbb{C}$.)

Anyways, as $\varepsilon \rightarrow 0$, the quantity $\frac{f(z)-f(a)}{z-a}$ just approaches $f'(a)$, and so for small enough ε (i.e. z close to a) there's some upper bound M . Yet the length of γ_ε is just the circumference $2\pi\varepsilon$. So the *ML lemma* says that

$$\left| \oint_{\gamma_\varepsilon} \frac{f(z) - f(a)}{z-a} dz \right| \leq 2\pi\varepsilon \cdot M \rightarrow 0$$

as desired. □

§14.6 Holomorphic functions are analytic

Prototypical example for this section: Imagine a formal series $\sum_k c_k x^k$!

In the setup of the previous problem, we have a circle $\gamma : [0, 2\pi] \rightarrow \mathbb{C}$ and a holomorphic function $U \rightarrow \mathbb{C}$ which contains the disk D . We can write

$$\begin{aligned} f(a) &= \frac{1}{2\pi i} \oint_{\gamma} \frac{f(z)}{z-a} dz \\ &= \frac{1}{2\pi i} \oint_{\gamma} \frac{f(z)/z}{1 - \frac{a}{z}} dz \\ &= \frac{1}{2\pi i} \oint_{\gamma} f(z)/z \cdot \sum_{k \geq 0} \left(\frac{a}{z}\right)^k dz \end{aligned}$$

You can prove (using the so-called Weierstrass M-test) that the summation order can be switched:

$$\begin{aligned}
 &= \frac{1}{2\pi i} \sum_{k \geq 0} \oint_{\gamma} \frac{f(z)}{z} \cdot \left(\frac{a}{z}\right)^k dz \\
 &= \frac{1}{2\pi i} \sum_{k \geq 0} \oint_{\gamma} a^k \cdot \frac{f(z)}{z^{k+1}} dz \\
 &= \sum_{k \geq 0} \left(\frac{1}{2\pi i} \oint_{\gamma} \frac{f(z)}{z^{k+1}} dz \right) a^k
 \end{aligned}$$

Letting $c_k = \frac{1}{2\pi i} \oint_{\gamma} \frac{f(z)}{z^{k+1}} dz$, and noting this is independent of a , this is

$$= \sum_{k \geq 0} c_k a^k$$

and that's the miracle: holomorphic functions are given by a Taylor series! This is one of the biggest results in complex analysis. Moreover, we also know that

$$c_k = \frac{f^{(k)}(0)}{k!}$$

from AP Calculus (if you don't see this, just take the derivative k times), and this gives us $f^{(k)}(0) = k! \cdot c_k$.

Naturally, we can do this with any circle (not just one centered at zero). So let's state the full result below, with arbitrary center p .

Theorem 14.6.1 (Cauchy's differentiation formula)

Let $f : U \rightarrow \mathbb{C}$ be a holomorphic function and let D be a disk centered at point p bounded by a circle γ . Suppose D is contained inside U . Then f is given everywhere in D by a Taylor series

$$f(z) = c_0 + c_1(z - p) + c_2(z - p)^2 + \dots$$

where

$$c_k = \frac{f^{(k)}(p)}{k!} = \frac{1}{2\pi i} \oint_{\gamma} \frac{f(w - p)}{(w - p)^{k+1}} dw$$

In particular,

$$f^{(k)}(p) = k! c_k = \frac{k!}{2\pi i} \oint_{\gamma} \frac{f(w - p)}{(w - p)^{k+1}} dw.$$

Most importantly,

Over any disk, a holomorphic function is given exactly by a Taylor series.

This establishes a result we stated at the beginning of the chapter: that a function being complex differentiable once means it is not only infinitely differentiable, but in fact equal to its Taylor series.

I should maybe emphasize a small subtlety of the result: the Taylor series centered at p is only valid in a disk centered at p which lies entirely in the domain U . If $U = \mathbb{C}$ this is no issue, since you can make the disk big enough to accommodate any point you want.

It's more subtle in the case that U is, for example, a square; you can't cover the entire square with a disk centered at some point without going outside the square. However, since U is open we can at any rate at least find some neighborhood for which the Taylor series is correct – in stark contrast to the real case. Indeed, as you'll see in the problems, the existence of a Taylor series is incredibly powerful.

§14.7 Problems to think about

These aren't olympiad problems, but I think they're especially nice! In the next complex analysis chapter we'll see some more nice applications.

The first few results are the most important.



Problem 14A* (Liouville's theorem). Let $f : \mathbb{C} \rightarrow \mathbb{C}$ be an entire function. Suppose that $|f(z)| < 1000$ for all complex numbers z . Prove that f is a constant function.

Problem 14B* (Zeros are isolated). An **isolated set** in the complex plane is a set of points S such that around each point in S , one can draw a neighborhood not intersecting any other point of S .

Show that the zero set of any nonzero holomorphic function $f : U \rightarrow \mathbb{C}$ is an isolated set, unless there exists a neighborhood of f which is identically zero.



Problem 14C* (Identity theorem). Let $f, g : U \rightarrow \mathbb{C}$ be holomorphic, and assume that U is connected. Prove that if f and g agree on some small neighborhood, then $f = g$.



Figure 14.2: Complex functions are unbounded. (Source: [Ge].)

Problem 14D[†] (Maximums Occur On Boundaries). Let $f : U \rightarrow \mathbb{C}$ be holomorphic, let $Y \subseteq U$ be compact, and let ∂Y be boundary³ of Y . Show that

$$\max_{z \in Y} |f(z)| = \max_{z \in \partial Y} |f(z)|.$$

In other words, the maximum values of $|f|$ occur on the boundary. (Such maximums exist by compactness.)

Problem 14E (Harvard quals). Let $f : \mathbb{C} \rightarrow \mathbb{C}$ be a nonconstant entire function. Prove that $f(\mathbb{C})$ is dense in $\mathbb{C}$. (In fact, a much stronger result is true: Little Picard's theorem says that the image of a nonconstant entire function omits at most one point.)

³ The boundary ∂Y is the set of points p such that no open neighborhood of p is contained in Y . It is also a compact set if Y is compact.

15 Meromorphic functions

§15.1 The second nicest functions on earth

If holomorphic functions are like polynomials, then *meromorphic* functions are like rational functions. Basically, a meromorphic function is a function of the form $\frac{A(z)}{B(z)}$ where $A, B : U \rightarrow \mathbb{C}$ are holomorphic and B is not zero. The most important example of a meromorphic function is $\frac{1}{z}$.

We are going to see that meromorphic functions behave like “almost-holomorphic” functions. Specifically, a meromorphic function A/B will be holomorphic at all points except the zeros of B (called *poles*). By the identity theorem, there cannot be too many zeros of B ! So meromorphic functions can be thought of as “almost holomorphic” (like $\frac{1}{z}$, which is holomorphic everywhere but the origin). We saw that

$$\frac{1}{2\pi i} \oint_{\gamma} \frac{1}{z} dz = 1$$

for $\gamma(t) = e^{it}$ the unit circle. We will extend our results on contours to such situations.

It turns out that, instead of just getting $\oint_{\gamma} f(z) dz = 0$ like we did in the holomorphic case, the contour integrals will actually be used to *count the number of poles* inside the loop γ . It’s ridiculous, I know.

§15.2 Meromorphic functions

Prototypical example for this section: $\frac{1}{z}$, with a pole of order 1 and residue 1 at $z = 0$.

Let U be an open subset of $\mathbb{C}$ again.

Definition 15.2.1. A function $f : U \rightarrow \mathbb{C}$ is **meromorphic** if it can be expressed as A/B for $A, B : U \rightarrow \mathbb{C}$ holomorphic, with B not identically zero in any neighborhood.

Let’s see how this function f behaves. If $z \in U$ has $B(z) \neq 0$, then in some small neighborhood the function B isn’t zero at all, and thus A/B is in fact *holomorphic*; thus f is holomorphic at z . (Concrete example: $\frac{1}{z}$ is holomorphic in any disk not containing 0.)

On the other hand, suppose $p \in U$ has $B(p) = 0$: without loss of generality, $p = 0$ to ease notation. By using the Taylor series at $p = 0$ we can put

$$B(z) = c_k z^k + c_{k+1} z^{k+1} + \dots$$

with $c_k \neq 0$ (certainly some coefficient is nonzero since B is not identically zero!). Then we can write

$$\frac{1}{B(z)} = \frac{1}{z^k} \cdot \frac{1}{c_k + c_{k+1}z + \dots}.$$

But the fraction on the right is a holomorphic function in this neighborhood! So all that’s happened is that we have an extra z^{-k} kicking around.

This gives us an equivalent way of viewing meromorphic functions:

Definition 15.2.2. Let $f : U \rightarrow \mathbb{C}$ as usual. A **meromorphic** function is a function which is holomorphic on U except at an isolated set S of points (meaning it is holomorphic as a function $U \setminus S \rightarrow \mathbb{C}$). For each $p \in S$, called a **pole** of f , the function f must admit a **Laurent series**, meaning that

$$f(z) = \frac{c_{-m}}{(z-p)^m} + \frac{c_{-m+1}}{(z-p)^{m-1}} + \cdots + \frac{c_{-1}}{z-p} + c_0 + c_1(z-p) + \cdots$$

for all z in some neighborhood of p , other than $z = p$. Here m is a positive integer (and $c_{-m} \neq 0$).

Note that the trailing end *must* terminate. By “isolated set”, I mean that we can draw neighborhoods around each pole in S , in such a way that no two neighborhoods intersect.

Example 15.2.3 (Example of a meromorphic function)

Consider the function

$$\frac{z+1}{\sin z}.$$

It is meromorphic, because it is holomorphic everywhere except at the zeros of $\sin z$. At each of these points we can put a Laurent series: for example at $z = 0$ we have

$$\begin{aligned} \frac{z+1}{\sin z} &= (z+1) \cdot \frac{1}{z - \frac{z^3}{3!} + \frac{z^5}{5!} - \cdots} \\ &= \frac{1}{z} \cdot \frac{z+1}{1 - \left(\frac{z^2}{3!} - \frac{z^4}{5!} + \frac{z^6}{7!} - \cdots \right)} \\ &= \frac{1}{z} \cdot (z+1) \sum_{k \geq 0} \left(\frac{z^2}{3!} - \frac{z^4}{5!} + \frac{z^6}{7!} - \cdots \right)^k. \end{aligned}$$

If we expand out the horrible sum (which I won't do), then you get $\frac{1}{z}$ times a perfectly fine Taylor series, i.e. a Laurent series.

Abuse of Notation 15.2.4. We'll often say something like “consider the function $f : \mathbb{C} \rightarrow \mathbb{C}$ by $z \mapsto \frac{1}{z}$ ”. Of course this isn't completely correct, because f doesn't have a value at $z = 0$. If I was going to be completely rigorous I would just set $f(0) = 2015$ or something and move on with life, but for all intents let's just think of it as “undefined at $z = 0$ ”.

Why don't I just write $g : \mathbb{C} \setminus \{0\} \rightarrow \mathbb{C}$? The reason I have to do this is that it's still important for f to remember it's “trying” to be holomorphic on $\mathbb{C}$, even if isn't assigned a value at $z = 0$. As a function $\mathbb{C} \setminus \{0\} \rightarrow \mathbb{C}$ the function $\frac{1}{z}$ is actually holomorphic.

Remark 15.2.5. I have shown that any function $A(z)/B(z)$ has this characterization with poles, but an important result is that the converse is true too: if $f : U \setminus S \rightarrow \mathbb{C}$ is holomorphic for some isolated set S , and moreover f admits a Laurent series at each point in S , then f can be written as a rational quotient of holomorphic functions. I won't prove this here, but it is good to be aware of.

Definition 15.2.6. Let p be a pole of a meromorphic function f , with Laurent series

$$f(z) = \frac{c_{-m}}{(z-p)^m} + \frac{c_{-m+1}}{(z-p)^{m-1}} + \cdots + \frac{c_{-1}}{z-p} + c_0 + c_1(z-p) + \cdots$$

The integer m is called the **order** of the pole. A pole of order 1 is called a **simple pole**.

We also give the coefficient c_{-1} a name, the **residue** of f at p , which we write $\text{Res}(f; p)$.

The order of a pole tells you how “bad” the pole is. The order of a pole is the “opposite” concept of the **multiplicity** of a **zero**. If f has a pole at zero, then its Taylor series near $z = 0$ might look something like

$$f(z) = \frac{1}{z^5} + \frac{8}{z^3} - \frac{2}{z^2} + \frac{4}{z} + 9 - 3z + 8z^2 + \dots$$

and so f has a pole of order five. By analogy, if g has a zero at $z = 0$, it might look something like

$$g(z) = 3z^3 + 2z^4 + 9z^5 + \dots$$

and so g has a zero of multiplicity three. These orders are additive: $f(z)g(z)$ still has a pole of order $5 - 3 = 2$, but $f(z)g(z)^2$ is completely patched now, and in fact has a **simple zero** now (that is, a zero of degree 1).

Question 15.2.7. Prove that orders are additive as described above. (This is obvious once you understand that you are multiplying Taylor/Laurent series.)

Metaphorically, poles can be thought of as “negative zeros”. We can now give many more examples.

Example 15.2.8 (Examples of meromorphic functions)

- (a) Any holomorphic function is a meromorphic function which happens to have no poles. Stupid, yes.
- (b) The function $\mathbb{C} \rightarrow \mathbb{C}$ by $z \mapsto 100z^{-1}$ for $z \neq 0$ but undefined at zero is a meromorphic function. Its only pole is at zero, which has order 1 and residue 100.
- (c) The function $\mathbb{C} \rightarrow \mathbb{C}$ by $z \mapsto z^{-3} + z^2 + z^9$ is also a meromorphic function. Its only pole is at zero, and it has order 3, and residue 0.
- (d) The function $\mathbb{C} \rightarrow \mathbb{C}$ by $z \mapsto \frac{e^z}{z^2}$ is meromorphic, with the Laurent series at $z = 0$ given by

$$\frac{e^z}{z^2} = \frac{1}{z^2} + \frac{1}{z} + \frac{1}{2} + \frac{z}{6} + \frac{z^2}{24} + \frac{z^3}{120} + \dots$$

Hence the pole $z = 0$ has order 2 and residue 1.

Example 15.2.9 (A rational meromorphic functions)

Consider the function $\mathbb{C} \rightarrow \mathbb{C}$ given by

$$\begin{aligned} z \mapsto \frac{z^4 + 1}{z^2 - 1} &= z^2 + 1 + \frac{2}{(z-1)(z+1)} \\ &= z^2 + 1 + \frac{1}{z-1} \cdot \frac{1}{1 + \frac{z-1}{2}} \\ &= \frac{2}{z-1} + \frac{3}{2} + \frac{9}{4}(z-1) + \frac{7}{8}(z-1)^2 - \dots \end{aligned}$$

It has a pole of order 1 and residue 2 at $z = 1$. (It also has a pole of order 1 at $z = -1$; you are invited to compute the residue.)

Example 15.2.10 (Function with infinitely many poles)

The function $\mathbb{C} \rightarrow \mathbb{C}$ by

$$z \mapsto \frac{1}{\sin(z)}$$

has infinitely many poles: the numbers $z = 2\pi k$, where k is an integer. Let's compute the Laurent series at just $z = 0$:

$$\begin{aligned} \frac{1}{\sin(2\pi z)} &= \frac{1}{\frac{z}{1!} - \frac{z^3}{3!} + \frac{z^5}{5!} - \dots} \\ &= \frac{1}{z} \cdot \frac{1}{1 - \left(\frac{z^2}{3!} - \frac{z^4}{5!} + \dots\right)} \\ &= \frac{1}{z} \sum_{k \geq 0} \left(\frac{z^2}{3!} - \frac{z^4}{5!} + \dots\right)^k. \end{aligned}$$

which is a Laurent series, though I have no clue what the coefficients are. You can at least see the residue; the constant term of that huge sum is 1, so the residue is 1. Also, the pole has order 1.

The Laurent series, if it exists, is unique (as you might have guessed), and by our result on holomorphic functions it is actually valid for *any* disk centered at p (minus the point p). The part $\frac{c_{-1}}{z-p} + \dots + \frac{c_{-m}}{(z-p)^m}$ is called the **principal part**, and the rest of the series $c_0 + c_1(z-p) + \dots$ is called the **analytic part**.

§15.3 Winding numbers and the residue theorem

Recall that for a counterclockwise circle γ and a point p inside it, we had

$$\oint_{\gamma} (z-p)^m dz = \begin{cases} 0 & m \neq -1 \\ 2\pi i & m = -1 \end{cases}$$

where m is an integer. One can extend this result to in fact show that $\oint_{\gamma} (z-p)^m dz = 0$ for *any* loop γ , where $m \neq -1$. So we associate a special name for the nonzero value at $m = -1$.

Definition 15.3.1. For a point $p \in \mathbb{C}$ and a loop γ not passing through it, we define the **winding number**, denoted $\mathbf{I}(p, \gamma)$, by

$$\mathbf{I}(\gamma, p) = \frac{1}{2\pi i} \oint_{\gamma} \frac{1}{z-p} dz$$

For example, by our previous results we see that if γ is a circle, we have

$$\mathbf{I}(\text{circle}, p) = \begin{cases} 1 & p \text{ inside the circle} \\ 0 & p \text{ outside the circle.} \end{cases}$$

If you've read the chapter on fundamental groups, then this is just the fundamental group associated to $\mathbb{C} \setminus \{p\}$. In particular, the winding number is always an integer (the proof of this requires the complex logarithm, so we omit it here). In the simplest case the winding numbers are either 0 or 1.

Definition 15.3.2. We say a loop γ is **regular** if $\mathbf{I}(p, \gamma) = 1$ for all points p in the interior of γ (for example, if γ is a counterclockwise circle).

With all these ingredients we get a stunning generalization of the Cauchy-Goursat theorem:

Theorem 15.3.3 (Cauchy's residue theorem)

Let $f : \Omega \rightarrow \mathbb{C}$ be meromorphic, where Ω is simply connected. Then for any loop γ not passing through any of its poles, we have

$$\frac{1}{2\pi i} \oint_{\gamma} f(z) dz = \sum_{\text{pole } p} \mathbf{I}(\gamma, p) \operatorname{Res}(f; p).$$

In particular, if γ is regular then the contour integral is the sum of all the residues, in the form

$$\frac{1}{2\pi i} \oint_{\gamma} f(z) dz = \sum_{\substack{\text{pole } p \\ \text{inside } \gamma}} \operatorname{Res}(f; p).$$

Question 15.3.4. Verify that this result coincides with what you expect when you integrate $\oint_{\gamma} cz^{-1} dz$ for γ a counter-clockwise circle.

The proof from here is not really too impressive – the “work” was already done in our statements about the winding number.

Proof. Let the poles with nonzero winding number be $p_1, \dots, p_k$ (the others do not affect the sum).¹ Then we can write f in the form

$$f(z) = g(z) + \sum_{i=1}^k P_i \left(\frac{1}{z - p_i} \right)$$

where $P_i \left(\frac{1}{z - p_i} \right)$ is the principal part of the pole p_i . (For example, if $f(z) = \frac{z^3 - z + 1}{z(z+1)}$ we would write $f(z) = (z - 1) + \frac{1}{z} - \frac{1}{1+z}$.)

The point of doing so is that the function g is holomorphic (we’ve removed all the “bad” parts), so

$$\oint_{\gamma} g(z) dz = 0$$

by Cauchy-Goursat.

On the other hand, if $P_i(x) = c_1x + c_2x^2 + \dots + c_dx^d$ then

$$\begin{aligned} \oint_{\gamma} P_i \left(\frac{1}{z - p_i} \right) dz &= \oint_{\gamma} c_1 \cdot \left(\frac{1}{z - p_i} \right) dz + \oint_{\gamma} c_2 \cdot \left(\frac{1}{z - p_i} \right)^2 dz + \dots \\ &= c_1 \cdot \mathbf{I}(\gamma, p_i) + 0 + 0 + \dots \\ &= \mathbf{I}(\gamma, p_i) \operatorname{Res}(f; p_i). \end{aligned}$$

which gives the conclusion. $\square$

¹ To show that there must be finitely many such poles: recall that all our contours $\gamma : [a, b] \rightarrow \mathbb{C}$ are in fact bounded, so there is some big closed disk D which contains all of γ . The poles outside D thus have winding number zero. Now we cannot have infinitely many poles inside the disk D , for D is compact and the set of poles is a closed and isolated set!

§15.4 Argument principle

One tricky application is as follows. Given a polynomial $P(x) = (x-a_1)^{e_1}(x-a_2)^{e_2} \dots (x-a_n)^{e_n}$, you might know that we have

$$\frac{P'(x)}{P(x)} = \frac{e_1}{x-a_1} + \frac{e_2}{x-a_2} + \dots + \frac{e_n}{x-a_n}.$$

The quantity P'/P is called the **logarithmic derivative**, as it is the derivative of $\log P$. This trick allows us to convert zeros of P into poles of P'/P with order 1; moreover the residues of these poles are the multiplicities of the roots.

In an analogous fashion, we can obtain a similar result for any meromorphic function f .

Proposition 15.4.1 (The logarithmic derivative)

Let $f : U \rightarrow \mathbb{C}$ be a meromorphic function. Then the logarithmic derivative f'/f is meromorphic as a function from U to $\mathbb{C}$; its zeros and poles are:

- (i) A pole at each zero of f whose residue is the multiplicity, and
- (ii) A pole at each pole of f whose residue is the negative of the pole's order.

Again, you can almost think of a pole as a zero of negative multiplicity. This spirit is exemplified below.

Proof. Dead easy with Taylor series. Let a be a zero/pole of f , and WLOG set $a = 0$ for convenience. We take the Taylor series at zero to get

$$f(z) = c_k z^k + c_{k+1} z^{k+1} + \dots$$

where $k < 0$ if 0 is a pole and $k > 0$ if 0 is a zero. Taking the derivative gives

$$f'(z) = k c_k z^{k-1} + (k+1) c_{k+1} z^k + \dots$$

Now look at f'/f ; with some computation, it equals

$$\frac{f'(z)}{f(z)} = \frac{1}{z} \frac{k c_k + (k+1) c_{k+1} z + \dots}{c_k + c_{k+1} z + \dots}.$$

So we get a simple pole at $z = 0$, with residue k . □

Using this trick you can determine the number of zeros and poles inside a regular closed curve, using the so-called Argument Principle.

Theorem 15.4.2 (Argument principle)

Let γ be a regular curve. Suppose $f : U \rightarrow \mathbb{C}$ is meromorphic inside and on γ , and none of its zeros or poles lie on γ . Then

$$\frac{1}{2\pi i} \oint_{\gamma} \frac{f'}{f} dz = Z - P$$

where Z is the number of zeros inside γ (counted with multiplicity) and P is the number of poles inside γ (again with multiplicity).

Proof. Immediate by applying Cauchy's residue theorem alongside the preceding proposition. In fact you can generalize to any curve γ via the winding number: the integral is

$$\frac{1}{2\pi i} \oint_{\gamma} \frac{f'}{f} dz = \sum_{\text{zero } z} \mathbf{I}(\gamma, z) - \sum_{\text{pole } p} \mathbf{I}(\gamma, p)$$

where the sums are with multiplicity. $\square$

Thus the Argument Principle allows one to count zeros and poles inside any region of choice.

Computers can use this to get information on functions whose values can be computed but whose behavior as a whole is hard to understand. Suppose you have a holomorphic function f , and you want to understand where its zeros are. Then just start picking various circles γ . Even with machine rounding error, the integral will be close enough to the true integer value that we can decide how many zeros are in any given circle. Numerical evidence for the Riemann Hypothesis (concerning the zeros of the Riemann zeta function) can be obtained in this way.

§15.5 Philosophy: why are holomorphic functions so nice?

All the fun we've had with holomorphic and meromorphic functions comes down to the fact that complex differentiability is such a strong requirement. It's a small miracle that $\mathbb{C}$, which *a priori* looks only like $\mathbb{R}^2$, is in fact a field. Moreover, $\mathbb{R}^2$ has the nice property that one can draw nontrivial loops (it's also true for real functions that $\int_a^a f dx = 0$, but this is not so interesting!), and this makes the theory much more interesting.

As another piece of intuition from Siu²: If you try to get (left) differentiable functions over *quaternions*, you find yourself with just linear functions.

§15.6 Problems to think about

Problem 15A (Fundamental theorem of algebra). Prove that if f is a nonzero polynomial of degree n then it has n roots.

Problem 15B[†] (Rouché's theorem). Let $f, g : U \rightarrow \mathbb{C}$ be holomorphic functions, where U contains the unit disk. Suppose that $|f(z)| > |g(z)|$ for all z on the unit circle. Prove that f and $f + g$ have the same number of zeros which lie strictly inside the unit circle (counting multiplicities).



Problem 15C. Prove that the integral

$$\int_{-\infty}^{\infty} \frac{\cos x}{x^2 + 1} dx$$

converges and determine its value.



Problem 15D*. Let $f : U \rightarrow \mathbb{C}$ be a nonconstant holomorphic function.

- (Open mapping theorem) Prove that $f(U)$ is open in $\mathbb{C}$.³
- (Maximum Modulus Principle) Show that $|f|$ cannot have a maximum over U . That is, show that for any $z \in U$, there is some $z' \in U$ such that $|f(z)| < |f(z')|$.

²Harvard professor

³Thus the image of *any* open set $V \subseteq U$ is open in $\mathbb{C}$ (by repeating the proof for $f|_V$).

16 Holomorphic square roots and logarithms

In this chapter we'll make sense of a holomorphic square root and logarithm. The main results are Theorem 16.3.2, Theorem 16.4.2, Corollary 16.5.1, and Theorem 16.5.2. If you like, you can read just these four results, and skip the discussion of how they came to be.

Let $f : U \rightarrow \mathbb{C}$ be a holomorphic function. A **holomorphic n th root** of f is a function $g : U \rightarrow \mathbb{C}$ such that $f(z) = g(z)^n$ for all $z \in U$. A **logarithm** of f is a function $g : U \rightarrow \mathbb{C}$ such that $f(z) = e^{g(z)}$ for all $z \in U$. The main question we'll try to figure out is: when do these exist? In particular, what if $f = \text{id}$?

§16.1 Motivation: square root of a complex number

To start us off, can we define $\sqrt{z}$ for any complex number z ?

The first obvious problem that comes up is that for any z , there are *two* numbers w such that $w^2 = z$. How can we pick one to use? For our ordinary square root function, we had a notion of “positive”, and so we simply took the positive root.


Let's expand on this: given $z = r(\cos \theta + i \sin \theta)$ (here $r \geq 0$) we should take the root to be

$$w = \sqrt{r}(\cos \alpha + i \sin \alpha).$$

such that $2\alpha \equiv \theta \pmod{2\pi}$; there are two choices for $\alpha \pmod{2\pi}$, differing by π .

For complex numbers, we don't have an obvious way to pick α . Nonetheless, perhaps we can also get away with an arbitrary distinction: let's see what happens if we just choose the α with $-\frac{1}{2}\pi < \alpha \leq \frac{1}{2}\pi$.


Pictured below are some points (in red) and their images (in blue) under this “upper-half” square root. The condition on α means we are forcing the blue points to lie on the right-half plane.




Here, $w_i^2 = z_i$ for each i , and we are constraining the w_i to lie in the right half of the complex plane. We see there is an obvious issue: there is a big discontinuity near

the points w_5 and w_7 ! The nearby point w_6 has been mapped very far away. This discontinuity occurs since the points on the negative real axis are at the “boundary”. For example, given -4 , we send it to $-2i$, but we have hit the boundary: in our interval $-\frac{1}{2}\pi \leq \alpha < \frac{1}{2}\pi$, we are at the very left edge.

The negative real axis that we must not touch is what we will later call a *branch cut*, but for now I call it a **ray of death**. It is a warning to the red points: if you cross this line, you will die! However, if we move the red circle just a little upwards (so that it misses the negative real axis) this issue is avoided entirely, and we get what seems to be a “nice” square root.



In fact, the ray of death is fairly arbitrary: it is the set of “boundary issues” that arose when we picked $-\frac{1}{2}\pi < \alpha \leq \frac{1}{2}\pi$. Suppose we instead insisted on the interval $0 \leq \alpha < \pi$; then the ray of death would be the *positive* real axis instead. The earlier circle we had now works just fine.



What we see is that picking a particular α -interval leads to a different set of edge cases, and hence a different ray of death. The only thing these rays have in common is their starting point of zero. In other words, given a red circle and a restriction of α , I can make a nice “square rooted” blue circle as long as the ray of death misses it.

So, what exactly is going on?

§16.2 Square roots of holomorphic functions


To get a picture of what's happening, we would like to consider a more general problem: let $f : U \rightarrow \mathbb{C}$ be holomorphic. Then we want to decide whether there is a $g : U \rightarrow \mathbb{C}$ such that

$$f(z) = g(z)^2.$$

Our previous discussion with $f = \text{id}$ tells us we cannot hope to achieve this for $U = \mathbb{C}$; there is a “half-ray” which causes problems. However, there are certainly functions $f : \mathbb{C} \rightarrow \mathbb{C}$ such that a g exists. As a simplest example, $f(z) = z^2$ should definitely have a square root!

Now let's see if we can fudge together a square root. Earlier, what we did was try to specify a rule to force one of the two choices at each point. This is unnecessarily strict. Perhaps we can do something like: start at a point in $z_0 \in U$, pick a square root w_0 of $f(z_0)$, and then try to “fudge” from there the square roots of the other points. What do I mean by fudge? Well, suppose z_1 is a point very close to z_0 , and we want to pick a square root w_1 of $f(z_1)$. While there are two choices, we also would expect w_0 to be close to w_1 . Unless we are highly unlucky, this should tell us which choice of w_1 to pick. (Stupid concrete example: if I have taken the square root $-4.12i$ of -17 and then ask you to continue this square root to -16 , which sign should you pick for $\pm 4i$?)

There are two possible ways we could get unlucky in the scheme above: first, if $w_0 = 0$, then we're sunk. But even if we avoid that, we have to worry that if we run a full loop in the complex plane, we might end up in a different place from where we started. For concreteness, consider the following situation, again with $f = \text{id}$:



We started at the point z_0 , with one of its square roots as w_0 . We then wound a full red circle around the origin, only to find that at the end of it, the blue arc is at a different place where it started!

The interval construction from earlier doesn't work either: no matter how we pick the interval for α , any ray of death must hit our red circle. The problem somehow lies with the fact that we have enclosed the very special point 0.

Nevertheless, we know that if we take $f(z) = z^2$, then we don't run into any problems with our “make it up as you go” procedure. So, what exactly is going on?

§16.3 Covering projections

By now, if you have read the part on algebraic topology, this should all seem quite familiar. The “fudging” procedure exactly describes the idea of a lifting.

More precisely, recall that there is a covering projection

$$(-)^2 : \mathbb{C} \setminus \{0\} \rightarrow \mathbb{C} \setminus \{0\}.$$

Let $V = \{z \in U \mid f(z) \neq 0\}$. For $z \in U \setminus V$, we already have the square root $g(z) = \sqrt{f(z)} = \sqrt{0} = 0$. So the burden is completing $g : V \rightarrow \mathbb{C}$.

Then essentially, what we are trying to do is construct a lifting g in the diagram

$$\begin{array}{ccc} & E = \mathbb{C} \setminus \{0\} & \\ \nearrow g & \downarrow p=(-)^2 & \\ V & \xrightarrow{f} & B = \mathbb{C} \setminus \{0\}. \end{array}$$

Our map p can be described as “winding around twice”. Our Theorem 22.2.5 now tells us that this lifting exists if and only if

$$f_* (\pi_1(V)) \subseteq p_* (\pi_1(E))$$

is a subset of the image of $\pi_1(E)$ by p . Since B and E are both punctured planes, we can identify them with S^1 .

Question 16.3.1. Show that the image under p is exactly $2\mathbb{Z}$ once we identify $\pi_1(B) = \mathbb{Z}$.

That means that for any loop γ in V , we need $f \circ \gamma$ to have an *even* winding number around $0 \in B$. This amounts to

$$\frac{1}{2\pi} \oint_{\gamma} \frac{f'}{f} dz \in 2\mathbb{Z}$$

since f has no poles.

Replacing 2 with n and carrying over the discussion gives the first main result.

Theorem 16.3.2 (Existence of holomorphic n th roots)

Let $f : U \rightarrow \mathbb{C}$ be holomorphic. Then f has a holomorphic n th root if and only if

$$\frac{1}{2\pi i} \oint_{\gamma} \frac{f'}{f} dz \in n\mathbb{Z}$$

for every contour γ in U .

§16.4 Complex logarithms

The multivalued nature of the complex logarithm comes from the fact that

$$\exp(z + 2\pi i) = \exp(z).$$

So if $e^w = z$, then any complex number $w + 2\pi ik$ is also a solution.

We can handle this in the same way as before: it amounts to a lifting of the following diagram.

$$\begin{array}{ccc} & & E = \mathbb{C} \\ & \nearrow g & \downarrow p=\exp \\ U & \xrightarrow{f} & B = \mathbb{C} \setminus \{0\} \end{array}$$

There is no longer a need to work with a separate V since:

Question 16.4.1. Show that if f has any zeros then g possibly can't exist.

In fact, the map $\exp : \mathbb{C} \rightarrow \mathbb{C} \setminus \{0\}$ is a universal cover, since $\mathbb{C}$ is simply connected. Thus, $p^*(\pi_1(\mathbb{C}))$ is *trivial*. So in addition to being zero-free, f cannot have any winding number around $0 \in B$ at all. In other words:

Theorem 16.4.2 (Existence of logarithms)

Let $f : U \rightarrow \mathbb{C}$ be holomorphic. Then f has a logarithm if and only if

$$\frac{1}{2\pi i} \oint_{\gamma} \frac{f'}{f} dz = 0$$

for every contour γ in U .

§16.5 Some special cases

The most common special case is

Corollary 16.5.1 (Nonvanishing functions from simply connected domains)

Let $f : \Omega \rightarrow \mathbb{C}$ be continuous, where Ω is simply connected. If $f(z) \neq 0$ for every $z \in \Omega$, then f has both a logarithm and holomorphic n th root.

Finally, let's return to the question of $f = \text{id}$ from the very beginning. What's the best domain U such that

$$\sqrt{-} : U \rightarrow \mathbb{C}$$

is well-defined? Clearly $U = \mathbb{C}$ cannot be made to work, but we can do almost as well. For note that the only zero of $f = \text{id}$ is at the origin. Thus if we want to make a logarithm exist, all we have to do is make an incision in the complex plane that renders it impossible to make a loop around the origin. The usual choice is to delete negative half of the real axis, our very first ray of death; we call this a **branch cut**, with **branch point** at $0 \in \mathbb{C}$ (the point which we cannot circle around). This gives

Theorem 16.5.2 (Branch cut functions)

There exist holomorphic functions

$$\log : \mathbb{C} \setminus (-\infty, 0] \rightarrow \mathbb{C}$$

$$\sqrt[n]{-} : \mathbb{C} \setminus (-\infty, 0] \rightarrow \mathbb{C}$$

satisfying the obvious properties.

There are many possible choices of such functions (n choices for the n th root and infinitely many for $\log$); a choice of such a function is called a **branch**. So this is what is meant by a “branch” of a logarithm.

The **principal branch** is the “canonical” branch, analogous to the way we arbitrarily pick the positive branch to define $\sqrt{-} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$. For $\log$, we take the w such that $e^w = z$ and the imaginary part of w lies in $(-\pi, \pi]$ (since we can shift by integer multiples of $2\pi i$). Often, authors will write $\text{Log } z$ to emphasize this choice.

§16.6 Problems to think about

Problem 16A. Show that a holomorphic function $f : U \rightarrow \mathbb{C}$ has a holomorphic logarithm if and only if it has a holomorphic n th root for every integer n .

Problem 16B. Show that the function $f : U \rightarrow \mathbb{C}$ by $z \mapsto z(z-1)$ has a holomorphic square root, where U is the entire complex plane minus the closed interval $[0, 1]$.



Quantum Algorithms

17 Quantum states and measurements	178
17.1 Bra-ket notation	178
17.2 The state space	179
17.3 Observations	179
17.4 Entanglement	182
17.5 Problems to think about	185
18 Quantum circuits	186
18.1 Classical logic gates	186
18.2 Reversible classical logic	187
18.3 Quantum logic gates	189
18.4 Deutsch-Jozsa algorithm	191
18.5 Problems to think about	192
19 Shor's algorithm	194
19.1 The classical (inverse) Fourier transform	194
19.2 The quantum Fourier transform	195
19.3 Shor's algorithm	197

17 Quantum states and measurements

In this chapter we'll explain how to set up quantum states using linear algebra. This will allow me to talk about quantum *circuits* in the next chapter, which will set the stage for Shor's algorithm.

I won't do very much physics (read: none at all). That is, I'll only state what the physical reality is in terms of linear algebras, and defer the philosophy of why this is true to your neighborhood "Philosophy of Quantum Mechanics" class (which is a "social science" class at MIT!).

§17.1 Bra-ket notation

Physicists have their own notation for vectors: whereas I previously used something like v , e_1 , and so on, in this chapter you'll see the infamous **bra-ket** notation: a vector will be denoted by $|\bullet\rangle$, where $\bullet$ is some variable name: unlike in math or Python, this can include numbers, symbols, Unicode characters, whatever you like. This is called a "ket". To pay a homage to physicists everywhere, we'll use this notation for this chapter too.

For us:

Definition 17.1.1. A **Hilbert space** H is a finite-dimensional complex inner product space.

Abuse of Notation 17.1.2. In actual physics, almost all Hilbert spaces used are infinite-dimensional, in which case there is the additional requirement that H be a complete metric space with respect to its norm. But finite-dimensional spaces will suffice for the purposes here (and automatically satisfy the completeness property).

If $\dim H = n$, then its orthonormal basis elements are often denoted

$$|0\rangle, |1\rangle, \dots, |n-1\rangle$$

(instead of e_i) and a generic element of H denoted by

$$|\psi\rangle, |\phi\rangle, \dots$$

and various other Greek letters.

Now for any $|\psi\rangle \in H$, we can consider the canonical dual element in $H^\vee$ (since H has an inner form), which we denote by $\langle\psi|$ (a "bra"). For example, if $\dim H = 2$ then we can write

$$|\psi\rangle = \begin{bmatrix} \alpha \\ \beta \end{bmatrix}$$

in an orthonormal basis, in which case

$$\langle\psi| = [\bar{\alpha} \quad \bar{\beta}].$$

We even can write dot products succinctly in this notation: if $|\phi\rangle = \begin{bmatrix} \gamma \\ \delta \end{bmatrix}$, then the dot product of $|\phi\rangle$ and $|\psi\rangle$ is given by

$$\langle\psi|\phi\rangle = [\bar{\alpha} \quad \bar{\beta}] \begin{bmatrix} \gamma \\ \delta \end{bmatrix} = \bar{\alpha}\gamma + \bar{\beta}\delta.$$

So we will use the notation $\langle\psi|\phi\rangle$ instead of the more mathematical $\langle|\psi\rangle, |\phi\rangle\rangle$. In particular, the squared norm of $|\psi\rangle$ is just $\langle\psi|\psi\rangle$. Concretely, for $\dim H = 2$ we have $\langle\psi|\psi\rangle = |\alpha|^2 + |\beta|^2$.

§17.2 The state space

If you think that's weird, well, it gets worse.

In classical computation, a bit is either 0 or 1. More generally, we can think of a classical space of n possible states $0, \dots, n-1$. Thus in the classical situation, the space of possible states is just a discrete set with n elements.

In quantum computation, a **qubit** is instead any *complex linear combination* of 0 and 1. To be precise, consider the normed complex vector space

$$H = \mathbb{C}^{\oplus 2}$$

and denote the orthonormal basis elements by $|0\rangle$ and $|1\rangle$. Then a *qubit* is a nonzero element $|\psi\rangle \in H$, so that it can be written in the form

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle$$

where α and β are not both zero. Typically, we normalize so that $|\psi\rangle$ has norm 1:

$$\langle\psi|\psi\rangle = 1 \iff |\alpha|^2 + |\beta|^2 = 1.$$

In particular, we can recover the “classical” situation with $|0\rangle \in H$ and $|1\rangle \in H$, but now we have some “intermediate” states, such as

$$\frac{1}{\sqrt{2}} (|0\rangle + |1\rangle).$$

Philosophically, what has happened is that:

Instead of allowing just the states $|0\rangle$ and $|1\rangle$, we allow any complex linear combination of them.

More generally, if $\dim H = n$, then the possible states are nonzero elements

$$c_0 |0\rangle + c_1 |1\rangle + \dots + c_{n-1} |n-1\rangle$$

which we usually normalize so that $|c_0|^2 + |c_1|^2 + \dots + |c_{n-1}|^2 = 1$.

§17.3 Observations

Prototypical example for this section: id corresponds to not making a measurement since all its eigenvalues are equal, but any operator with distinct eigenvalues will cause collapse.

If you think that's weird, well, it gets worse. First, some linear algebra:

Definition 17.3.1. Let V be a finite-dimensional inner product space. For a map $T : V \rightarrow V$, the following conditions are equivalent:

- $\langle Tx, y \rangle = \langle x, Ty \rangle$ for any $x, y \in V$.
- $T = T^\dagger$.

A map T satisfying these conditions is called **Hermitian**.

Question 17.3.2. Show that T is normal.

Thus, we know that T is diagonalizable with respect to the inner form, so for a suitable basis we can write it in an orthonormal basis as

$$T = \begin{bmatrix} \lambda_0 & 0 & \dots & 0 \\ 0 & \lambda_1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \lambda_{n-1} \end{bmatrix}.$$

As we've said, this is fantastic: not only do we have a basis of eigenvectors, but the eigenvectors are pairwise orthogonal, and so they form an orthonormal basis of V .

Question 17.3.3. Show that all eigenvalues of T are real. ($T = T^\dagger$.)

Back to quantum computation. Suppose we have a state $|\psi\rangle \in H$, where $\dim H = 2$; we haven't distinguished a particular basis yet, so we just have a nonzero vector. Then the way observations work (and this is physics, so you'll have to take my word for it) is as follows:

Pick a Hermitian operator $T : H \rightarrow H$; then observations of T return eigenvalues of T .

To be precise:

- Pick a Hermitian operator $T : H \rightarrow H$, which is called the **observable**.
- Consider its eigenvalues $\lambda_0, \dots, \lambda_{n-1}$ and corresponding eigenvectors $|0\rangle_T, \dots, |n-1\rangle_T$. Tacitly we may assume that $|0\rangle_T, \dots, |n-1\rangle_T$ form an orthonormal basis of H . (The subscript T is here to distinguish the eigenvectors of T from the basis elements of H .)
- Write $|\psi\rangle$ in the orthonormal basis as

$$c_0 |0\rangle_T + c_1 |1\rangle_T + \dots + c_{n-1} |n-1\rangle_T.$$

- Then the probability of observing λ_i is

$$\frac{|c_i|^2}{|c_0|^2 + \dots + |c_{n-1}|^2}.$$

This is called making an **observation along T** .

Note that in particular, for any nonzero constant c , $|\psi\rangle$ and $c|\psi\rangle$ are indistinguishable, which is why we like to normalize $|\psi\rangle$. But the queerest thing of all is what happens to $|\psi\rangle$: by measuring it, we actually destroy information. This behavior is called **quantum collapse**.

- Suppose for simplicity that we observe $|\psi\rangle$ with T and obtain an eigenvalue λ , and that $|i\rangle_T$ is the only eigenvector with this eigenvalue. Then, the state $|\psi\rangle$ *collapses* to just the state $c_i |i\rangle_T$: all the other information is destroyed. (In fact, we may as well say it collapses to $|i\rangle_T$, since again constant factors are not relevant.)

- More generally, if we observe λ , consider the generalized eigenspace H_λ (i.e. the span of eigenvectors with the same eigenvalue). Then the physical state $|\psi\rangle$ has been changed as well: it has now been projected onto the eigenspace H_λ . In still other words, after observation, the state collapses to

$$\sum_{\substack{0 \leq i \leq n \\ \lambda_i = \lambda}} c_i |i\rangle_T.$$

In other words,

When we make a measurement, the coefficients from different eigenspaces are destroyed.

Why does this happen? Beats me... physics (and hence real life) is weird. But anyways, an example.

Example 17.3.4 (Quantum measurement of a state $|\psi\rangle$)

Let $H = \mathbb{C}^{\oplus 2}$ with orthonormal basis $|0\rangle$ and $|1\rangle$ and consider the state

$$|\psi\rangle = \frac{i}{\sqrt{5}} |0\rangle + \frac{2}{\sqrt{5}} |1\rangle = \begin{bmatrix} i/\sqrt{5} \\ 2/\sqrt{5} \end{bmatrix} \in H.$$

(a) Let

$$T = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}.$$

This has eigenvectors $|0\rangle = |0\rangle_T$ and $|1\rangle = |1\rangle_T$, with eigenvalues $+1$ and -1 . So if we measure $|\psi\rangle$ to T , we get $+1$ with probability $1/5$ and -1 with probability $4/5$. After this measurement, the original state collapses to $|0\rangle$ if we measured 0, and $|1\rangle$ if we measured 1. So we never learn the original probabilities.

(b) Now consider $T = \text{id}$, and arbitrarily pick two orthonormal eigenvectors $|0\rangle_T, |1\rangle_T$; thus $\psi = c_0 |0\rangle_T + c_1 |1\rangle_T$. Since all eigenvalues of T are $+1$, our measurement will always be $+1$ no matter what we do. But there is also no collapsing, because none of the coefficients get destroyed.

(c) Now consider

$$T = \begin{bmatrix} 0 & 7 \\ 7 & 0 \end{bmatrix}.$$

The two normalized eigenvectors are

$$|0\rangle_T = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix} \quad |1\rangle_T = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -1 \end{bmatrix}$$

with eigenvalues $+7$ and -7 respectively. In this basis, we have

$$|\psi\rangle = \frac{2+i}{\sqrt{10}} |0\rangle_T + \frac{-2+i}{\sqrt{10}} |1\rangle_T.$$

So we get $+7$ with probability $\frac{1}{2}$ and -7 with probability $\frac{1}{2}$, and after the measurement, $|\psi\rangle$ collapses to one of $|0\rangle_T$ and $|1\rangle_T$.

Question 17.3.5. Suppose we measure $|\psi\rangle$ with T and get λ . What happens if we measure with T again?

For $H = \mathbb{C}^{\oplus 2}$ we can come up with more classes of examples using the so-called **Pauli matrices**. These are the three Hermitian matrices

$$\sigma_z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \quad \sigma_x = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad \sigma_y = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}.$$

These matrices are important because:

Question 17.3.6. Show that these three matrices, plus the identity matrix, form a basis for the set of Hermitian 2×2 matrices.

So the Pauli matrices are a natural choice of basis.

Their normalized eigenvectors are

$$\begin{aligned} |\uparrow\rangle &= |0\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix} & |\downarrow\rangle &= |1\rangle = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \\ |\rightarrow\rangle &= \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix} & |\leftarrow\rangle &= \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -1 \end{bmatrix} \\ |\otimes\rangle &= \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ i \end{bmatrix} & |\odot\rangle &= \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -i \end{bmatrix} \end{aligned}$$

which we call “z-up”, “z-down”, “x-up”, “x-down”, “y-up”, “y-down”. (The eigenvalues are +1 for “up” and −1 for “down”.) So, given a state $|\psi\rangle \in \mathbb{C}^{\oplus 2}$ we can make a measurement with respect to any of these three bases by using the corresponding Pauli matrix.

In light of this, the previous examples were (a) measuring along σ_z , (b) measuring along id, and (c) measuring along $7\sigma_x$.

Notice that if we are given a state $|\psi\rangle$, and are told in advance that it is either $|\rightarrow\rangle$ or $|\leftarrow\rangle$ (or any other orthogonal states) then we are in what is more or less a classical situation. Specifically, if we make a measurement along σ_x , then we find out which state that $|\psi\rangle$ was in (with 100% certainty), and the state does not undergo any collapse. Thus, orthogonal states are reliably distinguishable.

§17.4 Entanglement

Prototypical example for this section: Singlet state: spooky action at a distance.

If you think that’s weird, well, it gets worse.

Qubits don’t just act independently: they can talk to each other by means of a *tensor product*. Explicitly, consider

$$H = \mathbb{C}^{\oplus 2} \otimes \mathbb{C}^{\oplus 2}$$

endowed with the norm described in Problem 9C*. One should think of this as a qubit A in a space H_A along with a second qubit B in a different space H_B , which have been allowed to interact in some way, and $H = H_A \otimes H_B$ is the set of possible states of *both* qubits. Thus

$$|0\rangle_A \otimes |0\rangle_B, \quad |0\rangle_A \otimes |1\rangle_B, \quad |1\rangle_A \otimes |0\rangle_B, \quad |1\rangle_A \otimes |1\rangle_B$$

is an orthonormal basis of H ; here $|i\rangle_A$ is the basis of the first $\mathbb{C}^{\oplus 2}$ while $|i\rangle_B$ is the basis of the second $\mathbb{C}^{\oplus 2}$, so these vectors should be thought of as “unrelated” just as with any tensor product. The pure tensors mean exactly what you want: for example $|0\rangle_A \otimes |1\rangle_B$ means “0 for qubit A and 1 for qubit B ”.

As before, a measurement of a state in H requires a Hermitian map $H \rightarrow H$. In particular, if we only want to measure the qubit B along M_B , we can use the operator

$$\text{id}_A \otimes M_B.$$

The eigenvalues of this operator coincide with the ones for M_B , and the eigenspace for λ will be the $H_A \otimes (H_B)_\lambda$, so when we take the projection the A qubit will be unaffected.

This does what you would hope for pure tensors in H :

Example 17.4.1 (Two non-entangled qubits)

Suppose we have qubit A in the state $\frac{i}{\sqrt{5}}|0\rangle_A + \frac{2}{\sqrt{5}}|1\rangle_A$ and qubit B in the state $\frac{1}{\sqrt{2}}|0\rangle_B + \frac{1}{\sqrt{2}}|1\rangle_B$. So, the two qubits in tandem are represented by the pure tensor

$$|\psi\rangle = \left(\frac{i}{\sqrt{5}}|0\rangle_A + \frac{2}{\sqrt{5}}|1\rangle_A \right) \otimes \left(\frac{1}{\sqrt{2}}|0\rangle_B + \frac{1}{\sqrt{2}}|1\rangle_B \right).$$

Suppose we measure $|\psi\rangle$ along

$$M = \text{id}_A \otimes \sigma_z^B.$$

The eigenspace decomposition is

- +1 for the span of $|0\rangle_A \otimes |0\rangle_B$ and $|1\rangle_A \otimes |0\rangle_B$, and
- −1 for the span of $|0\rangle_A \otimes |1\rangle_B$ and $|1\rangle_A \otimes |1\rangle_B$.

(We could have used other bases, like $|\rightarrow\rangle_A \otimes |0\rangle_B$ and $|\leftarrow\rangle_A \otimes |0\rangle_B$ for the first eigenspace, but it doesn’t matter.) Expanding $|\psi\rangle$ in the four-element basis, we find that we’ll get the first eigenspace with probability

$$\left| \frac{i}{\sqrt{10}} \right|^2 + \left| \frac{2}{\sqrt{10}} \right|^2 = \frac{1}{2}.$$

and the second eigenspace with probability $\frac{1}{2}$ as well. (Note how the coefficients for A don’t do anything!) After the measurement, we destroy the coefficients of the other eigenspace; thus (after re-normalization) we obtain the collapsed state

$$\left(\frac{i}{\sqrt{5}}|0\rangle_A + \frac{2}{\sqrt{5}}|1\rangle_A \right) \otimes |0\rangle_B \quad \text{or} \quad \left(\frac{i}{\sqrt{5}}|0\rangle_A + \frac{2}{\sqrt{5}}|1\rangle_A \right) \otimes |1\rangle_B$$

again with 50% probability each.

So this model lets us more or less work with the two qubits independently: when we make the measurement, we just make sure to not touch the other qubit (which corresponds to the identity operator).

Exercise 17.4.2. Show that if $\text{id}_A \otimes \sigma_x^B$ is applied to the $|\psi\rangle$ in this example, there is no collapse at all. What’s the result of this measurement?

Since the $\otimes$ is getting cumbersome to write, we say:

Abuse of Notation 17.4.3. From now on $|0\rangle_A \otimes |0\rangle_B$ will be abbreviated to just $|00\rangle$, and similarly for $|01\rangle$, $|10\rangle$, $|11\rangle$.

Example 17.4.4 (Simultaneously measuring a general 2-Qubit state)

Consider a normalized state $|\psi\rangle$ in $H = \mathbb{C}^{\oplus 2} \otimes \mathbb{C}^{\oplus 2}$, say

$$|\psi\rangle = \alpha |00\rangle + \beta |01\rangle + \gamma |10\rangle + \delta |11\rangle.$$

We can make a measurement along the diagonal matrix $T : H \rightarrow H$ with

$$T(|00\rangle) = 0 |00\rangle, \quad T(|01\rangle) = 1 |01\rangle, \quad T(|10\rangle) = 2 |10\rangle, \quad T(|11\rangle) = 3 |11\rangle.$$

Thus we get each of the eigenvalues 0, 1, 2, 3 with probability $|\alpha|^2, |\beta|^2, |\gamma|^2, |\delta|^2$. So if we like we can make “simultaneous” measurements on two qubits in the same way that we make measurements on one qubit.

However, some states behave very weirdly.

Example 17.4.5 (The singlet state)

Consider the state

$$|\Psi_-\rangle = \frac{1}{\sqrt{2}} |01\rangle - \frac{1}{\sqrt{2}} |10\rangle$$

which is called the **singlet state**. One can see that $|\Psi_-\rangle$ is not a simple tensor, which means that it doesn’t just consist of two qubits side by side: the qubits in H_A and H_B have become *entangled*.

Now, what happens if we measure just the qubit A ? This corresponds to making the measurement

$$T = \sigma_z^A \otimes \text{id}_B.$$

The eigenspace decomposition of T can be described as:

- The span of $|00\rangle$ and $|01\rangle$, with eigenvalue $+1$.
- The span of $|10\rangle$ and $|11\rangle$, with eigenvalue -1 .

So one of two things will happen:

- With probability $\frac{1}{2}$, we measure $+1$ and the collapsed state is $|01\rangle$.
- With probability $\frac{1}{2}$, we measure -1 and the collapsed state is $|10\rangle$.

But now we see that measurement along A has told us what the state of the bit B is completely!

By solely looking at measurements on A , we learn B ; this paradox is called *spooky action at a distance*, or in Einstein’s tongue, **spukhafte Fernwirkung**. Thus,

In tensor products of Hilbert spaces, states which are not pure tensors correspond to “entangled” states.

What this really means is that the qubits cannot be described independently; the state of the system must be given as a whole. That's what entangled states mean: the qubits somehow depend on each other.

§17.5 Problems to think about

Problem 17A. We measure $|\Psi_-\rangle$ by $\sigma_x^A \otimes \text{id}_B$, and hence obtain either $+1$ or -1 . Determine the state of qubit B from this measurement.

Problem 17B (Greenberger-Horne-Zeilinger paradox). Consider the state in $(\mathbb{C}^{\oplus 2})^{\otimes 3}$

$$|\Psi\rangle_{\text{GHZ}} = \frac{1}{\sqrt{2}} (|0\rangle_A |0\rangle_B |0\rangle_C - |1\rangle_A |1\rangle_B |1\rangle_C).$$

Find the value of the measurements along each of

$$\sigma_y^A \otimes \sigma_y^B \otimes \sigma_x^C, \quad \sigma_y^A \otimes \sigma_x^B \otimes \sigma_y^C, \quad \sigma_x^A \otimes \sigma_y^B \otimes \sigma_y^C, \quad \sigma_x^A \otimes \sigma_x^B \otimes \sigma_x^C.$$

As for the paradox: what happens if you multiply all these measurement together?

18 Quantum circuits

Now that we've discussed qubits, we can talk about how to use them in circuits. The key change — and the reason that quantum circuits can do things that classical circuits cannot — is the fact that we are allowing linear combinations of 0 and 1.

§18.1 Classical logic gates

In classical logic, we build circuits which take in some bits for input, and output some more bits for input. These circuits are built out of individual logic gates. For example, the **AND gate** can be pictured as follows.



One can also represent the AND gate using the “truth table”:

A	B	A and B
0	0	0
0	1	0
1	0	0
1	1	1

Similarly, we have the **OR gate** and the **NOT gate**:

A	B	A or B
0	0	0
0	1	1
1	0	1
1	1	1

A	not A
0	1
1	0

We also have a so-called **COPY gate**, which duplicates a bit.



Of course, the first theorem you learn about these gates is that:

Theorem 18.1.1 (AND, OR, NOT, COPY are universal)

The set of four gates AND, OR, NOT, COPY is universal in the sense that any boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ can be implemented as a circuit using only these gates.

Proof. Somewhat silly: we essentially write down a circuit that OR's across all input strings in $f^{\text{pre}}(1)$. For example, suppose we have $n = 3$ and want to simulate the function $f(abc)$ with $f(011) = f(110) = 1$ and 0 otherwise. Then the corresponding Boolean expression for f is simply

$$f(abc) = [(\text{not } a) \text{ and } b \text{ and } c] \text{ or } [a \text{ and } b \text{ and } (\text{not } c)].$$

Clearly, one can do the same for any other f , and implement this logic into a circuit. $\square$

Remark 18.1.2. Since x and $y = \text{not}((\text{not } x) \text{ or } (\text{not } y))$, it follows that in fact, we can dispense with the AND gate.

§18.2 Reversible classical logic

Prototypical example for this section: CNOT gate, Toffoli gate.

For the purposes of quantum mechanics, this is not enough. To carry through the analogy we in fact need gates that are **reversible**, meaning the gates are bijections from the input space to the output space. In particular, such gates must take the same number of input and output gates.

Example 18.2.1 (Reversible gates)

- (a) None of the gates AND, OR, COPY are reversible for dimension reasons.
- (b) The NOT gate, however, is reversible: it is a bijection $\{0, 1\} \rightarrow \{0, 1\}$.

Example 18.2.2 (The CNOT gate)

The controlled-NOT gate, or the **CNOT** gate, is a reversible 2-bit gate with the following truth table.

In	Out
0 0	0 0
1 0	1 1
0 1	0 1
1 1	1 0

In other words, this gate XOR's the first bit to the second bit, while leaving the first bit unchanged. It is depicted as follows.

$$\begin{array}{c} x \text{ --- } \bullet \text{ --- } x \\ y \text{ --- } \oplus \text{ --- } x + y \pmod{2} \end{array}$$

The first dot is called the “control”, while the $\oplus$ is the “negation” operation: the first bit controls whether the second bit gets flipped or not. Thus, a typical application might be as follows.

$$\begin{array}{c} 1 \text{ --- } \bullet \text{ --- } 1 \\ 0 \text{ --- } \oplus \text{ --- } 1 \end{array}$$

So, NOT and CNOT are the only nontrivial reversible gates on two bits.

We now need a different definition of universal for our reversible gates.

Definition 18.2.3. A set of reversible gates can **simulate** a Boolean function $f(x_1 \dots x_n)$, if one can implement a circuit which takes

- As input, $x_1 \dots x_n$ plus some fixed bits set to 0 or 1, called **ancilla bits**¹.
- As output, the input bits $x_1, \dots, x_n$, the output bit $f(x_1, \dots, x_n)$, and possibly some extra bits (called **garbage bits**).

The gate(s) are **universal** if they can simulate any Boolean function.

For example, the CNOT gate can simulate the NOT gate, using a single ancilla bit 1 but with no garbage, according to the following circuit.

$$\begin{array}{c} x \text{ --- } \bullet \text{ --- } x \\ 1 \text{ --- } \oplus \text{ --- } \text{not } x \end{array}$$

¹The English word “ancilla” means “maid”.

Unfortunately, it is not universal.

Proposition 18.2.4 (CNOT $\nrightarrow$ AND)

The CNOT gate cannot simulate the boolean function “ x and y ”.

Sketch of Proof. One can see that any function simulated using only CNOT gates must be of the form

$$a_1x_1 + a_2x_2 + \cdots + a_nx_n \pmod{2}$$

because CNOT is the map $(x, y) \mapsto (x, x + y)$. Thus, even with ancilla bits, we can only create functions of the form $ax + by + c \pmod{2}$ for fixed a, b, c . The AND gate is not of this form. $\square$

So, we need at least a three-qubit gate. The most commonly used one is:

Definition 18.2.5. The three-bit **Toffoli gate**, also called the CCNOT gate, is given by

$$\begin{array}{ccc} x & \bullet & x \\ y & \bullet & y \\ z & \oplus & z + xy \pmod{2} \end{array}$$

So the Toffoli has two controls, and toggles the last bit if and only if both of the control bits are 1.

This replacement is sufficient.

Theorem 18.2.6 (Toffoli gate is universal)

The Toffoli gate is universal.

Proof. We will show it can *reversibly* simulate AND, NOT, hence OR, which we know is enough to show universality. (We don’t need COPY because of reversibility.)

For the AND gate, we draw the circuit

$$\begin{array}{ccc} x & \bullet & x \\ y & \bullet & y \\ 0 & \oplus & x \text{ and } y \end{array}$$

with one ancilla bit, and no garbage bits.

For the NOT gate, we use two ancilla 1 bits and two garbage bits:

$$\begin{array}{ccc} 1 & \bullet & 1 \\ z & \bullet & z \\ 1 & \oplus & \text{not } z \end{array}$$

This completes the proof. $\square$

Hence, in theory we can create any classical circuit we desire using the Toffoli gate alone. Of course, this could require exponentially many gates for even the simplest of functions. Fortunately, this is NO BIG DEAL because I’m a math major, and having 2^n gates is a problem best left for the CS majors.

§18.3 Quantum logic gates

In quantum mechanics, since we can have *linear combinations* of basis elements, our logic gates will instead consist of *linear maps*. Moreover, in quantum computation, gates are always reversible, which was why we took the time in the previous section to show that we can still simulate any function when restricted to reversible gates (e.g. using the Toffoli gate).

First, some linear algebra:

Definition 18.3.1. Let V be a finite dimensional inner product space. Then for a map $U : V \rightarrow V$, the following are equivalent:

- $\langle U(x), U(y) \rangle = \langle x, y \rangle$ for $x, y \in V$.
- $U^\dagger$ is the inverse of U .
- $\|x\| = \|U(x)\|$ for $x \in V$.

The map U is called **unitary** if it satisfies these equivalent conditions.

Then

Quantum logic gates are unitary matrices.

In particular, unlike the classical situation, quantum gates are always reversible (and hence they always take the same number of input and output bits).

For example, consider the CNOT gate. Its quantum analog should be a unitary map $U_{\text{CNOT}} : H \rightarrow H$, where $H = \mathbb{C}^{\oplus 2} \otimes \mathbb{C}^{\oplus 2}$, given on basis elements by

$$\begin{aligned} U_{\text{CNOT}}(|00\rangle) &= |00\rangle, & U_{\text{CNOT}}(|01\rangle) &= |01\rangle \\ U_{\text{CNOT}}(|10\rangle) &= |11\rangle, & U_{\text{CNOT}}(|11\rangle) &= |10\rangle. \end{aligned}$$

So pictorially, the quantum CNOT gate is given by

$$\begin{array}{cccc} \begin{array}{c} |0\rangle \\ |0\rangle \end{array} \begin{array}{c} \bullet \\ \oplus \end{array} \begin{array}{c} |0\rangle \\ |0\rangle \end{array} & \begin{array}{c} |0\rangle \\ |1\rangle \end{array} \begin{array}{c} \bullet \\ \oplus \end{array} \begin{array}{c} |0\rangle \\ |1\rangle \end{array} & \begin{array}{c} |1\rangle \\ |0\rangle \end{array} \begin{array}{c} \bullet \\ \oplus \end{array} \begin{array}{c} |1\rangle \\ |1\rangle \end{array} & \begin{array}{c} |1\rangle \\ |1\rangle \end{array} \begin{array}{c} \bullet \\ \oplus \end{array} \begin{array}{c} |1\rangle \\ |0\rangle \end{array} \end{array}$$

OK, so what? The whole point of quantum mechanics is that we allow linear qubits to be in linear combinations of $|0\rangle$ and $|1\rangle$, too, and this will produce interesting results. For example, let's take $|\leftarrow\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$ and plug it into the top, with $|1\rangle$ on the bottom, and see what happens:

$$U_{\text{CNOT}}(|\leftarrow\rangle \otimes |1\rangle) = U_{\text{CNOT}}\left(\frac{1}{\sqrt{2}}(|01\rangle - |11\rangle)\right) = \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle) = |\Psi_-\rangle$$

which is the fully entangled *singlet state*! Picture:

$$\begin{array}{c} |\leftarrow\rangle \begin{array}{c} \bullet \\ \oplus \end{array} |\Psi_-\rangle \\ |1\rangle \end{array}$$

Thus, when we input mixed states into our quantum gates, the outputs are often entangled states, even when the original inputs are not entangled.

Example 18.3.2 (More examples of quantum gates)

- (a) Every reversible classical gate that we encountered before has a quantum analog obtained in the same way as CNOT: by specifying the values on basis elements. For example, there is a quantum Tofoli gate which for example sends

$$\begin{array}{c} |1\rangle \\ |1\rangle \\ |0\rangle \end{array} \begin{array}{c} \bullet \\ \bullet \\ \oplus \end{array} \begin{array}{c} |1\rangle \\ |1\rangle \\ |1\rangle \end{array}$$

- (b) The **Hadamard gate** on one qubit is a rotation given by

$$\begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{bmatrix}.$$

Thus, it sends $|0\rangle$ to $|\rightarrow\rangle$ and $|1\rangle$ to $|\leftarrow\rangle$. Note that the Hadamard gate is its own inverse. It is depicted by an “ H ” box.

$$|0\rangle \xrightarrow{H} |\rightarrow\rangle$$

- (c) More generally, if U is a 2×2 unitary matrix (i.e. a map $\mathbb{C}^{\oplus 2} \rightarrow \mathbb{C}^{\oplus 2}$) then there is **U -rotation gate** similar to the previous one, which applies U to the input.

$$|\psi\rangle \xrightarrow{U} U|\psi\rangle$$

For example, the classical NOT gate is represented by $U = \sigma_x$.

- (d) A **controlled U -rotation gate** generalizes the CNOT gate. Let $U : \mathbb{C}^{\oplus 2} \rightarrow \mathbb{C}^{\oplus 2}$ be a rotation gate, and let $H = \mathbb{C}^{\oplus 2} \otimes \mathbb{C}^{\oplus 2}$ be a 2-qubit space. Then the controlled U gate has the following circuit diagrams.

$$\begin{array}{c} |0\rangle \\ |\psi\rangle \end{array} \begin{array}{c} \bullet \\ \boxed{U} \end{array} \begin{array}{c} |0\rangle \\ |\psi\rangle \end{array} \quad \begin{array}{c} |1\rangle \\ |\psi\rangle \end{array} \begin{array}{c} \bullet \\ \boxed{U} \end{array} \begin{array}{c} |1\rangle \\ U|\psi\rangle \end{array}$$

Thus, U is applied when the controlling bit is 1, and CNOT is the special case $U = \sigma_x$. As before, we get interesting behavior if the control is mixed.

And now, some more counterintuitive quantum behavior. Suppose we try to use CNOT as a copy, with truth table.

In	Out
0 0	0 0
1 0	1 1
0 1	0 1
1 1	1 0

The point of this gate is to be used with a garbage 0 at the bottom to try and simulate a “copy” operation. So indeed, one can check that

$$\begin{array}{c} |0\rangle \\ |0\rangle \end{array} \begin{array}{c} \boxed{U} \end{array} \begin{array}{c} |0\rangle \\ |0\rangle \end{array} \quad \begin{array}{c} |1\rangle \\ |0\rangle \end{array} \begin{array}{c} \boxed{U} \end{array} \begin{array}{c} |1\rangle \\ |1\rangle \end{array}$$

Thus we can copy $|0\rangle$ and $|1\rangle$. But as we’ve already seen if we input $|\leftarrow\rangle \otimes |0\rangle$ into U , we end up with the entangled state $|\Psi_-\rangle$ which is decisively *not* the $|\leftarrow\rangle \otimes |\leftarrow\rangle$ we wanted.

And in fact, the so-called **no-cloning theorem** implies that it's impossible to duplicate an arbitrary $|\psi\rangle$; the best we can do is copy specific orthogonal states as in the classical case. See also Problem 18B.

§18.4 Deutsch-Jozsa algorithm

The Deutsch-Jozsa algorithm is the first example of a nontrivial quantum algorithm which cannot be performed classically: it is a “proof of concept” that would later inspire Grover’s search algorithm and Shor’s factoring algorithm.

The problem is as follows: we’re given a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, and promised that the function f is either

- A constant function, or
- A balanced function, meaning that exactly half the inputs map to 0 and half the inputs map to 1.

The function f is given in the form of a reversible black box U_f which is the control of a NOT gate, so it can be represented as the circuit diagram

$$\begin{array}{c} |x_1 x_2 \dots x_n\rangle \xrightarrow{\text{---}^n \text{---}} \boxed{U_f} \text{---} |x_1 x_2 \dots x_n\rangle \\ |y\rangle \text{---} \boxed{U_f} \text{---} |y + f(x) \pmod{2}\rangle \end{array}$$

i.e. if $f(x_1, \dots, x_n) = 0$ then the gate does nothing, otherwise the gate flips the y bit at the bottom. The slash with the n indicates that the top of the input really consists of n qubits, not just the one qubit drawn, and so the black box U_f is a map on $n + 1$ qubits.

The problem is to determine, with as few calls to the black box U_f as possible, whether f is balanced or constant.


Question 18.4.1. Classically, show that in the worst case we may need up to $2^{n-1} + 1$ calls to the function f to answer the question.

So with only classical tools, it would take $O(2^n)$ queries to determine whether f is balanced or constant. However,

Theorem 18.4.2 (Deutsch-Jozsa)

The Deutsch-Jozsa problem can be determined in a quantum circuit with only a single call to the black box.

Proof. For concreteness, we do the case $n = 1$ explicitly; the general case is contained in Problem 18C. We claim that the necessary circuit is



Here the H ’s are Hadamard gates, and meter at the end of the rightmost wire indicates that we make a measurement along the usual $|0\rangle, |1\rangle$ basis. This is not a typo! Even though classically the top wire is just a repeat of the input information, we are about to see that it’s the top we want to measure.

Note that after the two Hadamard operations, the state we get is

$$\begin{aligned} |01\rangle &\xrightarrow{H^{\otimes 2}} \left(\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \right) \otimes \left(\frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) \right) \\ &= \frac{1}{2} \left(|0\rangle \otimes (|0\rangle - |1\rangle) + |1\rangle \otimes (|0\rangle - |1\rangle) \right). \end{aligned}$$

So after applying U_f , we obtain

$$\frac{1}{2} \left(|0\rangle \otimes (|0 + f(0)\rangle - |1 + f(0)\rangle) + |1\rangle \otimes (|0 + f(1)\rangle - |1 + f(1)\rangle) \right)$$

where the modulo 2 has been left implicit. Now, observe that the effect of going from $|0\rangle - |1\rangle$ to $|0 + f(x)\rangle - |1 + f(x)\rangle$ is merely to either keep the state the same (if $f(x) = 0$) or to negate it (if $f(x) = 1$). So we can simplify and factor to get

$$\frac{1}{2} \left((-1)^{f(0)} |0\rangle + (-1)^{f(1)} |1\rangle \right) \otimes (|0\rangle - |1\rangle).$$

Thus, the picture so far is:

$$\begin{array}{c} |0\rangle \xrightarrow{H} \boxed{U_f} \xrightarrow{\frac{1}{\sqrt{2}}((-1)^{f(0)}|0\rangle + (-1)^{f(1)}|1\rangle)} \\ |1\rangle \xrightarrow{H} \boxed{U_f} \xrightarrow{\frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)} \end{array}$$

In particular, the resulting state is not entangled, and we can simply discard the last qubit (!). Now observe:

- If f is constant, then the upper-most state is $\pm |\rightarrow\rangle$.
- If f is balanced, then the upper-most state is $\pm |\leftarrow\rangle$.

So simply doing a measurement along σ_x will give us the answer. Equivalently, perform another H gate (so that $H|\rightarrow\rangle = |0\rangle$, $H|\leftarrow\rangle = |1\rangle$) and measuring along σ_z in the usual $|0\rangle, |1\rangle$ basis. Thus for $n = 1$ we only need a single call to the oracle. $\square$

§18.5 Problems to think about

Problem 18A (Fredkin gate). The **Fredkin gate** (also called the controlled swap, or CSWAP gate) is the three-bit gate with the following truth table:


In	Out
0 0 0	0 0 0
0 0 1	0 0 1
0 1 0	0 1 0
0 1 1	0 1 1
1 0 0	1 0 0
1 0 1	1 1 0
1 1 0	1 0 1
1 1 1	1 1 1

Thus the gate swaps the last two input bits whenever the first bit is 1. Show that this gate is also reversible and universal.

Problem 18B (Baby no-cloning theorem). Show that there is no unitary map U on two qubits which sends $U(|\psi\rangle \otimes |0\rangle) = |\psi\rangle \otimes |\psi\rangle$ for any qubit $|\psi\rangle$, i.e. the following circuit diagram is impossible.

$$\begin{array}{c} |\psi\rangle \xrightarrow{\quad} \boxed{U} \xrightarrow{\quad} |\psi\rangle \\ |0\rangle \xrightarrow{\quad} \boxed{U} \xrightarrow{\quad} |\psi\rangle \end{array}$$

Problem 18C (Deutsch-Jozsa). Given the black box U_f described in the Deutsch-Jozsa algorithm, consider the following circuit.




That is, take n copies of $|0\rangle$, apply the Hadamard rotation to all of them, apply U_f , reverse the Hadamard to all n input bits (again discarding the last bit), then measure all n bits in the $|0\rangle/|1\rangle$ basis (as in Example 17.4.4).

Show that the probability of measuring $|0\dots 0\rangle$ is 1 if f is constant and 0 if f is balanced.

Problem 18D[†] (Barenco et al, 1995; arXiv:quant-ph/9503016v1). Let

$$P = \begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix} \quad Q = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & -i \\ -i & 1 \end{bmatrix}$$

Verify that the quantum Toffoli gate can be implemented using just controlled rotations via the circuit



This was a big surprise to researchers when discovered, because classical reversible logic requires three-bit gates (e.g. Toffoli, Fredkind).

19 Shor's algorithm

OK, now for Shor's Algorithm: how to factor $M = pq$ in $O((\log M)^2)$ time.

§19.1 The classical (inverse) Fourier transform

The “crux move” in Shor's algorithm is the so-called quantum Fourier transform. The Fourier transform is used to extract *periodicity* in data, and it turns out the quantum analogue is a lot faster than the classical one.

Let me throw the definition at you first. Let N be a positive integer, and let $\omega_N = \exp\left(\frac{2\pi i}{N}\right)$.

Definition 19.1.1. Given a tuple of complex numbers

$$(x_0, x_1, \dots, x_{N-1})$$

its **discrete inverse Fourier transform** is the sequence $(y_0, y_1, \dots, y_{N-1})$ defined by

$$y_k = \frac{1}{N} \sum_{j=0}^{N-1} \omega_N^{jk} x_j.$$

Equivalently, one is applying the matrix

$$\frac{1}{N} \begin{bmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & \omega_N & \omega_N^2 & \dots & \omega_N^{N-1} \\ 1 & \omega_N^2 & \omega_N^4 & \dots & \omega_N^{2(N-1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega_N^{N-1} & \omega_N^{2(N-1)} & \dots & \omega_N^{(N-1)^2} \end{bmatrix} \begin{bmatrix} x_0 \\ x_1 \\ \vdots \\ x_N \end{bmatrix} = \begin{bmatrix} y_0 \\ y_1 \\ \vdots \\ y_N \end{bmatrix}.$$

The reason this operation is important is because it lets us detect if the x_i are periodic:

Example 19.1.2 (Example of discrete inverse Fourier transform)

Let $N = 6$, $\omega = \omega_6 = \exp\left(\frac{2\pi i}{6}\right)$ and suppose $(x_0, x_1, x_2, x_3, x_4, x_5) = (0, 1, 0, 1, 0, 1)$ (hence x_i is periodic modulo 2). Thus,

$$\begin{aligned} y_0 &= \frac{1}{3} (\omega^0 + \omega^0 + \omega^0) = 1 \\ y_1 &= \frac{1}{3} (\omega^1 + \omega^3 + \omega^5) = 0 \\ y_2 &= \frac{1}{3} (\omega^2 + \omega^6 + \omega^{10}) = 0 \\ y_3 &= \frac{1}{3} (\omega^3 + \omega^9 + \omega^{15}) = -1 \\ y_4 &= \frac{1}{3} (\omega^4 + \omega^{12} + \omega^{20}) = 0 \\ y_5 &= \frac{1}{3} (\omega^5 + \omega^{15} + \omega^{25}) = 0. \end{aligned}$$

Thus, in the inverse transformation the “amplitudes” are all concentrated at multiples of 3; thus this reveals the periodicity of the original sequence by $\frac{N}{3} = 2$.

More generally, given a sequence of 1's appearing with period r , the amplitudes will peak at inputs which are divisible by $\frac{N}{\gcd(N,r)}$.

Remark 19.1.3. The fact that this operation is called the “inverse” Fourier transform is mostly a historical accident (as my understanding goes). Confusingly, the corresponding quantum operation is the (not-inverted) Fourier transform.

If we apply the definition as written, computing the transform takes $O(N^2)$ time. It turns out that by an algorithm called the **fast Fourier transform** (whose details we won't discuss), one can reduce this to $O(N \log N)$ time. However, for Shor's algorithm this is also insufficient; we need something like $O((\log N)^2)$ instead. This is where the quantum Fourier transform comes in.

§19.2 The quantum Fourier transform

Note that to compute a Fourier transform, we need to multiply an $N \times N$ matrix with an N -vector, so this takes $O(N^2)$ multiplications. However, we are about to show that with a quantum computer, one can do this using $O((\log N)^2)$ quantum gates when $N = 2^n$, on a system with n qubits.

First, some more notation:

Abuse of Notation 19.2.1. In what follows, $|x\rangle$ will refer to $|x_n\rangle \otimes |x_{n-1}\rangle \otimes \cdots \otimes |x_1\rangle$ where $x = x_n x_{n-1} \dots x_1$ in binary. For example, if $n = 3$ then $|6\rangle$ really means $|1\rangle \otimes |1\rangle \otimes |0\rangle$.

Observe that the n -qubit space now has an orthonormal basis $|0\rangle, |1\rangle, \dots, |N-1\rangle$

Definition 19.2.2. Consider an n -qubit state

$$|\psi\rangle = \sum_{k=0}^{N-1} x_k |k\rangle.$$

The **quantum Fourier transform** is defined by

$$U_{\text{QFT}}(|\psi\rangle) = \frac{1}{\sqrt{N}} \sum_{j=0}^{N-1} \left(\sum_{k=0}^{N-1} \omega_N^{jk} \right) |j\rangle.$$

In other words, using the basis $|0\rangle, \dots, |N-1\rangle$, U_{QFT} is given by the matrix

$$U_{\text{QFT}} = \frac{1}{\sqrt{N}} \begin{bmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & \omega_N & \omega_N^2 & \dots & \omega_N^{N-1} \\ 1 & \omega_N^2 & \omega_N^4 & \dots & \omega_N^{2(N-1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega_N^{N-1} & \omega_N^{2(N-1)} & \dots & \omega_N^{(N-1)^2} \end{bmatrix}$$

This is the exactly the same definition as before, except we have a $\sqrt{N}$ factor added so that U_{QFT} is unitary. But the trick is that in the quantum setup, the matrix can be rewritten:

Proposition 19.2.3 (Tensor representation)

Let $|x\rangle = |x_n x_{n-1} \dots x_1\rangle$. Then

$$\begin{aligned} U_{\text{QFT}}(|x_n x_{n-1} \dots x_1\rangle) &= \frac{1}{\sqrt{N}} (|0\rangle + \exp(2\pi i \cdot 0.x_1) |1\rangle) \\ &\quad \otimes (|0\rangle + \exp(2\pi i \cdot 0.x_2 x_1) |1\rangle) \\ &\quad \otimes \dots \\ &\quad \otimes (|0\rangle + \exp(2\pi i \cdot 0.x_n \dots x_1) |1\rangle) \end{aligned}$$


Proof. Direct (and quite annoying) computation. In short, expand everything. $\square$

So by using mixed states, we can deal with the quantum Fourier transform using this “multiplication by tensor product” trick that isn’t possible classically.

Now, without further ado, here’s the circuit. Define the rotation matrices

$$R_k = \begin{bmatrix} 1 & 0 \\ 0 & \exp(2\pi i / 2^k) \end{bmatrix}.$$

Then, for $n = 3$ the circuit is given by by using controlled R_k ’s as follows:




Exercise 19.2.4. Show that in this circuit, the image of $|x_3 x_2 x_1\rangle$ (for binary x_i) is

$$\left(|0\rangle + \exp(2\pi i \cdot 0.x_1) |1\rangle \right) \otimes \left(|0\rangle + \exp(2\pi i \cdot 0.x_2 x_1) |1\rangle \right) \otimes \left(|0\rangle + \exp(2\pi i \cdot 0.x_3 x_2 x_1) |1\rangle \right)$$

as claimed.

For general n , we can write this as inductively as



Question 19.2.5. Convince yourself that when $n = 3$ the two circuits displayed are equivalent.

Thus, the quantum Fourier transform is achievable with $O(n^2)$ gates, which is enormously better than the $O(N \log N)$ operations achieved by the classical fast Fourier transform (where $N = 2^n$).

§19.3 Shor's algorithm

The quantum Fourier transform is the key piece of Shor's algorithm. Now that we have it, we can solve the factoring problem.

Let $p, q > 3$ be odd primes, and assume $p \neq q$. The main idea is to turn factoring an integer $M = pq$ into a problem about finding the order of $x \pmod{M}$; the latter is a “periodicity” problem that the quantum Fourier transform will let us solve. Specifically, say that an $x \pmod{M}$ is *good* if

- (i) $\gcd(x, M) = 1$,
- (ii) The order r of $x \pmod{M}$ is even, and
- (iii) Factoring $0 \equiv (x^{r/2} - 1)(x^{r/2} + 1) \pmod{M}$, neither of the two factors is $0 \pmod{M}$.
Thus one of them is divisible by p , and the other is divisible by q .

Exercise 19.3.1 (For contest number theory practice). Show that for $M = pq$ at least half of the residues in $\mathbb{Z}_M^\times$ are good.

So if we can find the order of an arbitrary $x \in \mathbb{Z}_M^\times$, then we just keep picking x until we pick a good one (this happens more than half the time); once we do, we compute $\gcd(x^{r/2} - 1, M)$ using the Euclidean algorithm to extract one of the prime factors of M , and we're home free.

Now how do we do this? The idea is not so difficult: first we generate a sequence which is periodic modulo r .

Example 19.3.2 (Factoring 77: generating the periodic state)

Let's say we're trying to factor $M = 77$, and we randomly select $x = 2$, and want to find its order r . Let $n = 13$ and $N = 2^{13}$, and start by initializing the state

$$|\psi\rangle = \frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} |k\rangle.$$

Now, build a circuit U_x (depending on $x = 2$!) which takes $|k\rangle |0\rangle$ to $|k\rangle |2^k \pmod{M}\rangle$. Applying this to $|\psi\rangle \otimes |0\rangle$ gives

$$U(|\psi\rangle |0\rangle) = \frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} |k\rangle \otimes |2^k \pmod{M}\rangle.$$

Now suppose we measure the second qubit, and get a state of $|128\rangle$. That tells us that the collapsed state now, up to scaling, is

$$(|7\rangle + |7+r\rangle + |7+2r\rangle + \dots) \otimes |128\rangle.$$

The bottleneck is actually the circuit U_x ; one can compute $2^k \pmod{M}$ by using repeated squaring, but it's still the clumsy part of the whole operation.

In general, the operation is:

- Pick a sufficiently large $N = 2^n$ (say, $N \geq M^2$).
- Generate $|\psi\rangle = \sum_{k=0}^{2^n-1} |k\rangle$.

- Build a circuit U_x which computes $|2^k \bmod M\rangle$.
- Apply it to get a state $\frac{1}{\sqrt{N}} \sum_{k=0}^{2^n-1} |k\rangle \otimes |x^k \bmod M\rangle$.
- Measure the second qubit to cause the first qubit to collapse to something which is periodic modulo r . Let $|\phi\rangle$ denote the left qubit.

Suppose we apply the quantum Fourier transform to the left qubit $|\phi\rangle$ now: since the left bit is periodic modulo r , we expect the transform will tell us what r is. Unfortunately, this doesn't quite work out, since N is a power of two, but we don't expect r to be.

Nevertheless, consider a state

$$|\phi\rangle = |k_0\rangle + |k_0 + r\rangle + \dots$$

so for example previously we had $k_0 = 7$ if we measured 128 on $x = 2$. Applying the quantum Fourier transform, we see that the coefficient of $|j\rangle$ in the transformed image is equal to

$$\omega_N^{k_0 j} \cdot \left(\omega_N^0 + \omega_N^{jr} + \omega_N^{2jr} + \omega_N^{3jr} + \dots \right)$$

As this is a sum of roots of unity, we realize we have destructive interference unless $\omega_N^{jr} = 1$ (since N is large). In other words, we approximately have

$$U_{\text{QFT}}(|\phi\rangle) \approx \sum_{\substack{0 \leq j < N \\ jr/N \in \mathbb{Z}}} |j\rangle$$

up to scaling as usual. The bottom line is that

If we measure $U_{\text{QFT}}|\phi\rangle$ we obtain a $|j\rangle$ such that $\frac{jr}{N}$ is close to an $s \in \mathbb{Z}$.

And thus given sufficient luck we can use continued fractions to extract the value of r .

Example 19.3.3 (Finishing the factoring of $M = 77$)

As before, we made an observation to the second qubit, and thus the first qubit collapses to the state $|\phi\rangle = |7\rangle + |7+r\rangle + \dots$. Now we make a measurement and obtain $j = 4642$, which means that for some integer s we have

$$\frac{4642r}{2^{13}} \approx s.$$

Now, we analyze the continued fraction of $\frac{4642}{2^{13}}$; we find the first few convergents are

$$0, 1, \frac{1}{2}, \frac{4}{7}, \frac{13}{23}, \frac{17}{30}, \frac{1152}{2033}, \dots$$

So $\frac{17}{30}$ is a very good approximation, hence we deduce $s = 17$ and $r = 30$ as candidates. And indeed, one can check that $r = 30$ is the desired order.

This won't work all the time (for example, we could get unlucky and measure $j = 0$, i.e. $s = 0$, which would tell us no information at all).

But one can show that we succeed any time that

$$\gcd(s, r) = 1.$$

This happens at least $\frac{1}{\log r}$ of the time, and since $r < M$ this means that given sufficiently many trials, we will eventually extract the correct order r . This is Shor's algorithm.

VI

Algebraic Topology I: Homotopy

20	Some topological constructions	200
20.1	Spheres	200
20.2	Quotient topology	200
20.3	Product topology	201
20.4	Disjoint union and wedge sum	202
20.5	CW complexes	203
20.6	The torus, Klein bottle, $\mathbb{RP}^n$, $\mathbb{CP}^n$	204
20.7	Problems to think about	210
21	Fundamental groups	211
21.1	Fusing paths together	211
21.2	Fundamental groups	212
21.3	Fundamental groups are functorial	216
21.4	Higher homotopy groups	217
21.5	Homotopy equivalent spaces	218
21.6	The pointed homotopy category	220
21.7	Problems to think about	221
22	Covering projections	222
22.1	Even coverings and covering projections	222
22.2	Lifting theorem	223
22.3	Lifting correspondence	225
22.4	Regular coverings	227
22.5	The algebra of fundamental groups	228
22.6	Problems to think about	230

20 Some topological constructions

In this short chapter we briefly describe some common spaces and constructions in topology that we haven't yet discussed.

§20.1 Spheres

Recall that

$$S^n = \{(x_0, \dots, x_n) \mid x_0^2 + \dots + x_n^2 = 1\} \subset \mathbb{R}^{n+1}$$

is the surface of an n -sphere while

$$D^{n+1} = \{(x_0, \dots, x_n) \mid x_0^2 + \dots + x_n^2 \leq 1\} \subset \mathbb{R}^{n+1}$$

is the corresponding *closed ball* (So for example, D^2 is a disk in a plane while S^1 is the unit circle.)

Exercise 20.1.1. Show that the open ball $D^n \setminus S^{n-1}$ is homeomorphic to $\mathbb{R}^n$.

In particular, S^0 consists of two points, while D^1 can be thought of as the interval $[-1, 1]$.



§20.2 Quotient topology

Prototypical example for this section: $D^n/S^{n-1} = S^n$, or the torus.

Given a space X , we can *identify* some of the points together by any equivalence relation $\sim$; for an $x \in X$ we denote its equivalence class by $[x]$. Geometrically, this is the space achieved by welding together points equivalent under $\sim$.

Formally,


Definition 20.2.1. Let X be a topological space, and $\sim$ an equivalence relation on the points of X . Then $X/\sim$ is the space whose

- Points are equivalence classes of X , and
- $U \subseteq X/\sim$ is open if and only if $\{x \in X \text{ such that } [x] \in U\}$ is open in X .

As far as I can tell, this definition is mostly useless for intuition, so here are some examples.

Example 20.2.2 (Interval modulo endpoints)

Suppose we take $D^1 = [-1, 1]$ and quotient by the equivalence relation which identifies the endpoints -1 and 1 . (Formally, $x \sim y \iff (x = y) \text{ or } \{x, y\} = \{-1, 1\}$.) In that case, we simply recover S^1 :



Observe that a small neighborhood around $-1 \sim 1$ in the quotient space corresponds to two half-intervals at -1 and 1 in the original space D^1 . This should convince you the definition we gave is the right one.

Example 20.2.3 (More quotient spaces)

Convince yourself that:

- Generalizing the previous example, D^n modulo its boundary S^{n-1} is S^n .
- Given a square $ABCD$, suppose we identify segments AB and DC together. Then we get a cylinder. (Think elementary school, when you would tape up pieces of paper together to get cylinders.)
- In the previous example, if we also identify BC and DA together, then we get a torus. (Imagine taking our cylinder and putting the two circles at the end together.)
- Let $X = \mathbb{R}$, and let $x \sim y$ if $y - x \in \mathbb{Z}$. Then $X/\sim$ is S^1 as well.

One special case that we did above:

Definition 20.2.4. Let $A \subseteq X$. Consider the equivalence relation which identifies all the points of A with each other while leaving all remaining points inequivalent. (In other words, $x \sim y$ if $x = y$ or $x, y \in A$.) Then the resulting quotient space is denoted X/A .

So in this notation,

$$D^n/S^{n-1} = S^n.$$

Abuse of Notation 20.2.5. Note that I'm deliberately being sloppy, and saying " $D^n/S^{n-1} = S^n$ " or " D^n/S^{n-1} is S^n ", when I really ought to say " D^n/S^{n-1} is homeomorphic to S^n ". This is a general theme in mathematics: objects which are homeomorphic/isomorphic/etc. are generally not carefully distinguished from each other.

§20.3 Product topology

Prototypical example for this section: $\mathbb{R} \times \mathbb{R}$ is $\mathbb{R}^2$, $S^1 \times S^1$ is the torus.

Definition 20.3.1. Given topological spaces X and Y , the **product topology** on $X \times Y$ is the space whose

- Points are pairs (x, y) with $x \in X$, $y \in Y$, and
- Topology is given as follows: the *basis* of the topology for $X \times Y$ is $U \times V$, for $U \subseteq X$ open and $V \subseteq Y$ open.

Remark 20.3.2. It is not hard to show that, in fact, one need only consider basis elements for U and V . That is to say,


$$\{U \times V \mid U, V \text{ basis elements for } X, Y\}$$

is also a basis for $X \times Y$.

This does exactly what you think it would.

Example 20.3.3 (The unit square)

Let $X = [0, 1]$ and consider $X \times X$. We of course expect this to be the unit square. Pictured below is an open set of $X \times X$ in the basis.



Exercise 20.3.4. Convince yourself this basis gives the same topology as the product metric on $X \times X$. So this is the “right” definition.

Example 20.3.5 (More product spaces)

- $\mathbb{R} \times \mathbb{R}$ is the Euclidean plane.
- $S^1 \times [0, 1]$ is a cylinder.
- $S^1 \times S^1$ is a torus! (Why?)

§20.4 Disjoint union and wedge sum

Prototypical example for this section: $S^1 \vee S^1$ is the figure eight.

The disjoint union of two spaces is geometrically exactly what it sounds like: you just imagine the two spaces side by side. For completeness, here is the formal definition.

Definition 20.4.1. Let X and Y be two topological spaces. The **disjoint union**, denoted $X \amalg Y$, is defined by

- The points are the disjoint union $X \amalg Y$, and
- A subset $U \subseteq X \amalg Y$ is open if and only if $U \cap X$ and $U \cap Y$ are open.

Exercise 20.4.2. Show that the disjoint union of two nonempty spaces is disconnected.


More interesting is the wedge sum, where two topological spaces X and Y are fused together only at a single base point.

Definition 20.4.3. Let X and Y be topological spaces, and $x_0 \in X$ and $y_0 \in Y$ be points. We define the equivalence relation $\sim$ by declaring $x_0 \sim y_0$ only. Then the **wedge sum** of two spaces is defined as

$$X \vee Y = (X \amalg Y) / \sim.$$

Example 20.4.4 ($S^1 \vee S^1$ is a figure eight)

Let $X = S^1$ and $Y = S^1$, and let $x_0 \in X$ and $y_0 \in Y$ be any points. Then $X \vee Y$ is a “figure eight”: it is two circles fused together at one point.



Abuse of Notation 20.4.5. We often don’t mention x_0 and y_0 when they are understood (or irrelevant). For example, from now on we will just write $S^1 \vee S^1$ for a figure eight.

Remark 20.4.6. Annoyingly, in $\text{\LaTeX}$ `\wedge` gives $\wedge$ instead of $\vee$ (which is `\vee`). So this really should be called the “vee product”, but too late.

§20.5 CW complexes

Using this construction, we can start building some spaces. One common way to do so is using a so-called **CW complex**. Intuitively, a CW complex is built as follows:


- Start with a set of points X^0 .
- Define X^1 by taking some line segments (copies of D^1) and fusing the endpoints (copies of S^0) onto X^0 .
- Define X^2 by taking copies of D^2 (a disk) and welding its boundary (a copy of S^1) onto X^1 .
- Repeat inductively up until a finite stage n ; we say X is **n -dimensional**.

The resulting space X is the CW-complex. The set X^k is called the **k -skeleton** of X . Each D^k is called a **k -cell**; it is customary to denote it by e_α^k where α is some index. We say that X is **finite** if only finitely many cells were used.

Abuse of Notation 20.5.1. Technically, most sources (like [Ha02]) allow one to construct infinite-dimensional CW complexes. We will not encounter any such spaces in the Napkin.

Example 20.5.2 (D^2 with $2 + 2 + 1$ and $1 + 1 + 1$ cells)

- (a) First, we start with X^0 having two points e_a^0 and e_b^0 . Then, we join them with two 1-cells D^1 (green), call them e_c^1 and e_d^1 . The endpoints of each 1-cell (the copy of S^0) get identified with distinct points of X^0 ; hence $X^1 \cong S^1$. Finally, we take a single 2-cell e^2 (yellow) and weld it in, with its boundary fitting into the copy of S^1 that we just drew. This gives the figure on the left.
- (b) In fact, one can do this using just $1 + 1 + 1 = 3$ cells. Start with X^0 having a single point e^0 . Then, use a single 1-cell e^1 , fusing its two endpoints into the single point of X^0 . Then, one can fit in a copy of S^1 as before, giving D^2 as on the right.

**Example 20.5.3** (S^n as a CW complex)

- (a) One can obtain S^n (for $n \geq 1$) with just two cells. Namely, take a single point e^0 for X^0 , and to obtain S^n take D^n and weld its entire boundary into e^0 .

We already saw this example in the beginning with $n = 2$, when we saw that the sphere S^2 was the result when we fuse the boundary of a disk D^2 together.

- (b) Alternatively, one can do a “hemisphere” construction, by constructing S^n inductively using two cells in each dimension. So S^0 consists of two points, then S^1 is obtained by joining these two points by two segments (1-cells), and S^2 is obtained by gluing two hemispheres (each a 2-cell) with S^1 as its equator.

Definition 20.5.4. Formally, for each k -cell e_α^k we want to add to X^k , we take its boundary S_α^{k-1} and weld it onto X^{k-1} via an **attaching map** $S_\alpha^{k-1} \rightarrow X^{k-1}$. Then

$$X^k = X^{k-1} \amalg \left(\coprod_\alpha e_\alpha^k \right) / \sim$$


where $\sim$ identifies each boundary point of e_α^k with its image in X^{k-1} .

§20.6 The torus, Klein bottle, $\mathbb{RP}^n$, $\mathbb{CP}^n$

We now present four of the most important examples of CW complexes.

The torus

The **torus** can be formed by taking a square and identifying the opposite edges in the same direction: if you walk off the right edge, you re-appear at the corresponding point in on the left edge. (Think *Asteroids* from Atari!)



Thus the torus is $(\mathbb{R}/\mathbb{Z})^2 \cong S^1 \times S^1$.

Note that all four corners get identified together to a single point. One can realize the torus in 3-space by treating the square as a sheet of paper, taping together the left and right (red) edges to form a cylinder, then bending the cylinder and fusing the top and bottom (blue) edges to form the torus.






Image from [To]

The torus can be realized as a CW complex with

- A 0-skeleton consisting of a single point,
- A 1-skeleton consisting of two 1-cells e_a^1 , e_b^1 , and




- A 2-skeleton with a single 2-cell e^2 , whose circumference is divided into four parts, and welded onto the 1-skeleton “via $aba^{-1}b^{-1}$ ”. This means: wrap a quarter of the circumference around e_a^1 , then another quarter around e_b^1 , then the third quarter around e_a^1 but in the opposite direction, and the fourth quarter around e_b^1 again in the opposite direction as before.



We say that $aba^{-1}b^{-1}$ is the **attaching word**; this shorthand will be convenient later on.

The Klein bottle

The **Klein bottle** is defined similarly to the torus, except one pair of edges is identified in the opposite manner, as shown.



Unlike the torus one cannot realize this in 3-space without self-intersecting. One can tape together the red edges as before to get a cylinder, but to then fuse the resulting blue circles in opposite directions is not possible in 3D. Nevertheless, we often draw a picture in 3-dimensional space in which we tacitly allow the cylinder to intersect itself.










Image from [In; Fr]


Like the torus, the Klein bottle is realized as a CW complex with

- One 0-cell,


- Two 1-cells e_a^1 and e_b^1 , and
- A single 2-cell attached this time via the word $abab^{-1}$.

Real projective space

Let's start with $n = 2$. The space $\mathbb{RP}^2$ is obtained if we reverse both directions of the square from before, as shown.



However, once we do this the fact that the original polygon is a square is kind of irrelevant; we can combine a red and blue edge to get the single purple edge. Equivalently, one can think of this as a circle with half its circumference identified with the other half:



The resulting space should be familiar to those of you who do projective (Euclidean) geometry. Indeed, there are several possible geometric interpretations:

- One can think of $\mathbb{RP}^2$ as the set of lines through the origin in $\mathbb{R}^3$, with each line being a point in $\mathbb{RP}^2$.

Of course, we can identify each line with a point on the unit sphere S^2 , except for the property that two antipodal points actually correspond to the same line, so that $\mathbb{RP}^2$ can be almost thought of as “half a sphere”. Flattening it gives the picture above.

- Imagine $\mathbb{R}^2$, except augmented with “points at infinity”. This means that we add some lines “infinitely far away”, one for each possible direction of a line. Thus in $\mathbb{RP}^2$, any two lines indeed intersect (at a Euclidean point if they are not parallel, and at a point at infinity if they do).

This gives an interpretation of $\mathbb{RP}^2$, where the boundary represents the *line at infinity* through all of the points at infinity. Here we have used the fact that $\mathbb{R}^2$ and interior of D^2 are homeomorphic.

Exercise 20.6.1. Observe that these formulations are equivalent by consider the plane $z = 1$ in $\mathbb{R}^3$, and intersecting each line in the first formulation with this plane.

We can also express $\mathbb{RP}^2$ using coordinates: it is the set of triples $(x : y : z)$ of real numbers not all zero up to scaling, meaning that

$$(x : y : z) = (\lambda x : \lambda y : \lambda z)$$

for any $\lambda \neq 0$. Using the “lines through the origin in $\mathbb{R}^3$ ” interpretation makes it clear why this coordinate system gives the right space. The points at infinity are those with $z = 0$, and any point with $z \neq 0$ gives a Cartesian point since

$$(x : y : z) = \left(\frac{x}{z} : \frac{y}{z} : 1 \right)$$

hence we can think of it as the Cartesian point $(\frac{x}{z}, \frac{y}{z})$.


In this way we can actually define **real-projective n -space**, $\mathbb{RP}^n$ for any n , as either

- (i) The set of lines through the origin in $\mathbb{R}^{n+1}$,
- (ii) Using $n + 1$ coordinates as above, or
- (iii) As $\mathbb{R}^n$ augmented with points at infinity, which themselves form a copy of $\mathbb{RP}^{n-1}$.

As a possibly helpful example, we give all three pictures of $\mathbb{RP}^1$.

Example 20.6.2 (Real projective 1-Space)

$\mathbb{RP}^1$ can be thought of as S^1 modulo the relation the antipodal points are identified. Projecting onto a tangent line, we see that we get a copy of $\mathbb{R}$ plus a single point at infinity, corresponding to the parallel line (drawn in cyan below).



Thus, the points of $\mathbb{RP}^1$ have two forms:

- $(x : 1)$, which we think of as $x \in \mathbb{R}$ (in dark red above), and
- $(1 : 0)$, which we think of as $1/0 = \infty$, corresponding to the cyan line above.


So, we can literally write

$$\mathbb{RP}^1 = \mathbb{R} \cup \{\infty\}.$$

Note that $\mathbb{RP}^1$ is also the boundary of $\mathbb{RP}^2$. In fact, note also that topologically we have

$$\mathbb{RP}^1 \cong S^1$$

since it is the “real line with endpoints fused together”.



Since $\mathbb{RP}^n$ is just “ $\mathbb{R}^n$ (or D^n) with $\mathbb{RP}^{n-1}$ as its boundary”, we can construct $\mathbb{RP}^n$ as a CW complex inductively. Note that $\mathbb{RP}^n$ thus consists of **one cell in each dimension**.

Example 20.6.3 ($\mathbb{RP}^n$ as a cell complex)

- (a) $\mathbb{RP}^0$ is a single point.
- (b) $\mathbb{RP}^1 \cong S^1$ is a circle, which as a CW complex is a 0-cell plus a 1-cell.
- (c) $\mathbb{RP}^2$ can be formed by taking a 2-cell and wrapping its perimeter twice around a copy of $\mathbb{RP}^1$.

Complex projective space

The **complex projective space** $\mathbb{CP}^n$ is defined like $\mathbb{RP}^n$ with coordinates, i.e.

$$(z_0 : z_1 : \cdots : z_n)$$

under scaling; this time z_i are complex. As before, $\mathbb{CP}^n$ can be thought of as $\mathbb{C}^n$ augmented with some points at infinity (corresponding to $\mathbb{CP}^{n-1}$).


Example 20.6.4 (Complex projective space)

- (a) $\mathbb{CP}^0$ is a single point.
- (b) $\mathbb{CP}^1$ is $\mathbb{C}$ plus a single point at infinity (“complex infinity” if you will). That means as before we can think of $\mathbb{CP}^1$ as

$$\mathbb{CP}^1 = \mathbb{C} \cup \{\infty\}.$$

So, imagine taking the complex plane and then adding a single point to encompass the entire boundary. The result is just sphere S^2 .

Here is a picture of $\mathbb{CP}^1$ with its coordinate system, the **Riemann sphere**.



Remark 20.6.5 (For Euclidean geometers). You may recognize that while $\mathbb{RP}^2$ is the setting for projective geometry, inversion about a circle is done in $\mathbb{CP}^1$ instead. When

one does an inversion sending generalized circles to generalized circles, there is only one point at infinity: this is why we work in $\mathbb{CP}^n$.

Like $\mathbb{RP}^n$, $\mathbb{CP}^n$ is a CW complex, built inductively by taking $\mathbb{C}^n$ and welding its boundary onto $\mathbb{CP}^{n-1}$. The difference is that as topological spaces,

$$\mathbb{C}^n \cong \mathbb{R}^{2n} \cong D^{2n}.$$

Thus, we attach the cells D^0, D^2, D^4 and so on inductively to construct $\mathbb{CP}^n$. Thus we see that

$\mathbb{CP}^n$ consists of one cell in each *even* dimension.

§20.7 Problems to think about

Problem 20A. Show that a space X is Hausdorff if and only if the diagonal $\{(x, x) \mid x \in X\}$ is closed in the product space $X \times X$.

Problem 20B. Realize the following spaces as CW complexes:

- (a) Möbius strip.
- (b) $\mathbb{R}$.
- (c) $\mathbb{R}^n$.

Problem 20C[†]. Show that a finite CW complex is compact.

21 Fundamental groups

Topologists can't tell the difference between a coffee cup and a doughnut. So how do you tell *anything* apart?

This is a very hard question to answer, but one way we can try to answer it is to find some *invariants* of the space. To draw on the group analogy, two groups are clearly not isomorphic if, say, they have different orders, or if one is simple and the other isn't, etc. We'd like to find some similar properties for topological spaces so that we can actually tell them apart.

Two such invariants for a space X are


- Defining homology groups $H_1(X)$, $H_2(X)$, $\dots$
- Defining homotopy groups $\pi_1(X)$, $\pi_2(X)$, $\dots$

Homology groups are hard to define, but in general easier to compute. Homotopy groups are easier to define but harder to compute.

This chapter is about the fundamental group π_1 .

§21.1 Fusing paths together

Recall that a *path* in a space X is a function $[0, 1] \rightarrow X$. Suppose we have paths γ_1 and γ_2 such that $\gamma_1(1) = \gamma_2(0)$. We'd like to fuse¹ them together to get a path $\gamma_1 * \gamma_2$. Easy, right?



We unfortunately do have to hack the definition a tiny bit. In an ideal world, we'd have a path $\gamma_1 : [0, 1] \rightarrow X$ and $\gamma_2 : [1, 2] \rightarrow X$ and we could just merge them together to get $\gamma_1 * \gamma_2 : [0, 2] \rightarrow X$. But the "2" is wrong here. The solution is that we allocate $[0, \frac{1}{2}]$ for the first path and $[\frac{1}{2}, 1]$ for the second path; we run "twice as fast".

¹Almost everyone else in the world uses "gluing" to describe this and other types of constructs. But I was traumatized by Elmer's glue when I was in high school because I hated the stupid "make a poster" projects and hated having to use glue on them. So I refuse to talk about "gluing" paths together, referring instead to "fusing" them together, which sounds cooler anyways.


Definition 21.1.1. Given two paths $\gamma_1, \gamma_2 : [0, 1] \rightarrow X$ such that $\gamma_1(1) = \gamma_2(0)$, we define a path $\gamma_1 * \gamma_2 : [0, 1] \rightarrow X$ by

$$(\gamma_1 * \gamma_2)(t) = \begin{cases} \gamma_1(2t) & 0 \leq t \leq \frac{1}{2} \\ \gamma_2(2t - 1) & \frac{1}{2} \leq t \leq 1. \end{cases}$$

This hack unfortunately reveals a second shortcoming: this “product” is not associative. If we take $(\gamma_1 * \gamma_2) * \gamma_3$ for some suitable paths, then $[0, \frac{1}{4}]$, $[\frac{1}{4}, \frac{1}{2}]$ and $[\frac{1}{2}, 1]$ are the times allocated for $\gamma_1, \gamma_2, \gamma_3$.

Question 21.1.2. What are the times allocated for $\gamma_1 * (\gamma_2 * \gamma_3)$?

But I hope you’ll agree that even though this operation isn’t associative, the reason it fails to be associative is kind of stupid. It’s just a matter of how fast we run in certain parts.



So as long as we’re fusing paths together, we probably don’t want to think of $[0, 1]$ itself too seriously. And so we only consider everything up to (path) homotopy equivalence. (Recall that two paths α and β are homotopic if there’s a path homotopy $F : [0, 1]^2 \rightarrow X$ between them, which is a continuous deformation from α to β .) It is definitely true that

$$(\gamma_1 * \gamma_2) * \gamma_3 \simeq \gamma_1 * (\gamma_2 * \gamma_3).$$

It is also true that if $\alpha_1 \simeq \alpha_2$ and $\beta_1 \simeq \beta_2$ then $\alpha_1 * \beta_1 \simeq \alpha_2 * \beta_2$.


Naturally, homotopy is an equivalence relation, so paths γ lives in some “homotopy type”, the equivalence classes under $\simeq$. We’ll denote this $[\gamma]$. Then it makes sense to talk about $[\alpha] * [\beta]$. Thus, **we can think of $*$ as an operation on homotopy classes**.

§21.2 Fundamental groups

Prototypical example for this section: $\pi_1(\mathbb{R}^2)$ is trivial and $\pi_1(S^1) \cong \mathbb{Z}$.

At this point I'm a little annoyed at keeping track of endpoints, so now I'm going to specialize to a certain type of path.

Definition 21.2.1. A **loop** is a path with $\gamma(0) = \gamma(1)$.




Hence if we restrict our attention to paths starting at a single point x_0 , then we can stop caring about endpoints and start-points, since everything starts and stops at x_0 . We even have a very canonical loop: the “do-nothing” loop² given by standing at x_0 the whole time.

Definition 21.2.2. Denote the trivial “do-nothing loop” by 1. A loop γ is **nulhomotopic** if it is homotopic to 1; i.e. $\gamma \simeq 1$.

For homotopy of loops, you might visualize “reeling in” the loop, contracting it to a single point.

Example 21.2.3 (Loops in S^2 are nulhomotopic)

As the following picture should convince you, every loop in the simply connected space S^2 is nulhomotopic.



(Starting with the purple loop, we contract to the red-brown point.)

Hence to show that spaces are simply connected it suffices to understand the loops of that space. We are now ready to provide:

²Fatty.

Definition 21.2.4. The **fundamental group** of X with basepoint x_0 , denoted $\pi_1(X, x_0)$, is the set of homotopy classes

$$\{[\gamma] \mid \gamma \text{ a loop at } x_0\}$$

equipped with $*$ as a group operation.

It might come as a surprise that this has a group structure. For example, what is the inverse? Let's define it now.

Definition 21.2.5. Given a path $\alpha : [0, 1] \rightarrow X$ we can define a path $\bar{\alpha}$

$$\bar{\alpha}(t) = \alpha(1 - t).$$

In effect, this “runs α backwards”. Note that $\bar{\alpha}$ starts at the endpoint of α and ends at the starting point of α .

Exercise 21.2.6. Show that for any path α , $\alpha * \bar{\alpha}$ is homotopic to the “do-nothing” loop at $\alpha(0)$. (Draw a picture.)

Let's check it.

Proof that this is a group structure. Clearly $*$ takes two loops at x_0 and spits out a loop at x_0 . We also already took the time to show that $*$ is associative. So we only have to check that (i) there's an identity, and (ii) there's an inverse.

- We claim that the identity is the “do-nothing” loop 1 we described above. The reader can check that for any γ ,

$$\gamma \simeq \gamma * 1 = 1 * \gamma.$$

- For a loop γ , recall again we define its “backwards” loop $\bar{\gamma}$ by

$$\bar{\gamma}(t) = \gamma(1 - t).$$

Then we have $\gamma * \bar{\gamma} = \bar{\gamma} * \gamma = 1$.

Hence $\pi_1(X, x_0)$ is actually a group. □

Before going any further I had better give some examples.


Example 21.2.7 (Examples of fundamental groups)

Note that proving the following results is not at all trivial. For now, just try to see intuitively why the claimed answer “should” be correct.

- The fundamental group of $\mathbb{C}$ is the trivial group: in the plane, every loop is nullhomotopic. (Proof: imagine it's a piece of rope and reel it in.)
- On the other hand, the fundamental group of $\mathbb{C} - \{0\}$ (meteor example from earlier) with any base point is actually $\mathbb{Z}$! We won't be able to prove this for a while, but essentially a loop is determined by the number of times that it winds around the origin – these are so-called *winding numbers*. Think about it!
- Similarly, we will soon show that the fundamental group of S^1 (the boundary of the unit circle) is $\mathbb{Z}$.

Officially, I also have to tell you what the base point is, but by symmetry in these examples, it doesn't matter.

Here is the picture for $\mathbb{C} \setminus \{0\}$, with the hole exaggerated as the meteor from last chapter.




Question 21.2.8. Convince yourself that the fundamental group of S^1 is $\mathbb{Z}$, and understand why we call these “winding numbers”. (This will be the most important example of a fundamental group in later chapters, so it’s crucial you figure it out now.)

Example 21.2.9 (The figure eight)

Consider a figure eight $S^1 \vee S^1$, and let x_0 be the center. Then

$$\pi_1(S^1 \vee S^1, x_0) \cong \langle a, b \rangle$$

is the *free group* generated on two letters. The idea is that one loop of the eight is a , and the other loop is b , so we expect π_1 to be generated by this loop a and b (and its inverses $\bar{a}$ and $\bar{b}$). These loops don’t talk to each other.



Recall that in graph theory, we usually assume our graphs are connected, since otherwise we can just consider every connected component separately. Likewise, we generally want to restrict our attention to path-connected spaces, since if a space isn’t path-connected then it can be broken into a bunch of “path-connected components”. (Can you guess how to define this?) Indeed, you could imagine a space X that consists of the objects on my desk (but not the desk itself): π_1 of my phone has nothing to do with π_1 of my mug. They are just totally disconnected, both figuratively and literally.


But on the other hand we claim that in a path-connected space, the groups are very related!

Theorem 21.2.10 (Fundamental groups don’t depend on basepoint)

Let X be a path-connected space. Then for any $x_1 \in X$ and $x_2 \in X$, we have

$$\pi_1(X, x_1) \cong \pi_1(X, x_2).$$

Before you read the proof, see if you can guess the isomorphism based just on the picture below.



Proof. Let α be any path from x_1 to x_2 (possible by path-connectedness), and let $\bar{\alpha}$ be its reverse. Then we can construct a map

$$\pi_1(X, x_1) \rightarrow \pi_1(X, x_2) \text{ by } [\gamma] \mapsto [\bar{\alpha} * \gamma * \alpha].$$

In other words, given a loop γ at x_1 , we can start at x_2 , follow $\bar{\alpha}$ to x_1 , run γ , then run along α home to x_2 . Hence this is a map which builds a loop of $\pi_1(X, x_2)$ from every loop at $\pi_1(X, x_1)$. It is a *homomorphism* of the groups just because

$$(\bar{\alpha} * \gamma_1 * \alpha) * (\bar{\alpha} * \gamma_2 * \alpha) = \bar{\alpha} * \gamma_1 * \gamma_2 * \alpha$$

as $\alpha * \bar{\alpha}$ is nullhomotopic.

Similarly, there is a homomorphism

$$\pi_1(X, x_2) \rightarrow \pi_1(X, x_1) \text{ by } [\gamma] \mapsto [\alpha * \gamma * \bar{\alpha}].$$

As these maps are mutual inverses, it follows they must be isomorphisms. End of story. $\square$

This is a bigger reason why we usually only care about path-connected spaces.

Abuse of Notation 21.2.11. For a path-connected space X we will often abbreviate $\pi_1(X, x_0)$ to just $\pi_1(X)$, since it doesn't matter which $x_0 \in X$ we pick.

Finally, recall that we originally defined “simply connected” as saying that any two paths with matching endpoints were homotopic. It's possible to weaken this condition and then rephrase it using fundamental groups.

Exercise 21.2.12. Let X be a path-connected space. Prove that X is **simply connected** if and only if $\pi_1(X)$ is the trivial group. (One direction is easy; the other is a little trickier.)

This is the “usual” definition of simply connected.

§21.3 Fundamental groups are functorial

One quick shorthand I will introduce to clean up the discussion:

Definition 21.3.1. By $f : (X, x_0) \rightarrow (Y, y_0)$, we will mean that $f : X \rightarrow Y$ is a continuous function of spaces which also sends the point x_0 to y_0 .

Let X and Y be topological spaces and $f : (X, x_0) \rightarrow (Y, y_0)$. We now want to relate the fundamental groups of X and Y .

Recall that a loop γ in (X, x_0) is a map $\gamma : [0, 1] \rightarrow X$ with $\gamma(0) = \gamma(1) = x_0$. Then if we consider the composition

$$[0, 1] \xrightarrow{\gamma} (X, x_0) \xrightarrow{f} (Y, y_0)$$

then we get straight-away a loop in Y at y_0 ! Let's call this loop $f_{\#}\gamma$.

Lemma 21.3.2 ($f_{\#}$ is homotopy invariant)

If $\gamma_1 \simeq \gamma_2$ are path-homotopic, then in fact

$$f_{\#}\gamma_1 \simeq f_{\#}\gamma_2.$$

Proof. Just take the homotopy h taking γ_1 to γ_2 and consider $f \circ h$. □

It's worth noting at this point that if X and Y are homeomorphic, then their fundamental groups are all isomorphic. Indeed, let $f : X \rightarrow Y$ and $g : Y \rightarrow X$ be mutually inverse continuous maps. Then one can check that $f_{\#} : \pi_1(X, x_0) \rightarrow \pi_1(Y, y_0)$ and $g_{\#} : \pi_1(Y, y_0) \rightarrow \pi_1(X, x_0)$ are inverse maps between the groups (assuming $f(x_0) = y_0$ and $g(y_0) = x_0$).

§21.4 Higher homotopy groups

Why the notation π_1 for the fundamental group? And what are $\pi_2, \dots$? The answer lies in the following rephrasing:

Question 21.4.1. Convince yourself that a loop is the same thing as a continuous function $S^1 \rightarrow X$.

It turns out we can define homotopy for things other than paths. Two functions $f, g : Y \rightarrow X$ are **homotopic** if there exists a continuous function $Y \times [0, 1] \rightarrow X$ which continuously deforms f to g . So everything we did above was just the special case $Y = S^1$.

For general n , the group $\pi_n(X)$ is defined as the homotopy classes of the maps $S^n \rightarrow X$. The group operation is a little harder to specify. You have to show that S^n is homeomorphic to $[0, 1]^n$ with some endpoints fused together; for example S^1 is $[0, 1]$ with 0 fused to 1. Once you have these cubes, you can merge them together on a face. (Again, I'm being terribly imprecise, deliberately.)

For $n \neq 1$, π_n behaves somewhat differently than π_1 . (You might not be surprised, as S^n is simply connected for all $n \geq 2$ but not when $n = 1$.) In particular, it turns out that $\pi_n(X)$ is an abelian group for all $n \geq 2$.

Let's see some examples.

Example 21.4.2 ($\pi_n(S^n) \cong \mathbb{Z}$)

As we saw, $\pi_1(S^1) \cong \mathbb{Z}$; given the base circle S^1 , we can wrap a second circle around it as many times as we want. In general, it's true that $\pi_n(S^n) \cong \mathbb{Z}$.

Example 21.4.3 ($\pi_n(S^m) \cong \{1\}$ when $n < m$)

We saw that $\pi_1(S^2) \cong \{1\}$, because a circle in S^2 can just be reeled in to a point. It turns out that similarly, any smaller n -dimensional sphere can be reeled in on the surface of a bigger m -dimensional sphere. So in general, $\pi_n(S^m)$ is trivial for $n < m$.

However, beyond these observations, the groups behave quite weirdly. Here is a table of $\pi_n(S^m)$ for some values of m and n taken from Wikipedia.

$\pi_n(S^m)$	$m = 1$	2	3	4	5	6	7	8	9	10
$n = 1$	$\mathbb{Z}$	$\{1\}$	$\{1\}$	$\{1\}$	$\{1\}$	$\{1\}$	$\{1\}$	$\{1\}$	$\{1\}$	$\{1\}$
2		$\mathbb{Z}$	$\mathbb{Z}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_{12}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_3$	$\mathbb{Z}_{15}$
3			$\mathbb{Z}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_{12}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_3$	$\mathbb{Z}_{15}$
4				$\mathbb{Z}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z} \times \mathbb{Z}_{12}$	$\mathbb{Z}_2^2$	$\mathbb{Z}_2 \times \mathbb{Z}_2$	$\mathbb{Z}_{24} \times \mathbb{Z}_3$
5					$\mathbb{Z}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_{24}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$
6						$\mathbb{Z}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_{24}$	$\{1\}$
7							$\mathbb{Z}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_{24}$
8								$\mathbb{Z}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$

Actually, it turns out that if you can compute $\pi_n(S^m)$ for every m and n , then you can essentially compute *any* homotopy classes. Thus, computing $\pi_n(S^m)$ is sort of a lost cause in general, and the mixture of chaos and pattern in the above table is a testament to this.

§21.5 Homotopy equivalent spaces

Prototypical example for this section: A disk is homotopy equivalent to a point, an annulus is homotopy equivalent to S^1 .

Up to now I've abused notation and referred to "path homotopy" as just "homotopy" for two paths. I will unfortunately continue to do so (and so any time I say two paths are homotopic, you should assume I mean "path-homotopic"). But let me tell you what the general definition of homotopy is first.

Definition 21.5.1. Let $f, g : X \rightarrow Y$ be continuous functions. A **homotopy** is a continuous function $F : X \times [0, 1] \rightarrow Y$, which we'll write $F_s(x)$ for $s \in [0, 1]$, $x \in X$, such that

$$F_0(x) = f(x) \text{ and } F_1(x) = g(x) \text{ for all } x \in X.$$

If such a function exists, then f and g are **homotopic**.

Intuitively this is once again "deforming f to g ". You might notice this is almost exactly the same definition as path-homotopy, except that f and g are any functions instead of paths, and hence there's no restriction on keeping some "endpoints" fixed through the deformation.

This homotopy can be quite dramatic:

Example 21.5.2

The zero function $z \mapsto 0$ and the identity function $z \mapsto z$ are homotopic as functions $\mathbb{C} \rightarrow \mathbb{C}$. The necessary deformation is

$$[0, 1] \times \mathbb{C} \rightarrow \mathbb{C} \text{ by } (t, z) \mapsto tz.$$

I bring this up because I want to define:

Definition 21.5.3. Let X and Y be continuous spaces. They are **homotopy equivalent** if there exist functions $f : X \rightarrow Y$ and $g : Y \rightarrow X$ such that

- (i) $f \circ g : X \rightarrow X$ is homotopic to the identity map on X , and
- (ii) $g \circ f : Y \rightarrow Y$ is homotopic to the identity map on Y .

If a topological space is homotopy equivalent to a point, then it is said to be **contractible**.

Question 21.5.4. Why are two homeomorphic spaces also homotopy equivalent?

Intuitively, you can think of this as a more generous form of stretching and bending than homeomorphism: we are allowed to compress huge spaces into single points.

Example 21.5.5 ($\mathbb{C}$ is contractible)

Consider the topological spaces $\mathbb{C}$ and the space consisting of the single point $\{0\}$. We claim these spaces are homotopy equivalent (can you guess what f and g are?) Indeed, the two things to check are

- (i) $\mathbb{C} \rightarrow \{0\} \hookrightarrow \mathbb{C}$ by $z \mapsto 0 \mapsto 0$ is homotopy equivalent to the identity on $\mathbb{C}$, which we just saw, and
- (ii) $\{0\} \hookrightarrow \mathbb{C} \rightarrow \{0\}$ by $0 \mapsto 0 \mapsto 0$, which *is* the identity on $\{0\}$.

Here by $\hookrightarrow$ I just mean $\rightarrow$ in the special case that the function is just an “inclusion”.

Remark 21.5.6. $\mathbb{C}$ cannot be *homeomorphic* to a point because there is no bijection of sets between them.

Example 21.5.7 ($\mathbb{C} \setminus \{0\}$ is homotopy equivalent to S^1)

Consider the topological spaces $\mathbb{C} \setminus \{0\}$, the **punctured plane**, and the circle S^1 viewed as a subset of S^1 . We claim these spaces are actually homotopy equivalent! The necessary functions are the inclusion

$$S^1 \hookrightarrow \mathbb{C} \setminus \{0\}$$

and the function

$$\mathbb{C} \setminus \{0\} \rightarrow S^1 \quad \text{by} \quad z \mapsto \frac{z}{|z|}.$$

You can check that these satisfy the required condition.

Remark 21.5.8. On the other hand, $\mathbb{C} \setminus \{0\}$ cannot be *homeomorphic* to S^1 . One can make S^1 disconnected by deleting two points; the same is not true for $\mathbb{C} \setminus \{0\}$.

Example 21.5.9 (Disk = Point, Annulus = Circle.)

By the same token, a disk is homotopic to a point; an annulus is homotopic to a circle. (This might be a little easier to visualize, since it's finite.)

I bring these up because it turns out that

Algebraic topology can't distinguish between homotopy equivalent spaces.

More precisely,

Theorem 21.5.10 (Homotopy equivalent spaces have isomorphic fundamental groups)

Let X and Y be path-connected, homotopy-equivalent spaces. Then $\pi_n(X) \cong \pi_n(Y)$ for every positive integer n .

Proof. Let $\gamma : [0, 1] \rightarrow X$ be a loop. Let $f : X \rightarrow Y$ and $g : Y \rightarrow X$ be maps witnessing that X and Y are homotopy equivalent (meaning $f \circ g$ and $g \circ f$ are each homotopic to the identity). Then the composition

$$[0, 1] \xrightarrow{\gamma} X \xrightarrow{f} Y$$

is a loop in Y and hence f induces a natural homomorphism $\pi_1(X) \rightarrow \pi_1(Y)$. Similarly g induces a natural homomorphism $\pi_1(Y) \rightarrow \pi_1(X)$. The conditions on f and g now say exactly that these two homomorphisms are inverse to each other, meaning the maps are isomorphisms. $\square$

In particular,

Question 21.5.11. What are the fundamental groups of contractible spaces?

That means, for example, that algebraic topology can't tell the following homotopic subspaces of $\mathbb{R}^2$ apart.



§21.6 The pointed homotopy category

This section is meant to be read by those who know some basic category theory. Those of you that don't should come back after reading Chapters 23 and 24. Those of you that do will enjoy how succinctly we can summarize the content of this chapter using categorical notions.

Definition 21.6.1. The **pointed homotopy category** $\mathbf{hTop}_*$ is defined as follows.

- Objects: pairs (X, x_0) of spaces with a distinguished basepoint, and
- Morphisms: *homotopy classes* of continuous functions $(X, x_0) \rightarrow (Y, y_0)$.

In particular, two path-connected spaces are isomorphic in this category exactly when they are homotopy equivalent. Then we can summarize many of the preceding results as follows:

Theorem 21.6.2 (Functorial interpretation of fundamental groups)

There is a functor

$$\pi_1 : \mathbf{hTop}_* \rightarrow \mathbf{Grp}$$

sending

$$\begin{array}{ccc} (X, x_0) & \longrightarrow & \pi_1(X, x_0) \\ f \downarrow & & \downarrow f_{\#} \\ (Y, y_0) & \longrightarrow & \pi_1(Y, y_0) \end{array}$$

This implies several things, like

- The functor bundles the information of $f_{\#}$, including the fact that it respects composition. In the categorical language, $f_{\#}$ is $\pi_1(f)$.
- Homotopic spaces have isomorphic fundamental group (since the spaces are isomorphic in $\mathbf{hTop}$, and functors preserve isomorphism by Theorem 24.2.8). In fact, you'll notice that the proofs of Theorem 24.2.8 and Theorem 21.5.10 are secretly identical to each other.
- If maps $f, g : (X, x_0) \rightarrow (Y, y_0)$ are homotopic, then $f_{\#} = g_{\#}$. This is basically Lemma 21.3.2

Remark 21.6.3. In fact, $\pi_1(X, x_0)$ is the set of arrows $(S^1, 1) \rightarrow (X, x_0)$ in $\mathbf{hTop}_*$, so this is actually a covariant Yoneda functor (Example 24.2.6), except with target $\mathbf{Grp}$ instead of $\mathbf{Set}$.

§21.7 Problems to think about

Problem 21A (Harmonic fan). Exhibit a subspace X of the metric space $\mathbb{R}^2$ which is path-connected but for which a point p can be found such that any r -neighborhood of p with $r < 1$ is not path-connected.



Problem 21B[†] (Special case of Seifert-van Kampen). Let X be a topological space. Suppose U and V are connected open subsets of X , with $X = U \cup V$, so that $U \cap V$ is nonempty and path-connected.

Prove that if $\pi_1(U) = \pi_1(V) = \{1\}$ then $\pi_1(X) = 1$.

Remark 21.7.1. The **Seifert–van Kampen theorem** generalizes this for $\pi_1(U)$ and $\pi_1(V)$ any groups; it gives a formula for calculating $\pi_1(X)$ in terms of $\pi_1(U)$, $\pi_1(V)$, $\pi_1(U \cap V)$. The proof is much the same.

Unfortunately, this does not give us a way to calculate $\pi_1(S^1)$, because it is not possible to write $S^1 = U \cup V$ for $U \cap V$ connected.



Problem 21C (RMM 2013). Let $n \geq 2$ be a positive integer. A stone is placed at each vertex of a regular $2n$ -gon. A move consists of selecting an edge of the $2n$ -gon and swapping the two stones at the endpoints of the edge. Prove that if a sequence of moves swaps every pair of stones exactly once, then there is some edge never used in any move.

(This last problem doesn't technically have anything to do with the chapter, but the “gut feeling” which motivates the solution is very similar.)

22 Covering projections

A few chapters ago we talked about what a fundamental group was, but we didn't actually show how to compute any of them except for the most trivial case of a simply connected space. In this chapter we'll introduce the notion of a *covering projection*, which will let us see how some of these groups can be found.


§22.1 Even coverings and covering projections

Prototypical example for this section: $\mathbb{R}$ covers S^1 .

What we want now is a notion where a big space E , a “covering space”, can be projected down onto a base space B in a nice way. Here is the notion of “nice”:

Definition 22.1.1. Let $p : E \rightarrow B$ be a continuous function. Let U be an open set of B . We call U **evenly covered** (by p) if $p^{-1}(U)$ is a disjoint union of open sets (possibly infinite) such that p restricted to any of these sets is a homeomorphism.

Picture:



All we're saying is that U is evenly covered if its pre-image is a bunch of copies of it. (Actually, a little more: each of the pancakes is homeomorphic to U , but we also require that p is the homeomorphism.)

Definition 22.1.2. A **covering projection** $p : E \rightarrow B$ is a continuous map such that every base point $b \in B$ has a neighborhood $U \ni b$ which is evenly covered by p .

Question 22.1.3. Why must a covering projection be surjective?

Here is the most stupid example of a covering projection.

Example 22.1.4 (Tautological covering projection)

Let's take n disconnected copies of any space B : formally, $E = B \times \{1, \dots, n\}$ with the discrete topology on $\{1, \dots, n\}$. Then there exists a tautological covering projection $E \rightarrow B$ by $(x, m) \mapsto x$; we just project all n copies.

This is a covering projection because *every* open set in B is evenly covered.


This is not really that interesting because $B \times [n]$ is not path-connected.

A much more interesting example is that of $\mathbb{R}$ and S^1 .

Example 22.1.5 (Covering projection of S^1)

Take $p : \mathbb{R} \rightarrow S^1$ by $\theta \mapsto e^{2\pi i \theta}$. This is essentially wrapping the real line into a single helix and projecting it down.

We claim this is a covering projection. Indeed, consider the point $1 \in S^1$ (where we view S^1 as the unit circle in the complex plane). We can draw a small neighborhood of it whose pre-image is a bunch of copies in $\mathbb{R}$.



Note that not all neighborhoods work this time: notably, $U = S^1$ does not work because the pre-image would be the entire $\mathbb{R}$.

Example 22.1.6 (Covering of S^1 by itself)

The map $S^1 \rightarrow S^1$ by $z \mapsto z^3$ is also a covering projection. Can you see why?

Example 22.1.7 (Covering projections of $\mathbb{C} \setminus \{0\}$)

For those comfortable with complex arithmetic,

- (a) The exponential map $\exp : \mathbb{C} \rightarrow \mathbb{C} \setminus \{0\}$ is a covering projection.
- (b) For each n , the n th power map $-^n : \mathbb{C} \setminus \{0\} \rightarrow \mathbb{C} \setminus \{0\}$ is a covering projection.


§22.2 Lifting theorem

Prototypical example for this section: $\mathbb{R}$ covers S^1 .

Now here's the key idea: we are going to try to interpret loops in B as paths in $\mathbb{R}$. This is often much simpler. For example, we had no idea how to compute the fundamental group of S^1 , but the fundamental group of $\mathbb{R}$ is just the trivial group. So if we can interpret loops in S^1 as paths in $\mathbb{R}$, that might (and indeed it does!) make computing $\pi_1(S^1)$ tractable.

Definition 22.2.1. Let $\gamma : [0, 1] \rightarrow B$ be a path and $p : E \rightarrow B$ a covering projection. A **lifting** of γ is a path $\tilde{\gamma} : [0, 1] \rightarrow E$ such that $p \circ \tilde{\gamma} = \gamma$.


Picture:



Example 22.2.2 (Typical example of lifting)

Take $p : \mathbb{R} \rightarrow S^1 \subseteq \mathbb{C}$ by $\theta \mapsto e^{2\pi i \theta}$ (so S^1 is considered again as the unit circle). Consider the path γ in S^1 which starts at $1 \in \mathbb{C}$ and wraps around S^1 once, counterclockwise, ending at 1 again. In symbols, $\gamma : [0, 1] \rightarrow S^1$ by $t \mapsto e^{2\pi i t}$.

Then one lifting $\tilde{\gamma}$ is the path which walks from 0 to 1. In fact, for any integer n , walking from n to $n + 1$ works.



Similarly, the counterclockwise path from $1 \in S^1$ to $-1 \in S^1$ has a lifting: for some integer n , the path from n to $n + \frac{1}{2}$.

The above is the primary example of a lifting. It seems like we have the following structure: given a path γ in B starting at b_0 , we start at any point in the fiber $p^{\text{pre}}(b_0)$. (In our prototypical example, $B = S^1$, $b_0 = 1 \in \mathbb{C}$ and that's why we start at any integer n .) After that we just trace along the path in B , and we get a corresponding path in E .

Question 22.2.3. Take a path γ in S^1 with $\gamma(0) = 1 \in \mathbb{C}$. Convince yourself that once we select an integer $n \in \mathbb{R}$, then there is exactly one lifting starting at n .

It turns out this is true more generally.

Theorem 22.2.4 (Lifting paths)

Suppose $\gamma : [0, 1] \rightarrow B$ is a path with $\gamma(0) = b_0$, and $p : (E, e_0) \rightarrow (B, b_0)$ is a covering projection. Then there exists a *unique* lifting $\tilde{\gamma} : [0, 1] \rightarrow E$ such that $\tilde{\gamma}(0) = e_0$.

Proof. For every point $b \in B$, consider an evenly covered neighborhood U_b in B . Then the family of open sets

$$\{\gamma^{\text{pre}}(U_b) \mid b \in B\}$$

is an open cover of $[0, 1]$. As $[0, 1]$ is compact we can take a finite subcover. Thus we can chop $[0, 1]$ into finitely many disjoint closed intervals $[0, 1] = I_1 \sqcup I_2 \sqcup \cdots \sqcup I_N$ in that order, such that for every I_k , $\gamma(I_k)$ is contained in some U_b .

We'll construct $\tilde{\gamma}$ interval by interval now, starting at I_1 . Initially, place a robot at $e_0 \in E$ and a mouse at $b_0 \in B$. For each interval I_k , the mouse moves around according to however γ behaves on I_k . But the whole time it's in some evenly covered U_k ; the fact that p is a covering projection tells us that there are several copies of U_k living in E . Exactly one of them, say V_k , contains our robot. So the robot just mimics the mouse until it gets to the end of I_k . Then the mouse is in some new evenly covered U_{k+1} , and we can repeat. $\square$

The theorem can be generalized to a diagram

$$\begin{array}{ccc} & (E, e_0) & \\ \tilde{\gamma} \nearrow & & \downarrow p \\ (Y, y_0) & \xrightarrow{f} & (B, b_0) \end{array}$$

where Y is some general path-connected space, as follows.

Theorem 22.2.5 (General lifting criterion)

Let $f : (Y, y_0) \rightarrow (B, b_0)$ be continuous and consider a covering projection $p : (E, e_0) \rightarrow (B, b_0)$. (As usual, Y, B, E are path-connected.) Then a lifting $\tilde{f}$ with $\tilde{f}(y_0) = e_0$ exists if and only if

$$f_* (\pi_1(Y, y_0)) \subseteq p_* (\pi_1(E, e_0)),$$

i.e. the image of $\pi_1(Y, y_0)$ under f is contained in the image of $\pi_1(E, e_0)$ under p (both viewed as subgroups of $\pi_1(B, b_0)$). If this lifting exists, it is unique.

As p is injective, we actually have $p_* (\pi_1(E, e_0)) \cong \pi_1(E, e_0)$. But in this case we are interested in the actual elements, not just the isomorphism classes of the groups.

Question 22.2.6. What happens if we put $Y = [0, 1]$?

Remark 22.2.7 (Lifting homotopies). Here's another cool special case: Recall that a homotopy can be encoded as a continuous function $[0, 1] \times [0, 1] \rightarrow X$. But $[0, 1] \times [0, 1]$ is also simply connected. Hence given a homotopy $\gamma_1 \simeq \gamma_2$ in the base space B , we can lift it to get a homotopy $\tilde{\gamma}_1 \simeq \tilde{\gamma}_2$ in E .

Another nice application of this result is Chapter 16.

§22.3 Lifting correspondence

Prototypical example for this section: $(\mathbb{R}, 0)$ covers $(S^1, 1)$.

Let's return to the task of computing fundamental groups. Consider a covering projection $p : (E, e_0) \rightarrow (B, b_0)$.

A loop γ can be lifted uniquely to $\tilde{\gamma}$ in E which starts at e_0 and ends at some point e in the fiber $p^{\text{pre}}(b_0)$. You can easily check that this $e \in E$ does not change if we pick a different path γ' homotopic to E .

Question 22.3.1. Look at the picture in Example 22.2.2.

Put one finger at $1 \in S^1$, and one finger on $0 \in \mathbb{R}$. Trace a loop homotopic to γ in S^1 (meaning, you can go backwards and forwards but you must end with exactly one full counterclockwise rotation) and follow along with the other finger in $\mathbb{R}$.

Convince yourself that you have to end at the point $1 \in \mathbb{R}$.

Thus every homotopy class of a loop at b_0 (i.e. an element of $\pi_1(B, b_0)$) can be associated with some e in the fiber of b_0 . The below proposition summarizes this and more.

Proposition 22.3.2

Let $p : (E, e_0) \rightarrow (B, b_0)$ be a covering projection. Then we have a function of sets

$$\Phi : \pi_1(B, b_0) \rightarrow p^{\text{pre}}(b_0)$$

by $[\gamma] \mapsto \tilde{\gamma}(1)$, where $\tilde{\gamma}$ is the unique lifting starting at e_0 . Furthermore,

- If E is path-connected, then Φ is surjective.
- If E is simply connected, then Φ is injective.

Question 22.3.3. Prove that E path-connected implies Φ is surjective. (This is really offensively easy.)

Proof. To prove the proposition, we've done everything except show that E simply connected implies Φ injective. To do this suppose that γ_1 and γ_2 are loops such that $\Phi([\gamma_1]) = \Phi([\gamma_2])$.

Applying lifting, we get paths $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ both starting at some point $e_0 \in E$ and ending at some point $e_1 \in E$. Since E is simply connected that means they are *homotopic*, and we can write a homotopy $F : [0, 1] \times [0, 1] \rightarrow E$ which unites them. But then consider the composition of maps

$$[0, 1] \times [0, 1] \xrightarrow{F} E \xrightarrow{p} B.$$

You can check this is a homotopy from γ_1 to γ_2 . Hence $[\gamma_1] = [\gamma_2]$, done. $\square$

This motivates:

Definition 22.3.4. A **universal cover** of a space B is a covering projection $p : E \rightarrow B$ where E is simply connected (and in particular path-connected).

Abuse of Notation 22.3.5. When p is understood, we sometimes just say E is the universal cover.

Example 22.3.6 (Fundamental group of S^1)

Let's return to our standard $p : \mathbb{R} \rightarrow S^1$. Since $\mathbb{R}$ is simply connected, this is a universal cover of S^1 . And indeed, the fiber of any point in S^1 is a copy of the integers: naturally in bijection with loops in S^1 .

You can show (and it's intuitively obvious) that the bijection

$$\Phi : \pi_1(S^1) \leftrightarrow \mathbb{Z}$$

is in fact a group homomorphism if we equip $\mathbb{Z}$ with its additive group structure. Since it's a bijection, this leads us to conclude $\pi_1(S^1) \cong \mathbb{Z}$.

§22.4 Regular coverings

Prototypical example for this section: $\mathbb{R} \rightarrow S^1$ comes from $n \cdot x = n + x$

Here's another way to generate some coverings. Let X be a topological space and G a group acting on its points. Thus for every g , we get a map $X \rightarrow X$ by

$$x \mapsto g \cdot x.$$

We require that this map is continuous¹ for every $g \in G$, and that the stabilizer of each point in X is trivial. Then we can consider a quotient space X/G defined by fusing any points in the same orbit of this action. Thus the points of X/G are identified with the orbits of the action. Then we get a natural “projection”

$$X \rightarrow X/G$$

by simply sending every point to the orbit it lives in.


Definition 22.4.1. Such a projection is called **regular**. (Terrible, I know.)

Example 22.4.2 ($\mathbb{R} \rightarrow S^1$ is regular)

Let $G = \mathbb{Z}$, $X = \mathbb{R}$ and define the group action of G on X by

$$n \cdot x = n + x$$

You can then think of X/G as “real numbers modulo 1”, with $[0, 1)$ a complete set of representatives and $0 \sim 1$.



So we can identify X/G with S^1 and the associated regular projection is just our usual $\exp : \theta \mapsto e^{2i\pi\theta}$.

¹Another way of phrasing this: the action, interpreted as a map $X \times G \rightarrow X$, should be continuous, where G on the left-hand side is interpreted as a set with the discrete topology.

Example 22.4.3 (The torus)


Let $G = \mathbb{Z} \times \mathbb{Z}$ and $X = \mathbb{R}^2$, and define the group action of G on X by $(m, n) \cdot (x, y) = (m + x, n + y)$. As $[0, 1)^2$ is a complete set of representatives, you can think of it as a unit square with the edges identified. We obtain the torus $S^1 \times S^1$ and a covering projection $\mathbb{R}^2 \rightarrow S^1 \times S^1$.

Example 22.4.4 ($\mathbb{RP}^2$)

Let $G = \mathbb{Z}_2 = \langle T \mid T^2 = 1 \rangle$ and let $X = S^2$ be the surface of the sphere, viewed as a subset of $\mathbb{R}^3$. We'll let G act on X by sending $T \cdot \vec{x} = -\vec{x}$; hence the orbits are pairs of opposite points (e.g. North and South pole).

Let's draw a picture of a space. All the orbits have size two: every point below the equator gets fused with a point above the equator. As for the points on the equator, we can take half of them; the other half gets fused with the corresponding antipodes.

Now if we flatten everything, you can think of the result as a disk with half its boundary: this is $\mathbb{RP}^2$ from before. The resulting space has a name: *real projective 2-space*, denoted $\mathbb{RP}^2$.



This gives us a covering projection $S^2 \rightarrow \mathbb{RP}^2$ (note that the pre-image of a sufficiently small patch is just two copies of it on S^2 .)

Example 22.4.5 (Fundamental group of $\mathbb{RP}^2$)

As above, we saw that there was a covering projection $S^2 \rightarrow \mathbb{RP}^2$. Moreover the fiber of any point has size two. Since S^2 is simply connected, we have a natural bijection $\pi_1(\mathbb{RP}^2)$ to a set of size two; that is,

$$|\pi_1(\mathbb{RP}^2)| = 2.$$

This can only occur if $\pi_1(\mathbb{RP}^2) \cong \mathbb{Z}_2$, as there is only one group of order two!

Question 22.4.6. Show each of the continuous maps $x \mapsto g \cdot x$ is in fact a homeomorphism. (Name its continuous inverse).

§22.5 The algebra of fundamental groups

Prototypical example for this section: S^1 , with fundamental group $\mathbb{Z}$.

Next up, we're going to turn functions between spaces into homomorphisms of fundamental groups.

Let X and Y be topological spaces and $f : (X, x_0) \rightarrow (Y, y_0)$. Recall that we defined a group homomorphism

$$f_{\#} : \pi_1(X, x_0) \rightarrow \pi_1(Y, y_0) \quad \text{by} \quad [\gamma] \mapsto [f \circ \gamma].$$

More importantly, we have:

Proposition 22.5.1

Let $p : (E, e_0) \rightarrow (B, b_0)$ be a covering projection of path-connected spaces. Then the homomorphism $p_{\#} : \pi_1(E, e_0) \rightarrow \pi_1(B, b_0)$ is *injective*. Hence $p_{\#}(\pi_1(E, e_0))$ is an isomorphic copy of $\pi_1(E, e_0)$ as a subgroup of $\pi_1(B, b_0)$.

Proof. We'll show $\ker p_{\#}$ is trivial. It suffices to show if γ is a nulhomotopic loop in B then its lift is nulhomotopic.

By definition, there's a homotopy $F : [0, 1] \times [0, 1] \rightarrow B$ taking γ to the constant loop 1_B . We can lift it to a homotopy $\tilde{F} : [0, 1] \times [0, 1] \rightarrow E$ that establishes $\tilde{\gamma} \simeq \tilde{1}_B$. But 1_E is a lift of 1_B (duh) and lifts are unique. $\square$


Example 22.5.2 (Subgroups of $\mathbb{Z}$)

Let's look at the space S^1 with fundamental group $\mathbb{Z}$. The group $\mathbb{Z}$ has two types of subgroups:

- The trivial subgroup. This corresponds to the canonical projection $\mathbb{R} \rightarrow S^1$, since $\pi_1(\mathbb{R})$ is the trivial group ($\mathbb{R}$ is simply connected) and hence its image in $\mathbb{Z}$ is the trivial group.
- $n\mathbb{Z}$ for $n \geq 1$. This is given by the covering projection $S^1 \rightarrow S^1$ by $z \mapsto z^n$. The image of a loop in the covering S^1 is a “multiple of n ” in the base S^1 .

It turns out that these are the *only* covering projections of S^1 by path-connected spaces: there's one for each subgroup of $\mathbb{Z}$. (We don't care about disconnected spaces because, again, a covering projection via disconnected spaces is just a bunch of unrelated “good” coverings.) For this statement to make sense I need to tell you what it means for two covering projections to be equivalent.

Definition 22.5.3. Fix a space B . Given two covering projections $p_1 : E_1 \rightarrow B$ and $p_2 : E_2 \rightarrow B$ a **map of covering projections** is a continuous function $f : E_1 \rightarrow E_2$ such that $p_2 \circ f = p_1$.



Then two covering projections p_1 and p_2 are isomorphic if there are $f : E_1 \rightarrow E_2$ and $g : E_2 \rightarrow E_1$ such that $f \circ g = \text{id}_{E_2}$ and $g \circ f = \text{id}_{E_1}$.

Remark 22.5.4 (For category theorists). The set of covering projections forms a category in this way.

It's an absolute miracle that this is true more generally: the greatest triumph of covering spaces is the following result. Suppose a space X satisfies some nice conditions, like:

Definition 22.5.5. A space X is called **locally connected** if each point $x \in X$ has a connected neighborhood U .

Definition 22.5.6. A space X is **semi-locally simply connected** if for every point $x \in X$ there is a neighborhood U such that all loops in U are nullhomotopic. (But the contraction need not take place in U .)

Example 22.5.7 (These conditions are weak)

Pretty much every space I've shown you has these two properties. In other words, they are rather mild conditions, and you can think of them as just saying "the space is not too pathological".

Then we get:

Theorem 22.5.8 (Group theory via covering spaces)

Suppose B is a connected, locally connected, semi-locally simply connected space. Then:

- Every subgroup $H \subseteq \pi_1(B)$ corresponds to exactly one covering projection $p : E \rightarrow B$ with E path-connected (up to isomorphism).
(Specifically, H is the image of $\pi_1(E)$ in $\pi_1(B)$ through $p_\#$.)
- Moreover, the *normal* subgroups of $\pi_1(B)$ correspond exactly to the regular covering projections.

Hence it's possible to understand the group theory of $\pi_1(B)$ completely in terms of the covering projections.

Moreover, this is how the "universal cover" gets its name: it is the one corresponding to the trivial subgroup of $\pi_1(B)$. Actually, you can show that it really is universal in the sense that if $p : E \rightarrow B$ is another covering projection, then E is in turn covered by the universal space. More generally, if $H_1 \subseteq H_2 \subseteq G$ are subgroups, then the space corresponding to H_2 can be covered by the space corresponding to H_1 .

§22.6 Problems to think about

VII

Category Theory

23	Objects and morphisms	232
23.1	Motivation: isomorphisms	232
23.2	Categories, and examples thereof	232
23.3	Special objects in categories	236
23.4	Binary products	237
23.5	Equalizers	241
23.6	Monic and epic maps	242
23.7	Problems to think about	243
24	Functors and natural transformations	244
24.1	Many examples of functors	244
24.2	Covariant functors	245
24.3	Contravariant functors	247
24.4	(Optional) Natural transformations	248
24.5	(Optional) The Yoneda lemma	251
24.6	Problems to think about	253
25	Abelian categories	254
25.1	Zero objects, kernels, cokernels, and images	254
25.2	Additive and abelian categories	255
25.3	Exact sequences	257
25.4	The Freyd-Mitchell embedding theorem	258
25.5	Breaking long exact sequences	259
25.6	Problems to think about	260

23 Objects and morphisms

I can't possibly hope to do category theory any justice in these few chapters; thus I'll just give a very high-level overview of how many of the concepts we've encountered so far can be re-cast into categorical terms. So I'll say what a category is, give some examples, then talk about a few things that categories can do. For my examples, I'll be drawing from all the previous chapters; feel free to skip over the examples corresponding to things you haven't seen.

If you're interested in category theory (like I was!), perhaps in what surprising results are true for general categories, I strongly recommend [Le14].

§23.1 Motivation: isomorphisms

From earlier chapters let's recall the definition of an *isomorphism* of two objects:

- Two groups G and H are isomorphic if there was a bijective homomorphism: equivalently, we wanted homomorphisms $\phi : G \rightarrow H$ and $\psi : H \rightarrow G$ which were mutual inverses, meaning $\phi \circ \psi = \text{id}_H$ and $\psi \circ \phi = \text{id}_G$.
- Two metric (or topological) spaces X and Y are isomorphic if there is a continuous bijection $f : X \rightarrow Y$ such that f^{-1} is also continuous.
- Two vector spaces V and W are isomorphic if there is a bijection $T : V \rightarrow W$ which is a linear map. Again, this can be re-cast as saying that T and T^{-1} are continuous maps.
- Two rings R and S are isomorphic if there is a bijective ring homomorphism ϕ ; again, we can re-cast this as two mutually inverse ring homomorphisms.

In each case we have some collections of objects and some maps, and the isomorphisms can be viewed as just maps. Let's use this to motivate the definition of a general *category*.


§23.2 Categories, and examples thereof

Prototypical example for this section: Grp is possibly the most natural example.

Definition 23.2.1. A **category** $\mathcal{A}$ consists of:

- A class of **objects**, denoted $\text{obj}(\mathcal{A})$.
- For any two objects $A_1, A_2 \in \text{obj}(\mathcal{A})$, a class of **arrows** (also called **morphisms** or **maps**) between them. We'll denote the set of these arrows by $\text{Hom}_{\mathcal{A}}(A_1, A_2)$.
- For any $A_1, A_2, A_3 \in \text{obj}(\mathcal{A})$, if $f : A_1 \rightarrow A_2$ is an arrow and $g : A_2 \rightarrow A_3$ is an arrow, we can compose these arrows to get an arrow $g \circ f : A_1 \rightarrow A_3$.

We can represent this in a **commutative diagram**



where $h = g \circ f$. The composition operation $\circ$ is part of the data of $\mathcal{A}$; it must be associative. In the diagram above we say that h **factors** through A_2 .

- Finally, every object $A \in \text{obj}(\mathcal{A})$ has a special **identity arrow** id_A ; you can guess what it does.¹

Abuse of Notation 23.2.2. From now on, by $A \in \mathcal{A}$ we'll mean $A \in \text{obj}(\mathcal{A})$.

Abuse of Notation 23.2.3. You can think of “class” as just “set”. The reason we can't use the word “set” is because of some paradoxical issues with collections which are too large; Cantor's Paradox says there is no set of all sets. So referring to these by “class” is a way of sidestepping these issues.

Now and forever I'll be sloppy and assume all my categories are **locally small**, meaning that $\text{Hom}_{\mathcal{A}}(A_1, A_2)$ is a set for any $A_1, A_2 \in \mathcal{A}$. So elements of $\mathcal{A}$ may not form a set, but the set of morphisms between two *given* objects will always assumed to be a set.

Let's formalize the motivation we began with.

Example 23.2.4 (Basic examples of categories)

- There is a category of groups **Grp**. The data is
 - The objects of **Grp** are the groups.
 - The arrows of **Grp** are the homomorphisms between these groups.
 - The composition $\circ$ in **Grp** is function composition.
- In the same way we can conceive a category **CRing** of (commutative) rings.
- Similarly, there is a category **Top** of topological spaces, whose arrows are the continuous maps.
- There is a category **Top_{*}** of topological spaces with a *distinguished basepoint*; that is, a pair (X, x_0) where $x_0 \in X$. Arrows are continuous maps $f : X \rightarrow Y$ with $f(x_0) = y_0$.
- Similarly, there is a category **Vect_k** of vector spaces (possibly infinite-dimensional) over a field k , whose arrows are the linear maps. There is even a category **FDVect_k** of *finite-dimensional* vector spaces.
- We have a category **Set** of sets, where the arrows are *any* maps.

And of course, we can now define what an isomorphism is!

¹To be painfully explicit: if $f : A' \rightarrow A$ is an arrow then $\text{id}_A \circ f = f$; similarly, if $g : A \rightarrow A'$ is an arrow then $g \circ \text{id}_A = g$.

Definition 23.2.5. An arrow $A_1 \xrightarrow{f} A_2$ is an **isomorphism** if there exists $A_2 \xrightarrow{g} A_1$ such that $f \circ g = \text{id}_{A_2}$ and $g \circ f = \text{id}_{A_1}$. In that case we say A_1 and A_2 are **isomorphic**, hence $A_1 \cong A_2$.

Remark 23.2.6. Note that in Set , $X \cong Y \iff |X| = |Y|$.

Question 23.2.7. Check that every object in a category is isomorphic to itself. (This is offensively easy.)

More importantly, this definition should strike you as a little impressive. We're able to define whether two groups (rings, spaces, etc.) are isomorphic solely by the functions between the objects. Indeed, one of the key themes in category theory (and even algebra) is that

One can learn about objects by the functions between them. Category theory takes this to the extreme by *only* looking at arrows, and ignoring what the objects themselves are.

But there are some trickier interesting examples of categories.

Example 23.2.8 (Posets are categories)


Let $\mathcal{P}$ be a partially ordered set. We can construct a category P for it as follows:

- The objects of P are going to be the elements of $\mathcal{P}$.
- The arrows of P are defined as follows:
 - For every object $p \in P$, we add an identity arrow id_p , and
 - For any pair of distinct objects $p \leq q$, we add a single arrow $p \rightarrow q$.

There are no other arrows.

- There's only one way to do the composition. What is it?

For example, for the poset $\mathcal{P}$ on four objects $\{a, b, c, d\}$ with $a \leq b$ and $a \leq c \leq d$, we get:



This illustrates the point that

The arrows of a category can be totally different from functions.


In fact, in a way that can be made precise, the term “concrete category” refers to one where the arrows really are “structure-preserving maps between sets”, like $\mathbf{Grp}$, $\mathbf{Top}$, or $\mathbf{CRing}$.

Question 23.2.9. Check that no two distinct objects of a poset are isomorphic.

Here’s a second quite important example of a non-concrete category.

Example 23.2.10 (Important: groups are one-Object categories)

A group G can be interpreted as a category $\mathcal{G}$ with one object $*$, all of whose arrows are isomorphisms.



As [Le14] says:

The first time you meet the idea that a group is a kind of category, it’s tempting to dismiss it as a coincidence or a trick. It’s not: there’s real content. To see this, suppose your education had been shuffled and you took a course on category theory before ever learning what a group was. Someone comes to you and says:

“There are these structures called ‘groups’, and the idea is this: a group is what you get when you collect together all the symmetries of a given thing.”

“What do you mean by a ‘symmetry’?” you ask.

“Well, a symmetry of an object X is a way of transforming X or mapping X into itself, in an invertible way.”

“Oh,” you reply, “that’s a special case of an idea I’ve met before. A category is the structure formed by *lots* of objects and mappings between them – not necessarily invertible. A group’s just the very special case where you’ve only got one object, and all the maps happen to be invertible.”

Exercise 23.2.11. Verify the above! That is, show that the data of a one-object category with all isomorphisms is the same as the data of a group.

Finally, here are some examples of categories you can make from other categories.

Example 23.2.12 (Deriving categories)

- (a) Given a category $\mathcal{A}$, we can construct the **opposite category** $\mathcal{A}^{\text{op}}$, which is the same as $\mathcal{A}$ but with all arrows reversed.
- (b) Given categories $\mathcal{A}$ and $\mathcal{B}$, we can construct the **product category** $\mathcal{A} \times \mathcal{B}$ as follows: the objects are pairs (A, B) for $A \in \mathcal{A}$ and $B \in \mathcal{B}$, and the arrows from (A_1, B_1) to (A_2, B_2) are pairs

$$(A_1 \xrightarrow{f} A_2, B_1 \xrightarrow{g} B_2).$$

What do you think the composition and identities are?

§23.3 Special objects in categories

Prototypical example for this section: $\mathbf{Set}$ has initial object $\emptyset$ and final object $\{*\}$. An element of S corresponds to a map $\{*\} \rightarrow S$.

Certain objects in categories have special properties. Here are a couple examples.

Example 23.3.1 (Initial object)

An **initial object** of $\mathcal{A}$ is an object $A_{\text{init}} \in \mathcal{A}$ such that for any $A \in \mathcal{A}$ (possibly $A = A_{\text{init}}$), there is exactly one arrow from A_{init} to A . For example,

- (a) The initial object of $\mathbf{Set}$ is the empty set $\emptyset$.
- (b) The initial object of $\mathbf{Grp}$ is the trivial group $\{1\}$.
- (c) The initial object of $\mathbf{CRing}$ is the ring $\mathbb{Z}$ (recall that ring homomorphisms $R \rightarrow S$ map 1_R to 1_S).
- (d) The initial object of $\mathbf{Top}$ is the empty space.
- (e) The initial object of a partially ordered set is its smallest element, if one exists.

We will usually refer to “the” initial object of a category, since:

Exercise 23.3.2 (Important!). Show that any two initial objects A_1, A_2 of $\mathcal{A}$ are *uniquely isomorphic* meaning there is a unique isomorphism between them.

Remark 23.3.3. In mathematics, we usually neither know nor care if two objects are actually equal or whether they are isomorphic. For example, there are many competing ways to define $\mathbb{R}$, but we still just refer to it as “the” real numbers.

Thus when we define categorical notions, we would like to check they are unique up to isomorphism. This is really clean in the language of categories, and definitions often cause objects to be unique up to isomorphism for elegant reasons like the above.

One can take the “dual” notion, a terminal object.

Example 23.3.4 (Terminal object)

A **terminal object** of $\mathcal{A}$ is an object $A_{\text{final}} \in \mathcal{A}$ such that for any $A \in \mathcal{A}$ (possibly $A = A_{\text{final}}$), there is exactly one arrow from A to A_{final} . For example,

- (a) The terminal object of **Set** is the singleton set $\{*\}$. (There are many singleton sets, of course, but *as sets* they are all isomorphic!)
- (b) The terminal object of **Grp** is the trivial group $\{1\}$.
- (c) The terminal object of **CRing** is the zero ring (a ring with element $0 = 1$). (Recall that ring homomorphisms $R \rightarrow S$ must map 1_R to 1_S).
- (d) The terminal object of **Top** is the single-point space.
- (e) The terminal object of a partially ordered set is its maximal element, if one exists.

Again, terminal objects are unique up to isomorphism. The reader is invited to repeat the proof from the preceding exercise here. However, we can illustrate more strongly the notion of duality to give a short proof.

Question 23.3.5. Verify that terminal objects of $\mathcal{A}$ are equivalent to initial objects of $\mathcal{A}^{\text{op}}$. Thus terminal objects of $\mathcal{A}$ are unique up to isomorphism.

In general, one can consider in this way the dual of *any* categorical notion: properties of $\mathcal{A}$ can all be translated to dual properties of $\mathcal{A}^{\text{op}}$ (often by adding the prefix “co” in front).

One last neat construction: suppose we’re working in a concrete category, meaning (loosely) that the objects are “sets with additional structure”. Now suppose you’re sick of maps and just want to think about elements of these sets. Well, I won’t let you do that since you’re reading a category theory chapter, but I will offer you some advice:

- In **Set**, arrows from $\{*\}$ to S correspond to elements of S .
- In **Top**, arrows from $\{*\}$ to X correspond to points of X .
- In **Grp**, arrows from $\mathbb{Z}$ to G correspond to elements of G .
- In **CRing**, arrows from $\mathbb{Z}[x]$ to R correspond to elements of R .

and so on. So in most concrete categories, you can think of elements as functions from special sets to the set in question. In each of these cases we call the object in question a **free object**.

§23.4 Binary products

Prototypical example for this section: $X \times Y$ in most concrete categories is the set-theoretic product.

The “universal property” is a way of describing objects in terms of maps in such a way that it defines the object up to unique isomorphism (much the same as the initial and terminal objects).


To show how this works in general, let me give a concrete example. Suppose I’m in a category – let’s say **Set** for now. I have two sets X and Y , and I want to construct the

Cartesian product $X \times Y$ as we know it. The philosophy of category theory dictates that I should talk about maps only, and avoid referring to anything about the sets themselves. How might I do this?

Well, let's think about maps into $X \times Y$. The key observation is that


A function $A \xrightarrow{f} X \times Y$ amounts to a pair of functions $(A \xrightarrow{g} X, A \xrightarrow{h} Y)$.

Put another way, there is a natural projection map $X \times Y \rightarrow X$ and $X \times Y \rightarrow Y$:




(We have to do this in terms of projection maps rather than elements, because category theory forces us to talk about arrows.) Now how do I add A to this diagram? The point is that there is a bijection between functions $A \xrightarrow{f} X \times Y$ and pairs (g, h) of functions. Thus for every pair $A \xrightarrow{g} X$ and $A \xrightarrow{h} Y$ there is a *unique* function $A \xrightarrow{f} X \times Y$.

But $X \times Y$ is special in that it is “universal”: for any *other* set A , if you give me functions $A \rightarrow X$ and $A \rightarrow Y$, I can use it to build a *unique* function $A \rightarrow X \times Y$. Picture:



We can do this in any general category, defining a so-called product.

Definition 23.4.1. Let X and Y be objects in any category $\mathcal{A}$. The **product** consists of an object $X \times Y$ and arrows π_X, π_Y to X and Y (thought of as projection). We require that for any object A and arrows $A \xrightarrow{g} X, A \xrightarrow{h} Y$, there is a *unique* function $A \xrightarrow{f} X \times Y$ such that the diagram



commutes.


Abuse of Notation 23.4.2. Strictly speaking, the product should consist of *both* the object $X \times Y$ and the projection maps π_X and π_Y . However, if π_X and π_Y are understood, then we often use $X \times Y$ to refer to the object, and refer to it also as the product.

Products do not always exist; for example, take a category with just two objects and no non-identity morphisms. Nonetheless:

Proposition 23.4.3 (Uniqueness of products)

When they exist, products are unique up to isomorphism: given two products P_1 and P_2 of X and Y there is an isomorphism between the two objects.

Proof. This is very similar to the proof that initial objects are unique up to unique isomorphism. Consider two such objects P_1 and P_2 , and the associated projection maps. So, we have a diagram



There are unique morphisms f and g between P_1 and P_2 that make the entire diagram commute, according to the universal property.

On the other hand, look at $g \circ f$ and focus on just the outer square. Observe that $g \circ f$ is a map which makes the outer square commute, so by the universal property of P_1 it is the only one. But id_{P_1} works as well. Thus $\text{id}_{P_1} = g \circ f$. Similarly, $f \circ g = \text{id}_{P_2}$ so f and g are isomorphisms. $\square$

Abuse of Notation 23.4.4. Actually, this is not really the morally correct theorem; since we’ve only showed the objects P_1 and P_2 are isomorphic and have not made any assertion about the projection maps. But I haven’t (and won’t) define isomorphism of the entire product, and so in what follows if I say “ P_1 and P_2 are isomorphic” I really just mean the objects are isomorphic.

Exercise 23.4.5. In fact, show the products are unique up to *unique* isomorphism: the f and g above are the only isomorphisms between the objects P_1 and P_2 .

The nice fact about this “universal property” mindset is that we don’t have to give explicit constructions; assuming existence, the “universal property” allows us to bypass all this work by saying “the object with these properties is unique up to unique isomorphism”, thus we don’t need to understand the internal workings of the object to use its properties.


Of course, that’s not to say we can’t give concrete examples.

Example 23.4.6 (Examples of products)

- (a) In **Set**, the product of two sets X and Y is their Cartesian product.
- (b) In **Grp**, the product of G, H is the group product $G \times H$.
- (c) In $\mathbf{Vect}_k$, the product of V and W is $V \oplus W$.
- (d) Let $\mathcal{P}$ be a poset interpreted as a product. Then the product of two objects x and y is the **greatest lower bound**; for example,
 - If the poset is $(\mathbb{R}, \leq)$ then it's $\min\{x, y\}$.
 - If the poset is the subsets of a finite set by inclusion, then it's $x \cap y$.
 - If the poset is the positive integers ordered by division, then it's $\gcd(x, y)$.


Of course, we can define products of more than just one object. Consider a set of objects $(X_i)_{i \in I}$ in a category $\mathcal{A}$. We define a **cone** on the X_i to be an object A with some “projection” maps to each X_i . Then the **product** is a cone P which is “universal” in the same sense as before: given any other cone A there is a unique map $A \rightarrow P$ making the diagram commute. In short, a product is a “universal cone”.

The picture of this is



See also Problem 23B.

One can also do the dual construction to get a **coproduct**: given X and Y , it's the object $X + Y$ together with maps $X \xrightarrow{\iota_X} X + Y$ and $Y \xrightarrow{\iota_Y} X + Y$ (that's Greek iota, think inclusion) such that for any object A and maps $X \xrightarrow{g} A$, $Y \xrightarrow{h} A$ there is a unique f for which



commutes. We'll leave some of the concrete examples as an exercise this time, for example:

Exercise 23.4.7. Describe the coproduct in $\mathbf{Set}$.

Predictable terminology: a coproduct is a universal **cocone**.

§23.5 Equalizers

Prototypical example for this section: The equalizer of $f, g : X \rightarrow Y$ is the set of points with $f(x) = g(x)$.

Given two sets X and Y , and maps $X \xrightarrow{f,g} Y$, we define their **equalizer** to be

$$\{x \in X \mid f(x) = g(x)\}.$$

We would like a categorical way of defining this, too.

Consider two objects X and Y with two maps f and g between them. Stealing a page from [Le14], we call this a **fork**:

$$X \begin{array}{c} \xrightarrow{f} \\ \xrightarrow{g} \end{array} Y$$

A cone over this fork is an object A and arrows over X and Y which make the diagram commute, like so.

$$\begin{array}{ccc} A & & \\ \downarrow q & \searrow f \circ q = g \circ q & \\ X & \xrightarrow[f]{g} & Y \end{array}$$

Effectively, the arrow over Y is just forcing $f \circ q = g \circ q$. In any case, the **equalizer** of f and g is a “universal cone” over this fork: it is an object E and a map $E \xrightarrow{e} X$ such that for each $A \xrightarrow{q} X$ the diagram

$$\begin{array}{ccc} A & & \\ \downarrow q & \searrow \exists! h & \\ E & & \\ \downarrow e & \searrow f & \\ X & \xrightarrow[f]{g} & Y \end{array}$$

commutes for a unique $A \xrightarrow{h} E$. In other words, any map $A \xrightarrow{q} X$ as above must factor uniquely through E . Again, the dotted arrows can be omitted, and as before equalizers may not exist. But when they do exist:

Exercise 23.5.1. If $E \xrightarrow{e} X$ and $E' \xrightarrow{e'} X$ are equalizers, show that $E \cong E'$.

Example 23.5.2 (Examples of equalizers)

- (a) In **Set**, given $X \xrightarrow{f,g} Y$ the equalizer E can be realized as $E = \{x \mid f(x) = g(x)\}$, with the inclusion $e : E \hookrightarrow X$ as the morphism. As usual, by abuse we'll often just refer to E as the equalizer.
- (b) Ditto in **Top**, **Grp**. One has to check that the appropriate structures are preserved (e.g. one should check that $\{\phi(g) = \psi(g) \mid g \in G\}$ is a group).
- (c) In particular, given a homomorphism $\phi : G \rightarrow H$, the inclusion $\ker \phi \hookrightarrow G$ is an equalizer for the fork $G \rightarrow H$ by ϕ and the trivial homomorphism.

According to (c) equalizers let us get at the concept of a kernel if there is a distinguished “trivial map”, like the trivial homomorphism in **Grp**. We'll flesh this idea out in the chapter on abelian categories.

Remark 23.5.3 (Digression on limits). We've defined cones over discrete sets of X_i and over forks. It turns out you can also define a cone over any general **diagram** of objects and arrows; we specify a projection from A to each object and require that the projections from A commute with the arrows in the diagram. (For example, a cone over a fork is a diagram with two edges and two arrows.) If you then demand the cone be universal, you have the extremely general definition of a **limit**. As always, these are unique up to unique isomorphism. We can also define the dual notion of a **colimit** in the same way.

§23.6 Monic and epic maps

The notion of “injective” doesn't make sense in an arbitrary category since arrows need not be functions. The correct categorical notion is:

Definition 23.6.1. A map $X \xrightarrow{f} Y$ is **monic** (or a monomorphism) if for any commutative diagram

$$A \begin{array}{c} \xrightarrow{g} \\ \xrightarrow{h} \end{array} X \xrightarrow{f} Y$$

we must have $g = h$. In other words, $f \circ g = f \circ h \implies g = h$.

Question 23.6.2. Verify that in a *concrete* category, injective $\implies$ monic.

Question 23.6.3. Show that the composition of two monic maps is monic.

In most but not all situations, the converse is also true.

Exercise 23.6.4. Show that in **Set**, **Grp**, **CRing**, monic implies injective. (Take $A = \{*\}$, $A = \{1\}$, $A = \mathbb{Z}[x]$.)

More generally, as we said before there are many categories with a “free” object that you can use to think of as elements. An element of a set is a function $1 \rightarrow S$, and element of a ring is a function $\mathbb{Z}[x] \rightarrow R$, et cetera. In all these categories, the definition of monic literally reads “ f is injective on $\text{Hom}_A(A, X)$ ”. So in these categories, “monic” and “injective” coincide.

That said, here is the standard counterexample. An additive abelian group $G = (G, +)$ is called *divisible* if for every $x \in G$ and $n \in \mathbb{Z}$ there exists $y \in G$ with $ny = x$. Let DivAbGrp be the category of such groups.

Exercise 23.6.5. Show that the projection $\mathbb{Q} \rightarrow \mathbb{Q}/\mathbb{Z}$ is monic but not injective.

Of course, we can also take the dual notion.

Definition 23.6.6. A map $X \xrightarrow{f} Y$ is **epic** (or an epimorphism) if for any commutative diagram

$$X \xrightarrow{f} Y \begin{array}{c} \xrightarrow{g} \\ \xrightarrow{h} \end{array} A$$

we must have $g = h$. In other words, $g \circ f = h \circ f \implies g = h$.

This is kind of like surjectivity, although it's a little farther than last time. Note that in concrete categories, surjective $\implies$ epic.

Exercise 23.6.7. Show that in Set , Grp , Ab , Vect_k , Top , the notions of epic and surjective coincide. (For Set , take $A = \{0, 1\}$.)

However, there are more cases where it fails. Most notably:

Example 23.6.8 (Epic but not surjective)

- (a) In CRing , for instance, the inclusion $\mathbb{Z} \hookrightarrow \mathbb{Q}$ is epic (and not surjective).. Indeed, if two homomorphisms $\mathbb{Q} \rightarrow A$ agree on every integer then they agree everywhere (why?),
- (b) In the category of *Hausdorff* topological spaces (every two points have disjoint neighborhoods), in fact epic $\iff$ dense image (like $\mathbb{Q} \hookrightarrow \mathbb{R}$).

Thus failures arise when a function $f : X \rightarrow Y$ can be determined by just some of the points of X .

§23.7 Problems to think about

Problem 23A[†]. What is the coproduct $X + Y$ in the categories Set , Vect_k , and a poset?

Problem 23B. In any category $\mathcal{A}$ where all products exist, show that

$$(X \times Y) \times Z \cong X \times (Y \times Z)$$

where X, Y, Z are arbitrary objects. (Here both sides refer to the objects, as in Abuse of Notation 23.4.2.)

Problem 23C^{*} (Equalizers are monic). Show that the equalizer of any fork is monic.



Problem 23D. Consider a category $\mathcal{A}$ with a **zero object**, meaning an object which is both initial and terminal. Given objects X and Y in $\mathcal{A}$, prove that the projection $X \times Y \rightarrow X$ is epic.

24 Functors and natural transformations

Functors are maps between categories; natural transformations are maps between functors.

§24.1 Many examples of functors

Prototypical example for this section: Forgetful functors; fundamental groups; $-^\vee$.

Here's the point of a functor:

Pretty much any time you make an object out of another object, you get a functor.

Before I give you a formal definition, let me list (informally) some examples. (You'll notice some of them have opposite categories $\mathcal{A}^{\text{op}}$ appearing in places. Don't worry about those for now; you'll see why in a moment.)

- Given a group G (or vector space, field, $\dots$), we can take its underlying set S ; this is a functor from $\text{Grp} \rightarrow \text{Set}$.
- Given a set S we can consider a vector space with basis S ; this is a functor from $\text{Set} \rightarrow \text{Vect}$.
- Given a vector space V we can consider its dual space $V^\vee$. This is a functor $\text{Vect}_k^{\text{op}} \rightarrow \text{Vect}_k$.
- Tensor products give a functor from $\text{Vect}_k \times \text{Vect}_k \rightarrow \text{Vect}_k$.
- Given a set S , we can build its power set, giving a functor $\text{Set} \rightarrow \text{Set}$.
- In algebraic topology, we take a topological space X and build several groups $H_1(X)$, $\pi_1(X)$, etc. associated to it. All these group constructions are functors $\text{Top} \rightarrow \text{Grp}$.
- Sets of homomorphisms: let $\mathcal{A}$ be a category.
 - Given two vector spaces V_1 and V_2 over k , we construct the abelian group of linear maps $V_1 \rightarrow V_2$. This is a functor from $\text{Vect}_k^{\text{op}} \times \text{Vect}_k \rightarrow \text{AbGrp}$.
 - More generally for any category $\mathcal{A}$ we can take pairs (A_1, A_2) of objects and obtain a set $\text{Hom}_{\mathcal{A}}(A_1, A_2)$. This turns out to be a functor $\mathcal{A}^{\text{op}} \times \mathcal{A} \rightarrow \text{Set}$.
 - The above operation has two “slots”. If we “pre-fill” the first slots, then we get a functor $\mathcal{A} \rightarrow \text{Set}$. That is, by fixing $A \in \mathcal{A}$, we obtain a functor (called H^A) from $\mathcal{A} \rightarrow \text{Set}$ by sending $A' \in \mathcal{A}$ to $\text{Hom}_{\mathcal{A}}(A, A')$. This is called the covariant Yoneda functor (explained later).
 - As we saw above, for every $A \in \mathcal{A}$ we obtain a functor $H^A : \mathcal{A} \rightarrow \text{Set}$. It turns out we can construct a category $[\mathcal{A}, \text{Set}]$ whose elements are functors $\mathcal{A} \rightarrow \text{Set}$; in that case, we now have a functor $\mathcal{A}^{\text{op}} \rightarrow [\mathcal{A}, \text{Set}]$.

§24.2 Covariant functors

Prototypical example for this section: Forgetful/free functors, ...

Category theorists are always asking “what are the maps?”, and so we can now think about maps between categories.

Definition 24.2.1. Let $\mathcal{A}$ and $\mathcal{B}$ be categories. Of course, a **functor** F takes every object of $\mathcal{A}$ to an object of $\mathcal{B}$. In addition, though, it must take every arrow $A_1 \xrightarrow{f} A_2$ to an arrow $F(A_1) \xrightarrow{F(f)} F(A_2)$. You can picture this as follows.

$$\begin{array}{ccc} & A_1 & B_1 = F(A_1) \\ & \downarrow f & \downarrow F(f) \\ \mathcal{A} \ni & \cdots \xrightarrow{F} & \cdots \in \mathcal{B} \\ & \downarrow & \downarrow \\ & A_2 & B_2 = F(A_2) \end{array}$$

(I’ll try to use dotted arrows for functors, which cross different categories, for emphasis.) It needs to satisfy the “naturality” requirements:

- Identity arrows get sent to identity arrows: for each identity arrow id_A , we have $F(\text{id}_A) = \text{id}_{F(A)}$.
- The functor respects composition: if $A_1 \xrightarrow{f} A_2 \xrightarrow{g} A_3$ are arrows in $\mathcal{A}$, then $F(g \circ f) = F(g) \circ F(f)$.

So the idea is:

Whenever we naturally make an object $A \in \mathcal{A}$ into an object of $\mathcal{B} \in \mathcal{B}$, there should usually be a natural way to transform a map $A_1 \rightarrow A_2$ into a map $B_1 \rightarrow B_2$.

Let’s see some examples of this.

Example 24.2.2 (Free and forgetful functors)

Note that these are both informal terms, and don’t have a rigid definition.

- (a) We talked about a **forgetful functor** earlier, which takes the underlying set of a category like Vect_k . Let’s call it $U : \text{Vect}_k \rightarrow \text{Set}$.

Now, given a map $T : V_1 \rightarrow V_2$ in Vect_k , there is an obvious $U(T) : U(V_1) \rightarrow U(V_2)$ which is just the set-theoretic map corresponding to T .

Similarly there are forgetful functors from Grp , CRing , etc. to Set . There is even a forgetful functor $\text{CRing} \rightarrow \text{Grp}$: send a ring R to the abelian group $(R, +)$. The common theme is that we are “forgetting” structure from the original category.

- (b) We also talked about a **free functor** in the example. A free functor $F : \text{Set} \rightarrow \text{Vect}_k$ can be taken by considering $F(S)$ to be the vector space with basis S . Now, given a map $f : S \rightarrow T$, what is the obvious map $F(S) \rightarrow F(T)$? Simple: take each basis element $s \in S$ to the basis element $f(s) \in T$.

Similarly, we can define $F : \text{Set} \rightarrow \text{Grp}$ by taking the free group generated by a set S .

Remark 24.2.3. There is also a notion of “injective” and “surjective” for functors (on arrows) as follows. A functor $F : \mathcal{A} \rightarrow \mathcal{B}$ is **faithful** (resp. **full**) if for any A_1, A_2 , $F : \text{Hom}_{\mathcal{A}}(A_1, A_2) \rightarrow \text{Hom}_{\mathcal{B}}(FA_1, FA_2)$ is injective (resp. surjective).

We can use this to give an exact definition of concrete category: it’s a category with a faithful (forgetful) functor $U : \mathcal{A} \rightarrow \text{Set}$.

Example 24.2.4 (Functors from $\mathcal{G}$)

Let G be a group and $\mathcal{G} = \{*\}$ be the associated one-object category.

- (a) Consider a functor $F : \mathcal{G} \rightarrow \text{Set}$, and let $S = F(*)$. Then the data of F corresponds to putting a *group action* of G on S .
- (b) Consider a functor $F : \mathcal{G} \rightarrow \text{FDVect}_k$, and let $V = F(*)$ have dimension n . Then the data of F corresponds to embedding G as a subgroup of the $n \times n$ matrices (i.e. the linear maps $V \rightarrow V$). This is one way groups historically arose; the theory of viewing groups as matrices forms the field of representation theory.
- (c) Let H be a group and construct $\mathcal{H}$ the same way. Then functors $\mathcal{G} \rightarrow \mathcal{H}$ correspond to homomorphisms $G \rightarrow H$.

Exercise 24.2.5. Check the above group-based functors work as advertised.

Here’s a more involved example. If you find it confusing, skip it and come back after reading about its contravariant version.

Example 24.2.6 (Covariant Yoneda functor)

Fix an $A \in \mathcal{A}$. For a category $\mathcal{A}$, define the **covariant Yoneda functor** $H^A : \mathcal{A} \rightarrow \text{Set}$ by defining

$$H^A(A_1) \stackrel{\text{def}}{=} \text{Hom}_{\mathcal{A}}(A, A_1) \in \text{Set}.$$

Hence each A_1 is sent to the *arrows from A to A_1* ; so H^A **describes how A sees the world**.

Now we want to specify how H^A behaves on arrows. For each arrow $A_1 \xrightarrow{f} A_2$, we need to specify $\text{Set-map } \text{Hom}_{\mathcal{A}}(A, A_1) \rightarrow \text{Hom}_{\mathcal{A}}(A, A_2)$; in other words, we need to send an arrow $A \xrightarrow{p} A_1$ to an arrow $A \rightarrow A_2$. There’s only one reasonable way to do this: take the composition

$$A \xrightarrow{p} A_1 \xrightarrow{f} A_2.$$

In other words, $H_A(f)$ is $p \mapsto f \circ p$. In still other words, $H_A(f) = f \circ -$; the $-$ is a slot for the input to go into.

As another example:

Question 24.2.7. If $\mathcal{P}$ and $\mathcal{Q}$ are posets interpreted as categories, what does a functor from $\mathcal{P}$ to $\mathcal{Q}$ represent?

Now, let me give an explanation of why we might care. Consider the following “obvious” fact: if G and H are isomorphic groups, then they have the same size. We can formalize it by saying: if $G \cong H$ in Grp and $U : \text{Grp} \rightarrow \text{Set}$ is the forgetful functor (mapping each

group to its underlying set), then $U(G) \cong U(H)$. The beauty of category theory shows itself: this in fact works *for any functors and categories*, and the proof is done solely through arrows:

Theorem 24.2.8 (Functors preserve isomorphism)

If $A_1 \cong A_2$ are isomorphic objects in $\mathcal{A}$ and $F : \mathcal{A} \rightarrow \mathcal{B}$ is a functor then $F(A_1) \cong F(A_2)$.

Proof. Try it yourself! The picture is:

$$\mathcal{A} \ni \begin{array}{ccc} & A_1 & \\ \downarrow f & \uparrow g & \\ & A_2 & \end{array} \xrightarrow{F} \begin{array}{ccc} & B_1 = F(A_1) & \\ \downarrow F(f) & \uparrow F(g) & \\ & B_2 = F(A_2) & \end{array} \in \mathcal{B}$$

You'll need to use both key properties of functors: they preserve composition and the identity map. $\square$

This will give us a great intuition in the future, because

- (i) Almost every operation we do in our lifetime will be a functor, and
- (ii) We now know that functors take isomorphic objects to isomorphic objects.

Thus, we now automatically know that basically any “reasonable” operation we do will preserve isomorphism (where “reasonable” means that it’s a functor). This is super convenient in algebraic topology, for example; see Theorem 21.6.2, where we get for free that homotopic spaces have isomorphic fundamental groups.

§24.3 Contravariant functors

Prototypical example for this section: Dual spaces, contravariant Yoneda functor, etc.

Now I have to explain what the opposite categories were doing earlier. In all the previous examples, we took an arrow $A_1 \rightarrow A_2$, and it became an arrow $F(A_1) \rightarrow F(A_2)$. Sometimes, however, the arrow in fact goes the other way: we get an arrow $F(A_2) \rightarrow F(A_1)$ instead. In other words, instead of just getting a functor $\mathcal{A} \rightarrow \mathcal{B}$ we ended up with a functor $\mathcal{A}^{\text{op}} \rightarrow \mathcal{B}$.

These functors have a name:

Definition 24.3.1. A **contravariant functor** from $\mathcal{A}$ to $\mathcal{B}$ is a functor $F : \mathcal{A}^{\text{op}} \rightarrow \mathcal{B}$. (Note that we do *not* write “contravariant functor $F : \mathcal{A} \rightarrow \mathcal{B}$ ”, since that would be confusing; the function notation will always use the correct domain and codomain.)

Pictorially:

$$\mathcal{A} \ni \begin{array}{ccc} & A_1 & \\ \downarrow f & & \\ & A_2 & \end{array} \xrightarrow{F} \begin{array}{ccc} & B_1 = F(A_1) & \\ \uparrow F(f) & & \\ & B_2 = F(A_2) & \end{array} \in \mathcal{B}$$

For emphasis, a usual functor is often called a **covariant functor**. (The word “functor” with no adjective always refers to covariant.)

Let’s see why this might happen.

Example 24.3.2 ($V \mapsto V^\vee$ is contravariant)

Consider the functor $\mathbf{Vect}_k \rightarrow \mathbf{Vect}_k$ by $V \mapsto V^\vee$.

If we were trying to specify a covariant functor, we would need, for every linear map $T : V_1 \rightarrow V_2$, a linear map $T^\vee : V_1^\vee \rightarrow V_2^\vee$. But recall that $V_1^\vee = \text{Hom}(V_1, k)$ and $V_2^\vee = \text{Hom}(V_2, k)$: there’s no easy way to get an obvious map from left to right.

However, there *is* an obvious map from right to left: given $\xi_2 : V_2 \rightarrow k$, we can easily give a map from $V_1 \rightarrow k$: just compose with T ! In other words, there is a very natural map $V_2^\vee \rightarrow V_1^\vee$ according to the composition

$$V_1 \xrightarrow{T} V_2 \xrightarrow{\xi_2} k$$

In summary, a map $T : V_1 \rightarrow V_2$ induces naturally a map $T^\vee : V_2^\vee \rightarrow V_1^\vee$ in the opposite direction. So the contravariant functor looks like:

$$\begin{array}{ccc} V_1 & & V_1^\vee \\ \downarrow T & \xrightarrow{\quad -^\vee \quad} & \uparrow T^\vee \\ V_2 & & V_2^\vee \end{array}$$

We can generalize the example above in any category by replacing the field k with any chosen object $A \in \mathcal{A}$.

Example 24.3.3 (Contravariant Yoneda functor)

The **contravariant Yoneda functor** on $\mathcal{A}$, denoted $H_A : \mathcal{A}^{\text{op}} \rightarrow \mathbf{Set}$, is used to describe how objects of $\mathcal{A}$ see A . For each $X \in \mathcal{A}$ it puts

$$H_A(X) \stackrel{\text{def}}{=} \text{Hom}_{\mathcal{A}}(X, A) \in \mathbf{Set}.$$

For $X \xrightarrow{f} Y$ in $\mathcal{A}$, the map $H_A(f)$ sends each arrow $Y \xrightarrow{p} A \in \text{Hom}_{\mathcal{A}}(Y, A)$ to

$$X \xrightarrow{f} Y \xrightarrow{p} A \in \text{Hom}_{\mathcal{A}}(X, A)$$

as we did above. Thus $H_A(f)$ is an arrow from $\text{Hom}_{\mathcal{A}}(Y, A) \rightarrow \text{Hom}_{\mathcal{A}}(X, A)$. (Note the flipping!)

Exercise 24.3.4. Check now the claim that $\mathcal{A}^{\text{op}} \times \mathcal{A} \rightarrow \mathbf{Set}$ by $(A_1, A_2) \mapsto \text{Hom}(A_1, A_2)$ is in fact a functor.

§24.4 (Optional) Natural transformations


We made categories to keep track of objects and maps, then went a little crazy and asked “what are the maps between categories?” to get functors. Now we’ll ask “what are the maps between functors?” to get natural transformations.

It might sound terrifying that we're drawing arrows between functors, but this is actually an old idea. Recall that given two paths $\alpha, \beta : [0, 1] \rightarrow X$, we built a path-homotopy by “continuously deforming” the path α to β ; this could be viewed as a function $[0, 1] \times [0, 1] \rightarrow X$. The definition of a natural transformation is similar: we want to pull F to G along a series of arrows in the target space $\mathcal{B}$.

Definition 24.4.1. Let $F, G : \mathcal{A} \rightarrow \mathcal{B}$ be two functors. A **natural transformation** α from F to G , denoted

$$\mathcal{A} \begin{array}{c} \xrightarrow{F} \\ \Downarrow \alpha \\ \xrightarrow{G} \end{array} \mathcal{B}$$

consists of, for each $A \in \mathcal{A}$ an arrow $\alpha_A \in \text{Hom}_{\mathcal{B}}(F(A), G(A))$, which is called the **component** of α at A . Pictorially, it looks like this:




These α_A are subject to the “naturality” requirement that for any $A_1 \xrightarrow{f} A_2$, the diagram

$$\begin{array}{ccc} F(A_1) & \xrightarrow{F(f)} & F(A_2) \\ \downarrow \alpha_{A_1} & & \downarrow \alpha_{A_2} \\ G(A_1) & \xrightarrow{G(f)} & G(A_2) \end{array}$$

commutes.

The arrow α_A represents the path that $F(A)$ takes to get to $G(A)$ (just as in a path-homotopy from α to β each *point* $\alpha(t)$ gets deformed to the *point* $\beta(t)$ continuously). A picture might help: consider




Here $\mathcal{A}$ is the small category with three elements and two non-identity arrows f, g (I've omitted the identity arrows for simplicity). The images of $\mathcal{A}$ under F and G are the blue and green “subcategories” of $\mathcal{B}$. Note that $\mathcal{B}$ could potentially have many more objects and arrows in it (grey). The natural transformation α (red) selects an arrow of $\mathcal{B}$ from

each $F(A)$ to the corresponding $G(A)$, dragging the entire image of F to the image of G . Finally, we require that any diagram formed by the blue, red, and green arrows is commutative (naturality), so the natural transformation is really “natural”.

There is a second equivalent definition that looks much more like the homotopy.

Definition 24.4.2. Let $\mathbf{2}$ denote the category generated by a poset with two elements $0 \leq 1$, that is,




Then a *natural transformation* $\mathcal{A} \begin{matrix} \xrightarrow{F} \\ \Downarrow \alpha \\ \xrightarrow{G} \end{matrix} \mathcal{B}$ is just a functor $\alpha : \mathcal{A} \times \mathbf{2} \rightarrow \mathcal{B}$ satisfying

$$\alpha(A, 0) = F(A), \quad \alpha(f, 0) = F(f) \quad \text{and} \quad \alpha(A, 1) = G(A), \quad \alpha(f, 1) = G(f).$$


More succinctly, $\alpha(-, 0) = F$, $\alpha(-, 1) = G$.

The proof that these are equivalent is left as a practice problem.

Naturally, two natural transformations $\alpha : F \rightarrow G$ and $\beta : G \rightarrow H$ can get composed.



Now suppose α is a natural transformation such that α_A is an isomorphism for each A . In this way, we can construct an inverse arrow β_A to it.



In this case, we say α is a **natural isomorphism**. We can then say that $F(A) \cong G(A)$ **naturally** in A . (And β is an isomorphism too!) This means that the functors F and G are “really the same”: not only are they isomorphic on the level of objects, but these isomorphisms are “natural”. As a result of this, we also write $F \cong G$ to mean that the functors are naturally isomorphic.

This is what it really means when we say that “there is a natural / canonical isomorphism”. For example, I claimed earlier (in Section 7.4) that there was a canonical isomorphism $(V^\vee)^\vee \cong V$, and mumbled something about “not having to pick a basis”

and “God-given”. Category theory, amazingly, lets us formalize this: it just says that $(V^\vee)^\vee \cong \text{id}(V)$ naturally in $V \in \text{FDVect}_k$. Really, we have a natural transformation

$$\text{FDVect}_k \begin{array}{c} \xrightarrow{\text{id}} \\ \Downarrow \varepsilon \\ \xrightarrow{(-^\vee)^\vee} \end{array} \text{FDVect}_k .$$

where the component ε_V is given by $v \mapsto \text{ev}_v$ (as discussed earlier, the fact that it is an isomorphism follows from the fact that V and $(V^\vee)^\vee$ have equal dimensions and ε_V is injective).

§24.5 (Optional) The Yoneda lemma

Now that I have natural transformations, I can define:

Definition 24.5.1. The **functor category** of two categories $\mathcal{A}$ and $\mathcal{B}$, denoted $[\mathcal{A}, \mathcal{B}]$, is defined as follows:

- The objects of $[\mathcal{A}, \mathcal{B}]$ are (covariant) functors $F : \mathcal{A} \rightarrow \mathcal{B}$, and
- The morphisms are natural transformations $\alpha : F \rightarrow G$.

Question 24.5.2. When are two objects in the functor category isomorphic?

With this, I can make good on the last example I mentioned at the beginning:

Exercise 24.5.3. Construct the following functors:

- $\mathcal{A} \rightarrow [\mathcal{A}^{\text{op}}, \text{Set}]$ by $A \mapsto H_A$, which we call $H_\bullet$.
- $\mathcal{A}^{\text{op}} \rightarrow [\mathcal{A}, \text{Set}]$ by $A \mapsto H^A$, which we call $H^\bullet$.

Notice that we have opposite categories either way; even if you like H^A because it is covariant, the map $H^\bullet$ is contravariant. So for what follows, we’ll prefer to use $H_\bullet$.

The main observation now is that given a category $\mathcal{A}$, $H_\bullet$ provides some *special* functors $\mathcal{A}^{\text{op}} \rightarrow \text{Set}$ which are already “built” in to the category $\mathcal{A}$. In light of this, we define:

Definition 24.5.4. A **presheaf** X is just a contravariant functor $\mathcal{A}^{\text{op}} \rightarrow \text{Set}$. It is called **representable** if $X \cong H_A$ for some A .

In other words, when we think about representable, the question we’re asking is:

What kind of presheaves are already “built in” to the category $\mathcal{A}$?

One way to get at this question is: given a presheaf X and a particular H_A , we can look at the *set* of natural transformations $\alpha : X \Rightarrow H_A$, and see if we can learn anything about it. In fact, this set can be written explicitly:

Theorem 24.5.5 (Yoneda lemma)

Let $\mathcal{A}$ be a category, pick $A \in \mathcal{A}$, and let H_A be the contravariant Yoneda functor. Let $X : \mathcal{A}^{\text{op}} \rightarrow \mathbf{Set}$ be a contravariant functor. Then the map

$$\left\{ \text{Natural transformations } \mathcal{A}^{\text{op}} \begin{array}{c} \xrightarrow{H_A} \\ \Downarrow \alpha \\ \xrightarrow{X} \end{array} \mathbf{Set} \right\} \rightarrow X(A)$$

defined by $\alpha \mapsto \alpha_A(\text{id}_A) \in X(A)$ is an isomorphism of $\mathbf{Set}$ (i.e. a bijection). Moreover, if we view both sides of the equality as functors

$$\mathcal{A}^{\text{op}} \times [\mathcal{A}^{\text{op}}, \mathbf{Set}] \rightarrow \mathbf{Set}$$

then this isomorphism is natural.

This might be startling at first sight. Here's an unsatisfying explanation why this might not be too crazy: in category theory, a rule of thumb is that "two objects of the same type that are built naturally are probably the same". You can see this theme when we defined functors and natural transformations, and even just compositions. Now to look at the set of natural transformations, we took a pair of elements $A \in \mathcal{A}$ and $X \in [\mathcal{A}^{\text{op}}, \mathbf{Set}]$ and constructed a *set* of natural transformations. Is there another way we can get a set from these two pieces of information? Yes: just look at $X(A)$. The Yoneda lemma is telling us that our heuristic still holds true here.

Some consequences of the Yoneda lemma are recorded in [Le14]. Since this chapter is already a bit too long, I'll just write down the statements, and refer you to [Le14] for the proofs.

1. As we mentioned before, $H^\bullet$ provides a functor

$$\mathcal{A} \rightarrow [\mathcal{A}^{\text{op}}, \mathbf{Set}].$$

It turns out this functor is in fact *fully faithful*; it quite literally embeds the category $\mathcal{A}$ into the functor category on the right (much like Cayley's theorem embeds every group into a permutation group).

2. If $X, Y \in \mathcal{A}$ then

$$H_X \cong H_Y \iff X \cong Y \iff H^X \cong H^Y.$$

To see why this is expected, consider $\mathcal{A} = \mathbf{Grp}$ for concreteness. Suppose A, X, Y are groups such that $H_X(A) \cong H_Y(A)$ for all A . For example,

- If $A = \mathbb{Z}$, then $|X| = |Y|$.
- If $A = \mathbb{Z}/2\mathbb{Z}$, then X and Y have the same number of elements of order 2.
- ...

Each A gives us some information on how X and Y are similar, but the whole natural isomorphism is strong enough to imply $X \cong Y$.

3. Consider the functor $U : \mathbf{Grp} \rightarrow \mathbf{Set}$. It can be represented by $H^{\mathbb{Z}}$, in the sense that

$$\text{Hom}_{\mathbf{Grp}}(\mathbb{Z}, G) \cong U(G) \quad \text{by} \quad \phi \mapsto \phi(1).$$

That is, elements of G are in bijection with maps $\mathbb{Z} \rightarrow G$, determined by the image of $+1$ (or -1 if you prefer). So a representation of U was determined by looking at $\mathbb{Z}$ and picking $+1 \in U(\mathbb{Z})$.

The generalization of this is as follows: let $\mathcal{A}$ be a category and $X : \mathcal{A} \rightarrow \mathbf{Set}$ a covariant functor. Then a representation $H^A \cong X$ consists of an object $A \in \mathcal{A}$ and an element $u \in X(A)$ satisfying a certain condition. You can read this off the condition¹ if you know what the inverse map is in Theorem 24.5.5. In the above situation, $X = U$, $A = \mathbb{Z}$ and $u = \pm 1$.

§24.6 Problems to think about

Problem 24A. Show that the two definitions of natural transformation (one in terms of $\mathcal{A} \times \mathbf{2} \rightarrow \mathcal{B}$ and one in terms of arrows $F(A) \xrightarrow{\alpha_A} G(A)$) are equivalent.

Problem 24B. Let $\mathcal{A}$ be the category of finite sets whose arrows are bijections between sets. For $A \in \mathcal{A}$, let $F(A)$ be the set of *permutations* of A and let $G(A)$ be the set of *orderings* on A .²

- (a) Extend F and G to functors $\mathcal{A} \rightarrow \mathbf{Set}$.
- (b) Show that $F(A) \cong G(A)$ for every A , but this isomorphism is *not* natural.

Problem 24C (Proving the Yoneda lemma). In the context of Theorem 24.5.5:

- (a) Prove that the map described is in fact a bijection. (To do this, you will probably have to explicitly write down the inverse map.)



- (b) Prove that the bijection is indeed natural. (This is long-winded, but not difficult; from start to finish, there is only one thing you can possibly do.)

¹Just for completeness, the condition is: For all $A' \in \mathcal{A}$ and $x \in X(A')$, there's a unique $f : A \rightarrow A'$ with $(Xf)(u) = x$.

² A permutation is a bijection $A \rightarrow A$, and an ordering is a bijection $\{1, \dots, n\} \rightarrow A$, where n is the size of A .

25 Abelian categories

In this chapter I'll translate some more familiar concepts into categorical language; this will require some additional assumptions about our category, culminating in the definition of a so-called “abelian category”. Once that's done, I'll be able to tell you what this “diagram chasing” thing is all about.

Throughout this chapter, “ $\hookrightarrow$ ” will be used for monic maps and “ $\twoheadrightarrow$ ” for epic maps.

§25.1 Zero objects, kernels, cokernels, and images

Prototypical example for this section: In $\mathbf{Grp}$, the trivial group and homomorphism are the zero objects and morphisms. If G, H are abelian then $\phi : G \rightarrow H$ is $H/\text{im } \phi$.

A **zero object** of a category is an object 0 which is both final and terminal; of course, it's unique up to unique isomorphism. For example, in $\mathbf{Grp}$ the zero object is the trivial group, in $\mathbf{Vect}_k$ it's the zero-dimensional vector space consisting of one point, and so on.

Question 25.1.1. Show that $\mathbf{Set}$ and $\mathbf{Top}$ don't have zero objects.

For the rest of this chapter, all categories will have zero objects.

In a category $\mathcal{A}$ with zero objects, any two objects A and B thus have a distinguished morphism

$$A \rightarrow 0 \rightarrow B$$

which is called the **zero morphism** and also denoted 0 . For example, in $\mathbf{Grp}$ this is the trivial homomorphism.

We can now define:

Definition 25.1.2. Consider a map $A \xrightarrow{f} B$. The **kernel** is defined as the equalizer of this map and the map $A \xrightarrow{0} B$. Thus, it's a map $\ker f : \text{Ker } f \hookrightarrow A$ such that

$$\begin{array}{ccc} & \text{Ker } f & \\ \ker f \downarrow & \searrow 0 & \\ A & \xrightarrow{f} & B \end{array}$$

commutes, and moreover any other map with the same property factors uniquely through $\text{Ker } f$ (so it is universal with this property). By Problem 23C*, $\ker f$ is a monic morphism, which justifies the use of “ $\hookrightarrow$ ”.

Notice that we're using $\ker f$ to represent the map and $\text{Ker } f$ to represent the object. Similarly, we define the cokernel, the dual notion:

Definition 25.1.3. Consider a map $A \xrightarrow{f} B$. The **cokernel** of f is a map $\text{coker } f : B \twoheadrightarrow \text{Coker } f$ such that

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ & \searrow \scriptstyle o & \downarrow \scriptstyle \text{coker } f \\ & & \text{Coker } f \end{array}$$

commutes, and moreover any other map with the same property factors uniquely through $\text{Coker } f$ (so it is universal with this property). Thus it is the “coequalizer” of this map and the map $A \xrightarrow{0} B$. By the dual of Problem 23C*, $\text{coker } f$ is an epic morphism, which justifies the use of “ $\twoheadrightarrow$ ”.

Think of the cokernel of a map $A \xrightarrow{f} B$ as “ B modulo the image of f ”, e.g.

Example 25.1.4 (Cokernels)

Consider the map $\mathbb{Z}_6 \rightarrow D_{12} = \langle r, s \mid r^6 = s^2 = 1, rs = sr^{-1} \rangle$. Then the cokernel of this map in Grp is $D_{12}/\langle r \rangle \cong \mathbb{Z}_2$.

This doesn’t always work out quite the way we want since in general the image of a homomorphism need not be normal in the codomain. Nonetheless, we can use this to define:

Definition 25.1.5. The **image** of $A \xrightarrow{f} B$ is the kernel of $\text{coker } f$. We denote $\text{Im } f = \text{Ker}(\text{coker } f)$. This gives a unique map $\text{im } f : A \rightarrow \text{Im } f$.

When it exists, this coincides with our concrete notion of “image”. Picture:

$$\begin{array}{ccccc} A & \xrightarrow{f} & B & & \\ & \searrow \scriptstyle o & \downarrow \scriptstyle \text{coker } f & & \\ & & \text{Im } f & \xrightarrow{\scriptstyle 0} & \text{Coker } f \end{array}$$

Note that by universality of $\text{Im } f$, we find that there is a unique map $\text{im } f : A \rightarrow \text{Im } f$ that makes the entire diagram commute.

§25.2 Additive and abelian categories

Prototypical example for this section: Ab , Vect_k , or more generally Mod_R .

We can now define the notion of an additive and abelian category, which are the types of categories where this notion is most useful.

Definition 25.2.1. An **additive category** $\mathcal{A}$ is one such that:

- $\mathcal{A}$ has a zero object, and any two objects have a product.

- More importantly: every $\text{Hom}_{\mathcal{A}}(A, B)$ forms an *abelian group* (written additively) such that composition distributes over addition:

$$(g + h) \circ f = g \circ f + h \circ f \quad \text{and} \quad f \circ (g + h) = f \circ g + f \circ h.$$

The zero map serves as the identity element for each group.

Definition 25.2.2. An **abelian category** $\mathcal{A}$ is one with the additional properties that for any morphism $A \xrightarrow{f} B$,

- The kernel and cokernel exist, and
- The morphism factors through the image so that $\text{im}(f)$ is epic.

So, this yields a diagram

$$\text{Ker}(f) \xrightarrow{\text{ker}(f)} A \xrightarrow{\text{im}(f)} \text{Im}(f) \hookrightarrow B \xrightarrow{\text{coker}(f)} \text{Coker}(f).$$

Example 25.2.3 (Examples of abelian categories)

- (a) Vect_k , Ab are abelian categories, where $f + g$ takes its usual meaning.
- (b) Generalizing this, the category Mod_R of R -modules is abelian.
- (c) Grp is not even additive, because there is no way to assign a commutative addition to pairs of morphisms.

In general, once you assume a category is abelian, all the properties you would want of these kernels, cokernels, ... that you would guess hold true. For example,

Proposition 25.2.4 (Monic $\iff$ trivial kernel)

A map $A \xrightarrow{f} B$ is monic if and only if its kernel is $0 \rightarrow A$. Dually, $A \xrightarrow{f} B$ is epic if and only if its cokernel is $B \rightarrow 0$.

Proof. The easy direction is:

Exercise 25.2.5. Show that if $A \xrightarrow{f} B$ is monic, then $0 \rightarrow A$ is a kernel. (This holds even in non-abelian categories.)

Of course, since kernels are unique up to isomorphism, $\text{monic} \implies 0 \text{ kernel}$. On the other hand, assume that $0 \rightarrow A$ is a kernel of $A \xrightarrow{f} B$. For this we can exploit the group structure of the underlying homomorphisms now. Assume the diagram

$$\begin{array}{ccccc} Z & \xrightarrow{g} & A & \xrightarrow{f} & B \\ & \searrow h & & & \\ & & A & \xrightarrow{f} & B \end{array}$$

commutes. Then $(g - h) \circ f = g \circ f - h \circ f = 0$, and we've arrived at a commutative diagram.

$$\begin{array}{ccc} Z & & \\ \downarrow g-h & \searrow o & \\ A & \xrightarrow{f} & B \end{array}$$

But since $0 \rightarrow A$ is a kernel it follows that $g - h$ factors through 0, so $g - h = 0 \implies g = h$, which is to say that f is monic. $\square$

Proposition 25.2.6 (Isomorphism $\iff$ monic and epic)

In an abelian category, a map is an isomorphism if and only if it is monic and epic.

Proof. Omitted, because the Mitchell embedding theorem presented later implies this. $\square$

§25.3 Exact sequences

Prototypical example for this section: $0 \rightarrow G \rightarrow G \times H \rightarrow H \rightarrow 0$ is exact.

Exact sequences will seem exceedingly unmotivated until you learn about homology groups, which is one of the most natural places that exact sequences appear. In light of this, it might be worth trying to read the chapter on homology groups simultaneously with this one.

First, let me state the definition for groups, to motivate the general categorical definition. A sequence of groups

$$G_0 \xrightarrow{f_1} G_1 \xrightarrow{f_2} G_2 \xrightarrow{f_3} \dots \xrightarrow{f_n} G_n$$

is *exact* at G_k if the image of f_k is the kernel of f_{k+1} . We say the entire sequence is exact if it's exact at $k = 1, \dots, n - 1$.

Example 25.3.1 (Exact sequences)

(a) The sequence

$$0 \rightarrow \mathbb{Z}_3 \xrightarrow{\times 5} \mathbb{Z}_{15} \rightarrow \mathbb{Z}_5 \rightarrow 0$$

is exact. Actually, $0 \rightarrow G \hookrightarrow G \times H \twoheadrightarrow H \rightarrow 0$ is exact in general. (Here 0 denotes the trivial group.)

(b) For groups, the map $0 \rightarrow A \rightarrow B$ is exact if and only if $A \rightarrow B$ is injective.

(c) For groups, the map $A \rightarrow B \rightarrow 0$ is exact if and only if $A \rightarrow B$ is surjective.

Now, we want to mimic this definition in a general *abelian* category $\mathcal{A}$. So, let's write down a criterion for when $A \xrightarrow{f} B \xrightarrow{g} C$ is exact. First, we had better have that $g \circ f = 0$, which encodes the fact that $\text{im}(f) \subseteq \ker(g)$. Adding in all the relevant objects, we get the commutative diagram below.

$$\begin{array}{ccccc}
 A & & \xrightarrow{f} & B & \xrightarrow{g} & C \\
 \downarrow \text{im } f & & & \uparrow \iota & & \uparrow 0 \\
 \text{Im } f & \xrightarrow{\iota} & B & \xleftarrow{\pi} & \text{Ker } g & \xrightarrow{\pi} & 0
 \end{array}$$

$\exists!$

Here the map $A \twoheadrightarrow \text{Im } f$ is epic since we are assuming $\mathcal{A}$ is an abelian category. So, we have that

$$0 = (g \circ \iota) \circ \text{im } f = g \circ (\iota \circ \text{im } f) = g \circ f = 0$$

but since $\text{im } f$ is epic, this means that $g \circ \iota = 0$. So there is a *unique* map $\text{Im } f \rightarrow \text{Ker } g$, and we require that this diagram commutes. In short,

Definition 25.3.2. Let $\mathcal{A}$ be an abelian category. The sequence

$$\cdots \rightarrow A_{n-1} \xrightarrow{f_n} A_n \xrightarrow{f_{n+1}} A_{n+1} \rightarrow \cdots$$

is **exact** at A_n if $f_n \circ f_{n+1} = 0$ and the canonical map $\text{Im } f_n \rightarrow \text{Ker } f_{n+1}$ is an isomorphism. The entire sequence is exact if it is exact at each A_i . (For finite sequences we don't impose condition on the very first and very last object.)

Exercise 25.3.3. Show that, as before, $0 \rightarrow A \rightarrow B$ is exact $\iff A \rightarrow B$ is monic.

§25.4 The Freyd-Mitchell embedding theorem

We now introduce the Freyd-Mitchell embedding theorem, which essentially says that any abelian category can be realized as a concrete one.

Definition 25.4.1. A category is **small** if $\text{obj}(\mathcal{A})$ is a set (as opposed to a class), i.e. there is a “set of all objects in $\mathcal{A}$ ”. For example, **Set** is not small because there is no set of all sets.

Theorem 25.4.2 (Freyd-Mitchell embedding theorem)

Let $\mathcal{A}$ be a small abelian category. Then there exists a ring R (with 1 but possibly non-commutative) and a full, faithful, exact functor onto the category of left R -modules.

Here a functor is **exact** if it preserves exact sequences. This theorem is good because it means

You can basically forget about all the weird definitions that work in any abelian category.

Any time you're faced with a statement about an abelian category, it suffices to just prove it for a “concrete” category where injective/surjective/kernel/image/exact/etc. agree with your previous notions. A proof by this means is sometimes called *diagram chasing*.

Remark 25.4.3. The “small” condition is a technical obstruction that requires the objects $\mathcal{A}$ to actually form a class. I'll ignore this distinction, because one can almost always work around it by doing enough set-theoretic technicalities.

For example, let's prove:

Lemma 25.4.4 (Short five lemma)

In an abelian category, consider the commutative diagram

$$\begin{array}{ccccccc}
 0 & \longrightarrow & A & \xrightarrow{p} & B & \xrightarrow{q} & C \longrightarrow 0 \\
 & & \downarrow \cong \alpha & & \downarrow \beta & & \downarrow \cong \gamma \\
 0 & \longrightarrow & A' & \xrightarrow{p'} & B' & \xrightarrow{q'} & C' \longrightarrow 0
 \end{array}$$

and assume the top and bottom rows are exact. If α and γ are isomorphisms, then so is β .

Proof. We prove that β is epic (with a similar proof to get monic). By the embedding theorem we can treat the category as R -modules over some R . This lets us do a so-called “diagram chase” where we move elements around the picture, using the concrete interpretation of our category as R -modules.

Let $b' \in B'$. Then $q'(b') \in C'$, and since γ is surjective, we have a $c \in C$ such that $\gamma(c) = b'$, and finally a $b \in B$ such that $q(b) = c$. Picture:

$$\begin{array}{ccc}
 b \in B & \xrightarrow{q} & c \in C \\
 \downarrow \beta & & \downarrow \gamma \\
 b' \in B' & \xrightarrow{q'} & c' \in C'
 \end{array}$$

Now, it is not necessarily the case that $\beta(b) = b'$. However, since the diagram commutes we at least have that

$$q'(b') = q'(\beta(b))$$

so $b' - \beta(b) \in \text{Ker } q' = \text{Im } p'$, and there is an $a' \in A'$ such that $p'(a') = b' - \beta(b)$; use α now to lift it to $a \in A$. Picture:

$$\begin{array}{ccc}
 a \in A & & b \in B \\
 \downarrow & & \\
 a' \in A' & \xrightarrow{\quad} & b' - \beta(b) \in B' \xrightarrow{\quad} 0 \in C'
 \end{array}$$

Then, we have

$$\beta(b + q(a)) = \beta b + \beta p a = \beta b + p' \alpha a = \beta b + (b' - \beta b) = b'$$

so $b' \in \text{Im } \beta$ which completes the proof that β is surjective. $\square$

§25.5 Breaking long exact sequences

Prototypical example for this section: First isomorphism theorem.

In fact, it turns out that any exact sequence breaks into short exact sequences. This relies on:

Proposition 25.5.1 (“First isomorphism theorem” in abelian categories)

Let $A \xrightarrow{f} B$ be an arrow of an abelian category. Then there is an exact sequence

$$0 \rightarrow \text{Ker } f \xrightarrow{\text{ker } f} A \xrightarrow{\text{im } f} \text{Im } f \rightarrow 0.$$

Example 25.5.2

Let’s analyze this theorem in our two examples of abelian categories:

- (a) In the category of abelian groups, this is basically the first isomorphism theorem.
- (b) In the category $\mathbf{Vect}_k$, this amounts to the rank-nullity theorem, Theorem 6.7.7.

Thus, any exact sequence can be broken into short exact sequences, as

$$\begin{array}{ccccccc}
 & & 0 & & 0 & & 0 \\
 & & \swarrow & & \searrow & & \swarrow \\
 & & & C_n & & & C_{n+2} \\
 & & \nwarrow & & \nearrow & & \nwarrow \\
 \dots & \longrightarrow & A_{n-1} & \xrightarrow{f_{n-1}} & A_n & \xrightarrow{f_n} & A_{n+1} \xrightarrow{f_{n+1}} \dots \\
 & & \nwarrow & & \nearrow & & \nwarrow \\
 & & C_{n-1} & & C_{n+1} & & \\
 & & \swarrow & & \searrow & & \swarrow \\
 & & 0 & & 0 & & 0
 \end{array}$$

where $C_k = \text{im } f_{k-1} = \text{ker } f_k$ for every k .

§25.6 Problems to think about

Problem 25A (Five lemma). In an abelian category, consider the commutative diagram

$$\begin{array}{ccccccccc}
 A & \xrightarrow{p} & B & \xrightarrow{q} & C & \xrightarrow{r} & D & \xrightarrow{s} & E \\
 \downarrow \alpha & & \downarrow \cong \beta & & \downarrow \gamma & & \downarrow \cong \delta & & \downarrow \varepsilon \\
 A' & \xrightarrow{p'} & B' & \xrightarrow{q'} & C' & \xrightarrow{r'} & D' & \xrightarrow{s'} & E'
 \end{array}$$

where the first and second rows are exact. Prove that if α is epic, ε is monic, and β, δ are isomorphisms, then γ is an isomorphism as well. Thus this is a stronger version of the short five lemma.



Problem 25B* (Snake lemma). In an abelian category, consider the diagram

$$\begin{array}{ccccccc}
 A & \xrightarrow{f} & B & \xrightarrow{g} & C & \longrightarrow & 0 \\
 \downarrow a & & \downarrow b & & \downarrow c & & \\
 0 & \longrightarrow & A' & \xrightarrow{f'} & B' & \xrightarrow{g'} & C'
 \end{array}$$

where the first and second rows are exact sequences. Prove that there is an exact sequence

$$\operatorname{Ker} a \rightarrow \operatorname{Ker} b \rightarrow \operatorname{Ker} c \rightarrow \operatorname{Coker} a \rightarrow \operatorname{Coker} b \rightarrow \operatorname{Coker} c.$$

DRAFT (Evan Chen)
Updated February 15, 2017

VIII

Differential Geometry

26	Multivariable calculus done correctly	263
26.1	The total derivative	263
26.2	The projection principle	265
26.3	Total and partial derivatives	266
26.4	(Optional) A word on higher derivatives	268
26.5	Towards differential forms	269
26.6	Problems to think about	269
27	Differential forms	270
27.1	Pictures of differential forms	270
27.2	Pictures of exterior derivatives	272
27.3	Differential forms	273
27.4	Exterior derivatives	274
27.5	Closed and exact forms	276
27.6	Problems to think about	277
28	Integrating differential forms	278
28.1	Motivation: line integrals	278
28.2	Pullbacks	279
28.3	Cells	280
28.4	Boundaries	282
28.5	Stokes' theorem	284
28.6	Problems to think about	284
29	A bit of manifolds	285
29.1	Topological manifolds	285
29.2	Smooth manifolds	286
29.3	Differential forms on manifolds	287
29.4	Orientations	288
29.5	Stokes' theorem for manifolds	289
29.6	Problems to think about	289

26 Multivariable calculus done correctly

As I have ranted about before, linear algebra is done wrong by the extensive use of matrices to obscure the structure of a linear map. Similar problems occur with multivariable calculus, so here I would like to set the record straight.


Since we are doing this chapter using morally correct linear algebra, it's imperative you're comfortable with linear maps, and in particular the dual space $V^\vee$ which we will repeatedly use.

In this chapter, all vector spaces have norms and are finite-dimensional over $\mathbb{R}$. So in particular every vector space is also a metric space (with norm given by the metric), and we can talk about open sets as usual.

§26.1 The total derivative

Prototypical example for this section: If $f(x, y) = x^2 + y^2$, then $(Df)_{(x,y)} = 2xe_1^\vee + 2ye_2^\vee$.

First, let $f : [a, b] \rightarrow \mathbb{R}$. You might recall from high school calculus that for every point $p \in \mathbb{R}$, we defined $f'(p)$ as the derivative at the point p (if it existed), which we interpreted as the *slope* of the “tangent line”.



That's fine, but I claim that the “better” way to interpret the derivative at that point is as a *linear map*, that is, as a *function*. If $f'(p) = 1.5$, then the derivative tells me that if I move ε away from p then I should expect f to change by about 1.5ε . In other words,

The derivative of f at p approximates f near p by a *linear function*.

What about more generally? Suppose I have a function like $f : \mathbb{R}^2 \rightarrow \mathbb{R}$, say

$$f(x, y) = x^2 + y^2$$

for concreteness or something. For a point $p \in \mathbb{R}^2$, the “derivative” of f at p ought to represent a linear map that approximates f at that point p . That means I want a linear map $T : \mathbb{R}^2 \rightarrow \mathbb{R}$ such that

$$f(p + v) \approx f(p) + T(v)$$

for small displacements $v \in \mathbb{R}^2$.

Even more generally, if $f : U \rightarrow W$ with $U \subseteq V$ open (in the $\|\bullet\|_V$ metric as usual), then the derivative at $p \in U$ ought to be so that

$$f(p + v) \approx f(p) + T(v) \in W.$$

(We need U open so that for small enough v , $p + v \in U$ as well.) In fact this is exactly what we were doing earlier with $f'(p)$ in high school.




Image derived from [gk]

The only difference is that, by an unfortunate coincidence, a linear map $\mathbb{R} \rightarrow \mathbb{R}$ can be represented by just its slope. And in the unending quest to make everything a number so that it can be AP tested, we immediately forgot all about what we were trying to do in the first place and just defined the derivative of f to be a *number* instead of a *function*.

The fundamental idea of Calculus is the local approximation of functions by linear functions. The derivative does exactly this.

Jean Dieudonné as quoted in [Pu02] continues:

In the classical teaching of Calculus, this idea is immediately obscured by the accidental fact that, on a one-dimensional vector space, there is a one-to-one correspondence between linear forms and numbers, and therefore the derivative at a point is defined as a number instead of a linear form. This **slavish subservience to the shibboleth of numerical interpretation at any cost** becomes much worse . . .

So let's do this right. The only thing that we have to do is say what " $\approx$ " means, and for this we use the norm of the vector space.

Definition 26.1.1. Let $U \subseteq V$ be open. Let $f : U \rightarrow W$ be a continuous function, and $p \in U$. Suppose there exists a linear map $T : V \rightarrow W$ such that

$$\lim_{\|v\|_V \rightarrow 0} \frac{\|f(p + v) - f(p) - T(v)\|_W}{\|v\|_V} = 0.$$

Then T is the **total derivative** of f at p . We denote this by $(Df)_p$, and say f is **differentiable at p** .

If $(Df)_p$ exists at every point, we say f is **differentiable**.

Question 26.1.2. Check if that $V = W = \mathbb{R}$, this is equivalent to the single-variable definition. (What are the linear maps from V to W ?)

Example 26.1.3 (Total derivative of $f(x, y) = x^2 + y^2$)

Let $V = \mathbb{R}^2$ with standard basis $\mathbf{e}_1, \mathbf{e}_2$ and let $W = \mathbb{R}$, and let $f(x\mathbf{e}_1 + y\mathbf{e}_2) = x^2 + y^2$. Let $p = a\mathbf{e}_1 + b\mathbf{e}_2$. Then, we claim that

$$(Df)_p : \mathbb{R}^2 \rightarrow \mathbb{R} \quad \text{by} \quad v \mapsto 2a \cdot \mathbf{e}_1^\vee(v) + 2b \cdot \mathbf{e}_2^\vee(v).$$

Here, the notation $\mathbf{e}_1^\vee$ and $\mathbf{e}_2^\vee$ makes sense, because by definition $(Df)_p \in V^\vee$: these are functions from V to $\mathbb{R}$!

Let's check this manually with the limit definition. Set $v = x\mathbf{e}_1 + y\mathbf{e}_2$, and note that the norm on V is $\|(x, y)\|_V = \sqrt{x^2 + y^2}$ while the norm on W is just the absolute value $\|c\|_W = |c|$. Then we compute

$$\begin{aligned} \frac{\|f(p+v) - f(p) - T(v)\|_W}{\|v\|_V} &= \frac{|(a+x)^2 + (b+y)^2 - (a^2 + b^2) - (2ax + 2by)|}{\sqrt{x^2 + y^2}} \\ &= \frac{x^2 + y^2}{\sqrt{x^2 + y^2}} = \sqrt{x^2 + y^2} \\ &\rightarrow 0 \end{aligned}$$

as $\|v\| \rightarrow 0$. Thus, for $p = a\mathbf{e}_1 + b\mathbf{e}_2$ we indeed have $(Df)_p = 2a \cdot \mathbf{e}_1^\vee + 2b \cdot \mathbf{e}_2^\vee$.

Remark 26.1.4. As usual, differentiability implies continuity.

Remark 26.1.5. Although $U \subseteq V$, it might be helpful to think of vectors from U and V as different types of objects (in particular, note that it's possible for $0_V \notin U$). The vectors in U are “inputs” on our space while the vectors coming from V are “small displacements”. For this reason, I deliberately try to use $p \in U$ and $v \in V$ when possible.

§26.2 The projection principle

Before proceeding I need to say something really important.

Theorem 26.2.1 (Projection principle)

Let U be an open subset of the vector space V . Let W be an n -dimensional real vector space with basis $w_1, \dots, w_n$. Then there is a bijection between continuous functions $f : U \rightarrow W$ and n -tuples of continuous $f_1, f_2, \dots, f_n : U \rightarrow \mathbb{R}$ by projection onto the i th basis element, i.e.

$$f(v) = f_1(v)w_1 + \dots + f_n(v)w_n.$$

Proof. Obvious. □

The theorem remains true if one replaces “continuous” by “differentiable”, “smooth”, “arbitrary”, or most other reasonable words. Translation:

To think about a function $f : U \rightarrow \mathbb{R}^n$, it suffices to think about each coordinate separately.

For this reason, we'll most often be interested in functions $f : U \rightarrow \mathbb{R}$. That's why the dual space $V^\vee$ is so important.

§26.3 Total and partial derivatives

Prototypical example for this section: If $f(x, y) = x^2 + y^2$, then $(Df) : (x, y) \mapsto 2x \cdot \mathbf{e}_1^\vee + 2y \cdot \mathbf{e}_2^\vee$, and $\frac{\partial f}{\partial x} = 2x$, $\frac{\partial f}{\partial y} = 2y$.

Let $U \subseteq V$ be open and let V have a basis $e_1, \dots, e_n$. Suppose $f : U \rightarrow \mathbb{R}$ is a function which is differentiable everywhere, meaning $(Df)_p \in V^\vee$ exists for every p . In that case, one can consider Df as *itself* a function:

$$\begin{aligned} Df : U &\rightarrow V^\vee \\ p &\mapsto (Df)_p. \end{aligned}$$

This is a little crazy: to every *point* in U we associate a *function* in $V^\vee$. We say Df is the **total derivative** of f , to reflect how much information we're dealing with. We say $(Df)_p$ is the total derivative at p .

Let's apply the projection principle now to Df . Since we picked a basis $e_1, \dots, e_n$ of V , there is a corresponding dual basis $e_1^\vee, e_2^\vee, \dots, e_n^\vee$. The Projection Principle tells us that Df can thus be thought of as just n functions, so we can write

$$Df = \psi_1 e_1^\vee + \dots + \psi_n e_n^\vee.$$

In fact, we can even describe what the ψ_i are.

Definition 26.3.1. The i^{th} **partial derivative** of $f : U \rightarrow \mathbb{R}$, denoted

$$\frac{\partial f}{\partial e_i} : U \rightarrow \mathbb{R}$$

is defined by

$$\frac{\partial f}{\partial e_i}(p) \stackrel{\text{def}}{=} \lim_{t \rightarrow 0} \frac{f(p + te_i) - f(p)}{t}.$$

You can think of it as “ f' along e_i ”.

Question 26.3.2. Check that if Df exists, then

$$(Df)_p(e_i) = \frac{\partial f}{\partial e_i}(p).$$

Remark 26.3.3. Of course you can write down a definition of $\frac{\partial f}{\partial v}$ for any v (rather than just the e_i).

From the above remarks, we can derive that

$$Df = \frac{\partial f}{\partial e_1} \cdot e_1^\vee + \dots + \frac{\partial f}{\partial e_n} \cdot e_n^\vee.$$

and so given a basis of V , we can think of Df as just the n partials.

Remark 26.3.4. Keep in mind that each $\frac{\partial f}{\partial e_i}$ is a function from U to the *reals*. That is to say,

$$(Df)_p = \underbrace{\frac{\partial f}{\partial e_1}(p)}_{\in \mathbb{R}} \cdot e_1^\vee + \cdots + \underbrace{\frac{\partial f}{\partial e_n}(p)}_{\in \mathbb{R}} \cdot e_n^\vee \in V^\vee.$$

Example 26.3.5 (Partial derivatives of $f(x, y) = x^2 + y^2$)

Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ by $(x, y) \mapsto x^2 + y^2$. Then in our new language,

$$Df : (x, y) \mapsto 2x \cdot \mathbf{e}_1^\vee + 2y \cdot \mathbf{e}_2^\vee.$$

Thus the partials are

$$\frac{\partial f}{\partial x} : (x, y) \mapsto 2x \in \mathbb{R} \quad \text{and} \quad \frac{\partial f}{\partial y} : (x, y) \mapsto 2y \in \mathbb{R}$$

With all that said, I haven't really said much about how to find the total derivative itself. For example, if I told you

$$f(x, y) = x \sin y + x^2 y^4$$

you might want to be able to compute Df without going through that horrible limit definition I told you about earlier.

Fortunately, it turns out you already know how to compute partial derivatives, because you had to take AP Calculus at some point in your life. It turns out for most reasonable functions, this is all you'll ever need.

Theorem 26.3.6 (Continuous partials implies differentiable)

Let $U \subseteq V$ be open and pick any basis $e_1, \dots, e_n$. Let $f : U \rightarrow \mathbb{R}$ and suppose that $\frac{\partial f}{\partial e_i}$ is defined for each i and moreover is *continuous*. Then f is differentiable and Df is given by

$$Df = \sum_{i=1}^n \frac{\partial f}{\partial e_i} \cdot e_i^\vee.$$

Proof. Not going to write out the details, but... given $v = t_1 e_1 + \cdots + t_n e_n$, the idea is to just walk from p to $p + t_1 e_1$, $p + t_1 e_1 + t_2 e_2$, ..., up to $p + t_1 e_1 + t_2 e_2 + \cdots + t_n e_n = p + v$, picking up the partial derivatives on the way. Do some calculation. $\square$

Remark 26.3.7. The continuous condition cannot be dropped. The function

$$f(x, y) = \begin{cases} \frac{xy}{x^2 + y^2} & (x, y) \neq (0, 0) \\ 0 & (x, y) = (0, 0). \end{cases}$$

is the classic counterexample – the total derivative Df does not exist at zero, even though both partials do.

Example 26.3.8 (Actually computing a total derivative)

Let $f(x, y) = x \sin y + x^2 y^4$. Then

$$\begin{aligned}\frac{\partial f}{\partial x}(x, y) &= \sin y + y^4 \cdot 2x \\ \frac{\partial f}{\partial y}(x, y) &= x \cos y + x^2 \cdot 4y^3.\end{aligned}$$

So Theorem 26.3.6 applies, and $Df = \frac{\partial f}{\partial x} \mathbf{e}_1^\vee + \frac{\partial f}{\partial y} \mathbf{e}_2^\vee$, which I won't bother to write out.

The example $f(x, y) = x^2 + y^2$ is the same thing. That being said, who cares about $x \sin y + x^2 y^4$ anyways?

§26.4 (Optional) A word on higher derivatives

Let $U \subseteq V$ be open, and take $f : U \rightarrow W$, so that $Df : U \rightarrow \text{Hom}(V, W)$.

Well, $\text{Hom}(V, W)$ can also be thought of as a normed vector space in its own right: it turns out that one can define an operator norm on it by setting

$$\|T\| \stackrel{\text{def}}{=} \sup \left\{ \frac{\|T(v)\|_W}{\|v\|_V} \mid v \neq 0_V \right\}.$$

So $\text{Hom}(V, W)$ can be thought of as a normed vector space as well. Thus it makes sense to write

$$D(Df) : U \rightarrow \text{Hom}(V, \text{Hom}(V, W))$$

which we abbreviate as $D^2 f$. Dropping all doubt and plunging on,

$$D^3 f : U \rightarrow \text{Hom}(V, \text{Hom}(V, \text{Hom}(V, W))).$$

I'm sorry. As consolation, we at least know that $\text{Hom}(V, W) \cong V^\vee \otimes W$ in a natural way, so we can at least condense this to

$$D^k f : V \rightarrow (V^\vee)^{\otimes k} \otimes W$$

rather than writing a bunch of Hom's.

Remark 26.4.1. If $k = 2$, $W = \mathbb{R}$, then $D^2 f(v) \in (V^\vee)^{\otimes 2}$, so it can be represented as an $n \times n$ matrix, which for some reason is called a **Hessian**.

The most important property of the second derivative is that

Theorem 26.4.2 (Symmetry of $D^2 f$)

Let $f : U \rightarrow W$ with $U \subseteq V$ open. If $(D^2 f)_p$ exists at some $p \in U$, then it is symmetric, meaning

$$(D^2 f)_p(v_1, v_2) = (D^2 f)_p(v_2, v_1).$$

I'll just quote this without proof (see e.g. [Pu02, §5, theorem 16]), because double derivatives make my head spin. An important corollary of this theorem:

Corollary 26.4.3 (Clairaut's theorem: mixed partials are symmetric)

Let $f : U \rightarrow \mathbb{R}$ with $U \subseteq V$ open be twice differentiable. Then for any point p such that the quantities are defined,

$$\frac{\partial}{\partial e_i} \frac{\partial}{\partial e_j} f(p) = \frac{\partial}{\partial e_j} \frac{\partial}{\partial e_i} f(p).$$

§26.5 Towards differential forms

This concludes the exposition of what the derivative really is: the key idea I want to communicate in this chapter is that Df should be thought of as a map from $U \rightarrow V^\vee$.

The next natural thing to do is talk about *integration*. The correct way to do this is through a so-called *differential form*: you'll finally know what all those stupid dx 's and dy 's really mean. (They weren't just there for decoration!)

§26.6 Problems to think about

Problem 26A* (Chain rule). Let $U_1 \xrightarrow{f} U_2 \xrightarrow{g} U_3$ be differentiable maps between open sets of normed vector spaces V_i , and let $h = g \circ f$. Prove the Chain Rule: for any point $p \in U_1$, we have

$$(Dh)_p = (Dg)_{f(p)} \circ (Df)_p.$$

Problem 26B. Let $U \subseteq V$ be open, and $f : U \rightarrow \mathbb{R}$ be differentiable k times. Show that $(D^k f)_p$ is symmetric in its k arguments, meaning for any $v_1, \dots, v_k \in V$ and any permutation σ on $\{1, \dots, k\}$ we have

$$(D^k f)_p(v_1, \dots, v_k) = (D^k f)_p(v_{\sigma(1)}, \dots, v_{\sigma(k)}).$$

27 Differential forms

In this chapter, all vector spaces are finite-dimensional real inner product spaces. We first start by (non-rigorously) drawing pictures of all the things that we will define in this chapter. Then we re-do everything again in its proper algebraic context.


§27.1 Pictures of differential forms

Before defining a differential form, we first draw some pictures. The key thing to keep in mind is

“The definition of a differential form is: something you can integrate.”
— Joe Harris

We’ll assume that all functions are **smooth**, i.e. infinitely differentiable.

Let $U \subseteq V$ be an open set of a vector space V . Suppose that we have a function $f : U \rightarrow \mathbb{R}$, i.e. we assign a value to every point of U .




Definition 27.1.1. A **0-form** f on U is just a smooth function $f : U \rightarrow \mathbb{R}$.

Thus, if we specify a finite set S of points in U we can “integrate” over S by just adding up the values of the points:

$$0 + \sqrt{2} + 3 + (-1) = 2 + \sqrt{2}.$$

So, a **0-form** f lets us integrate over 0-dimensional “cells”.

But this is quite boring, because as we know we like to integrate over things like curves, not single points. So, by analogy, we want a 1-form to let us integrate over 1-dimensional cells: i.e. over curves. What information would we need to do that? To answer this, let’s draw a picture of a curve c , which can be thought of as a function $c : [0, 1] \rightarrow U$.



We might think that we could get away with just specifying a number on every point of U (i.e. a 0-form f), and then somehow “add up” all the values of f along the curve. We’ll use this idea in a moment, but we can in fact do something more general. Notice how when we walk along a smooth curve, at every point p we also have some extra information: a *tangent vector* v . So, we can define a 1-form α as follows. A 0-form just took a point and gave a real number, but **a 1-form will take both a point *and* a tangent vector at that point, and spit out a real number.** So a 1-form α is a smooth function on pairs (p, v) , where v is a tangent vector at p , to $\mathbb{R}$. Hence

$$\alpha : U \times V \rightarrow \mathbb{R}.$$

Actually, for any point p , we will require that $\alpha(p, -)$ is a linear function in terms of the vectors: i.e. we want for example that $\alpha(p, 2v) = 2\alpha(p, v)$. So it is more customary to think of α as:

Definition 27.1.2. A **1-form** α is a smooth function

$$\alpha : U \rightarrow V^\vee.$$


Like with Df , we’ll use α_p instead of $\alpha(p)$. So, at every point p , α_p is some linear functional that eats tangent vectors at p , and spits out a real number. Thus, we think of α_p as an element of $V^\vee$;

$$\alpha_p \in V^\vee.$$

Next, we draw pictures of 2-forms. This should, for example, let us integrate over a blob (a so-called 2-cell) of the form

$$c : [0, 1] \times [0, 1] \rightarrow U$$

i.e. for example, a square in U . In the previous example with 1-forms, we looked at tangent vectors to the curve c . This time, at points we will look at *pairs* of tangent vectors in U : in the same sense that lots of tangent vectors approximate the entire curve, lots of tiny squares will approximate the big square in U .



So what should a 2-form β be? As before, it should start by taking a point $p \in U$, so β_p is now a linear functional: but this time, it should be a linear map on two vectors v and w . Here v and w are not tangent so much as their span cuts out a small parallelogram. So, the right thing to do is in fact consider

$$\beta_p \in V^\vee \wedge V^\vee.$$


That is, to use the wedge product to get a handle on the idea that v and w span a parallelogram. Another valid choice would have been $(V \wedge V)^\vee$; in fact, the two are isomorphic, but it will be more convenient to write it in the former.

§27.2 Pictures of exterior derivatives

Next question:

How can we build a 0-form from a 1-form?

Let f be a 0-form on U ; thus, we have a function $f : U \rightarrow \mathbb{R}$. Then in fact there is a very natural 1-form on U arising from f , appropriately called df . Namely, given a point p and a tangent vector v , the differential form $(df)_p$ returns the *change in f along v* . In other words, it's just the total derivative $(Df)_p(v)$.



Thus, df measures “the change in f ”.

Now, even if I haven’t defined integration yet, given a curve c from a point a to b , what do you think

$$\int_c df$$

should be equal to? Remember that df is the 1-form that measures “infinitesimal change in f ”. So if we add up all the change in f along a path from a to b , then the answer we get should just be

$$\int_c df = f(b) - f(a).$$

This is the first case of something we call Stokes’ theorem.

Generalizing, how should we get from a 1-form to a 2-form? At each point p , the 2-form β gives a β_p which takes in a “parallelogram” and returns a real number. Now suppose we have a 1-form α . Then along each of the edges of a parallelogram, with an appropriate sign convention the 1-form α gives us a real number. So, given a 1-form α , we define $d\alpha$ to be the 2-form that takes in a parallelogram spanned by v and w , and returns **the measure of α along the boundary**.

Now, what happens if you integrate df along the entire square c ? The right picture is that, if we think of each little square as making up the big square, then the adjacent boundaries cancel out, and all we are left is the main boundary. This is again just a case of the so-called Stokes’ theorem.





Image from [Na]

§27.3 Differential forms

Prototypical example for this section: Algebraically, something that looks like $f\mathbf{e}_1^\vee \wedge \mathbf{e}_2^\vee + \dots$, and geometrically, see the previous section.

Let's now get a handle on what dx means. Fix a real vector space V of dimension n , and let $\mathbf{e}_1, \dots, \mathbf{e}_n$ be a standard basis. Let U be an open set.

Definition 27.3.1. We define a **differential k -form** α on U to be a smooth (infinitely differentiable) map $\alpha : U \rightarrow \Lambda^k(V^\vee)$. (Here $\Lambda^k(V^\vee)$ is the wedge product.)

Like with Df , we'll use α_p instead of $\alpha(p)$.

Example 27.3.2 (k -forms for $k = 0, 1$)

- (a) A 0-form is just a function $U \rightarrow \mathbb{R}$.
- (b) A 1-form is a function $U \rightarrow V^\vee$. For example, the total derivative Df of a function $V \rightarrow \mathbb{R}$ is a 1-form.
- (c) Let $V = \mathbb{R}^3$ with standard basis $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$. Then a typical 2-form is given by

$$\alpha_p = f(p) \cdot \mathbf{e}_1^\vee \wedge \mathbf{e}_2^\vee + g(p) \cdot \mathbf{e}_1^\vee \wedge \mathbf{e}_3^\vee + h(p) \cdot \mathbf{e}_2^\vee \wedge \mathbf{e}_3^\vee \in \Lambda^2(V)$$

where $f, g, h : V \rightarrow \mathbb{R}$ are smooth functions.

Now, by the projection principle (Theorem 26.2.1) we only have to specify a function on each of $\binom{n}{k}$ basis elements of $\Lambda^k(V^\vee)$. So, take any basis $\{e_i\}$ of V , and take the usual basis for $\Lambda^k(V^\vee)$ of elements

$$e_{i_1}^\vee \wedge e_{i_2}^\vee \wedge \dots \wedge e_{i_k}^\vee.$$

Thus, a general k -form takes the shape

$$\alpha_p = \sum_{1 \leq i_1 < \dots < i_k \leq n} f_{i_1, \dots, i_k}(p) \cdot e_{i_1}^\vee \wedge e_{i_2}^\vee \wedge \dots \wedge e_{i_k}^\vee.$$

Since this is a huge nuisance to write, we will abbreviate this to just

$$\alpha = \sum_I f_I \cdot de_I$$

where we understand the sum runs over $I = (i_1, \dots, i_k)$, and de_I represents $e_{i_1}^\vee \wedge \dots \wedge e_{i_k}^\vee$.

Now that we have an element $\Lambda^k(V^\vee)$, what can it do? Well, first let me get the definition on the table, then tell you what it's doing.

Definition 27.3.3. For linear functions $\xi_1, \dots, \xi_k \in V^\vee$ and vectors $v_1, \dots, v_k \in V$, set

$$(\xi_1 \wedge \dots \wedge \xi_k)(v_1, \dots, v_k) \stackrel{\text{def}}{=} \det \begin{bmatrix} \xi_1(v_1) & \dots & \xi_1(v_k) \\ \vdots & \ddots & \vdots \\ \xi_k(v_1) & \dots & \xi_k(v_k) \end{bmatrix}.$$

You can check that this is well-defined under e.g. $v \wedge w = -w \wedge v$ and so on.

Example 27.3.4 (Evaluation of a differential form)

Set $V = \mathbb{R}^3$. Suppose that at some point p , the 2-form α returns

$$\alpha_p = 2\mathbf{e}_1^\vee \wedge \mathbf{e}_2^\vee + \mathbf{e}_1^\vee \wedge \mathbf{e}_3^\vee.$$


Let $v_1 = 3\mathbf{e}_1 + \mathbf{e}_2 + 4\mathbf{e}_3$ and $v_2 = 8\mathbf{e}_1 + 9\mathbf{e}_2 + 5\mathbf{e}_3$. Then

$$\alpha_p(v_1, v_2) = 2 \det \begin{bmatrix} 3 & 1 \\ 8 & 9 \end{bmatrix} + \det \begin{bmatrix} 3 & 4 \\ 8 & 5 \end{bmatrix} = 21.$$

What does this definition mean? One way to say it is that

If I walk to a point $p \in U$, a k -form α will take in k vectors $v_1, \dots, v_k$ and spit out a number, which is to be interpreted as a (signed) volume.

Picture:



In other words, at every point p , we get a function α_p . Then I can feed in k vectors to α_p and get a number, which I interpret as a signed volume of the parallelepiped spanned by the $\{v_i\}$'s in some way (e.g. the flux of a force field). That's why α_p as a “function” is contrived to lie in the wedge product: this ensures that the notion of “volume” makes sense, so that for example, the equality $\alpha_p(v_1, v_2) = -\alpha_p(v_2, v_1)$ holds.

This is what makes differential forms so fit for integration.

§27.4 Exterior derivatives

Prototypical example for this section: Possibly $dx_1 = \mathbf{e}_1^\vee$.

We now define the exterior derivative df that we gave pictures of at the beginning of the section. It turns out that the exterior derivative is easy to compute given explicit coordinates to work with.

First, given a function $f : U \rightarrow \mathbb{R}$, we define

$$df \stackrel{\text{def}}{=} Df = \sum_i f_i \mathbf{e}_i^\vee$$

In particular, suppose $V = \mathbb{R}^n$ and $f(x_1, \dots, x_n) = x_1$ (i.e. $f = \mathbf{e}_1^\vee$). Then:

Question 27.4.1. Show that for any $p \in U$,

$$(d(\mathbf{e}_1^\vee))_p = \mathbf{e}_1^\vee.$$

Abuse of Notation 27.4.2. Unfortunately, someone somewhere decided it would be a good idea to use “ x_1 ” to denote $\mathbf{e}_1^\vee$ (because *obviously*¹ x_1 means “the function that takes $(x_1, \dots, x_n) \in \mathbb{R}^n$ to x_1 ”) and then decided that

$$dx_1 \stackrel{\text{def}}{=} \mathbf{e}_1^\vee.$$

This notation is so entrenched that I have no choice but to grudgingly accept it. Note that it’s not even right, since technically it’s $(dx_1)_p = \mathbf{e}_1^\vee$; dx_1 is a 1-form.

Remark 27.4.3. This is the reason why we use the notation $\frac{df}{dx}$ in calculus now: given, say, $f : \mathbb{R} \rightarrow \mathbb{R}$ by $f(x) = x^2$, it is indeed true that

$$df = 2x \cdot \mathbf{e}_1^\vee = 2x \cdot dx$$

and so by abuse of notation we write $df/dx = 2x$.

More generally, we can define the **exterior derivative** in terms of our basis $e_1, \dots, e_n$ as follows: if $\alpha = \sum_I f_I de_I$ then we set

$$d\alpha \stackrel{\text{def}}{=} \sum_I df_I \wedge de_I = \sum_I \sum_j \frac{\partial f_I}{\partial e_j} de_j \wedge de_I.$$

This doesn’t depend on the choice of basis.

Example 27.4.4 (Computing some exterior derivatives)

Let $V = \mathbb{R}^3$ with standard basis $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$. Let $f(x, y, z) = x^4 + y^3 + 2xz$. Then we compute

$$df = Df = (4x^3 + 2z) dx + 3y^2 dy + 2x dz.$$

Next, we can evaluate $d(df)$ as prescribed: it is

$$\begin{aligned} d^2 f &= (12x^2 dx + 2dz) \wedge dx + (6y dy) \wedge dy + 2(dx \wedge dz) \\ &= 12x^2(dx \wedge dx) + 4(dz \wedge dx) + 6y(dy \wedge dy) + 4(dx \wedge dz) \\ &= 4(dz \wedge dx) + 4(dx \wedge dz) \\ &= 0. \end{aligned}$$

So surprisingly, $d^2 f$ is the zero map. Here, we have exploited Abuse of Notation 27.4.2 for the first time, in writing dx, dy, dz .

And in fact, this is always true in general:

Theorem 27.4.5 (Exterior derivative vanishes)

Let α be any k -form. Then $d^2(\alpha) = 0$. Even more succinctly,

$$d^2 = 0.$$

The proof is left as Problem 27B.

¹Sarcasm.

Exercise 27.4.6. Compare the statement $d^2 = 0$ to the geometric picture of a 2-form given at the beginning of this chapter. Why does this intuitively make sense?

Here are some other properties of d :

- As we just saw, $d^2 = 0$.
- It is bilinear; $d(c\alpha + \beta) = c \cdot d\alpha + d\beta$.
- For a k -form α and ℓ -form β , one can show that

$$d(\alpha \wedge \beta) = d\alpha \wedge \beta + (-1)^k(\alpha \wedge d\beta).$$

In fact, one can show that df as defined in terms of bases above is the *unique* map sending k -forms to $(k+1)$ -forms. Actually, one way to define df is to take as axioms the bulleted properties above and then declare d to be the unique solution to this functional equation. In any case, this tells us that our definition of d does not depend on the basis chosen.

Recall that df measures the change in boundary. In that sense, $d^2 = 0$ is saying something like “the boundary of the boundary is empty”. We’ll make this precise when we see Stokes’ theorem in the next chapter.

§27.5 Closed and exact forms

Let α be a k -form.

Definition 27.5.1. We say α is **closed** if $d\alpha = 0$.

Definition 27.5.2. We say α is **exact** if for some $(k-1)$ -form β , $d\beta = \alpha$. If $k = 0$, α is exact only when $\alpha = 0$.

Question 27.5.3. Show that exact forms are closed.

A natural question arises: are there closed forms which are not exact? Surprisingly, the answer to this question is tied to topology. Here is one important example.

Example 27.5.4 (The angle form)

Let $U = \mathbb{R}^2 \setminus \{0\}$, and let $\theta(p)$ be the angle formed by the x -axis and the line from the origin to p .

The 1-form $\alpha : U \rightarrow (\mathbb{R}^2)^\vee$ defined by

$$\alpha = \frac{-y \, dx + x \, dy}{x^2 + y^2}$$

is called the **angle form**: given $p \in U$ it measures the change in angle $\theta(p)$ along a tangent vector. So intuitively, “ $\alpha = d\theta$ ”. Indeed, one can check directly that the angle form is closed.

However, α is not exact: there is no global smooth function $\theta : U \rightarrow \mathbb{R}$ having α as a derivative. This reflects the fact that one can actually perform a full 2π rotation around the origin, i.e. θ only makes sense mod 2π . Thus existence of the angle form α reflects the possibility of “winding” around the origin.

So the key idea is that the failure of a closed form to be exact corresponds quite well with “holes” in the space: the same information that homotopy and homology groups are trying to capture. To draw another analogy, in complex analysis Cauchy-Goursat only works when U is simply connected. The “hole” in U is being detected by the existence of a form α . The so-called de Rham cohomology will make this relation explicit.

§27.6 Problems to think about

Problem 27A. Show directly that the angle form

$$\alpha = \frac{-y \, dx + x \, dy}{x^2 + y^2}$$

is closed.

Problem 27B. Establish Theorem 27.4.5, which states that $d^2 = 0$.

DRAFT (Evan Chen)
Updated February 15, 2017

28 Integrating differential forms

We now show how to integrate differential forms over cells, and state Stokes' theorem in this context. In this chapter, all vector spaces are finite-dimensional and real.

§28.1 Motivation: line integrals

Given a function $g : [a, b] \rightarrow \mathbb{R}$, we know by the fundamental theorem of calculus that


$$\int_{[a,b]} g(t) dt = f(b) - f(a)$$

where f is a function such that $g = df/dt$. Equivalently, for $f : [a, b] \rightarrow \mathbb{R}$,

$$\int_{[a,b]} g dt = \int_{[a,b]} df = f(b) - f(a)$$

where df is the exterior derivative we defined earlier.

Cool, so we can integrate over $[a, b]$. Now suppose more generally, we have U an open subset of our real vector space V and a 1-form $\alpha : U \rightarrow V^\vee$. We consider a **parametrized curve**, which is a smooth function $c : [a, b] \rightarrow U$. Picture:



We want to define an $\int_c \alpha$ such that:

The integral $\int_c \alpha$ should add up all the α along the curve c .

Our differential form α first takes in a point p to get $\alpha_p \in V^\vee$. Then, it eats a tangent vector $v \in V$ to the curve c to finally give a real number $\alpha_p(v) \in \mathbb{R}$. We would like to “add all these numbers up”, using only the notion of an integral over $[a, b]$.

Exercise 28.1.1. Try to guess what the definition of the integral should be. (By type-checking, there's only one reasonable answer.)

So, the definition we give is

$$\int_c \alpha \stackrel{\text{def}}{=} \int_{[a,b]} \alpha_{c(t)}(c'(t)) dt.$$

Here, $c'(t)$ is shorthand for $(Dc)_{c(t)}(1)$. It represents the *tangent vector* to the curve c at the point $p = c(t)$, at time t . (Here we are taking advantage of the fact that $[a, b]$ is one-dimensional.)

Now that definition was a pain to write, so we will define a differential 1-form $c^*\alpha$ on $[a, b]$ to swallow that entire thing: specifically, in this case we define $c^*\alpha$ to be

$$(c^*\alpha)_t(\varepsilon) = (Dc)_{c(t)}(\varepsilon)$$

(here ε is some displacement in time). Thus, we can more succinctly write


$$\int_c \alpha \stackrel{\text{def}}{=} \int_{[a,b]} c^*\alpha.$$

This is a special case of a *pullback*: roughly, if $\phi : U \rightarrow U'$ (where $U \subseteq V$, $U' \subseteq V'$), we can change any differential k -form α on U' to a k -form on U . In particular, if $U = [a, b]$,¹ we can resort to our old definition of an integral. Let's now do this in full generality.

§28.2 Pullbacks

Let V and V' be finite dimensional real vector spaces (possibly different dimensions) and suppose U and U' are open subsets of each; next, consider a k -form α on U' .

Given a map $\phi : U \rightarrow U'$ we now want to define a pullback in much the same way as before. Picture:



Well, there's a total of about one thing we can do. Specifically: α accepts a point in U' and k tangent vectors in V' , and returns a real number. We want $\phi^*\alpha$ to accept a point in $p \in U$ and k tangent vectors $v_1, \dots, v_k$ in V , and feed the corresponding information to α .

¹ OK, so $[a, b]$ isn't actually open, sorry. I ought to write $(a - \varepsilon, b + \varepsilon)$, or something.

Clearly we give the point $q = \phi(p)$. As for the tangent vectors, since we are interested in volume, we take the derivative of ϕ at p , $(D\phi)_p$, which will scale each of our vectors v_i into some vector in the target V' . To cut a long story short:

Definition 28.2.1. Given $\phi : U \rightarrow U'$ and α a k -form, we define the **pullback**

$$(\phi^*\alpha)_p(v_1, \dots, v_k) \stackrel{\text{def}}{=} \alpha_{\phi(p)}((D\phi)_p(v_1), \dots, (D\phi)_p(v_k)).$$

There is a more concrete way to define the pullback using bases. Suppose $w_1, \dots, w_n$ is a basis of V' and $e_1, \dots, e_m$ is a basis of V . Thus, by the projection principle (Theorem 26.2.1) the map $\phi : V \rightarrow V'$ can be thought of as

$$\phi(v) = \phi_1(v)w_1 + \dots + \phi_n(v)w_n$$

where each ϕ_i takes in a $v \in V$ and returns a real number. We know also that α can be written concretely as

$$\alpha = \sum_{J \subseteq \{1, \dots, n\}} f_J w_J.$$

Then, we define

$$\phi^*\alpha = \sum_{I \subseteq \{1, \dots, m\}} (f_I \circ \phi)(D\phi_{i_1} \wedge \dots \wedge D\phi_{i_k}).$$

A diligent reader can check these definitions are equivalent.

Example 28.2.2 (Computation of a pullback)

Let $V = \mathbb{R}^2$ with basis $\mathbf{e}_1$ and $\mathbf{e}_2$, and suppose $\phi : V \rightarrow V'$ is given by sending

$$\phi(a\mathbf{e}_1 + b\mathbf{e}_2) = (a^2 + b^2)w_1 + \log(a^2 + 1)w_2 + b^3w_3$$

where w_1, w_2, w_3 is a basis for V' . Consider the form $\alpha_q = f(q)w_1 \wedge w_3$, where $f : V' \rightarrow \mathbb{R}$. Then

$$(\phi^*\alpha)_p = f(\phi(p)) \cdot (2a\mathbf{e}_1^\vee + 2b\mathbf{e}_2^\vee) \wedge (3b^2\mathbf{e}_2^\vee) = f(\phi(p)) \cdot 6ab^2 \cdot \mathbf{e}_1^\vee \wedge \mathbf{e}_2^\vee.$$

It turns out that the pullback basically behaves nicely as possible, e.g.

- $\phi^*(c\alpha + \beta) = c\phi^*\alpha + \phi^*\beta$ (bilinearity)
- $\phi^*(\alpha \wedge \beta) = (\phi^*\alpha) \wedge (\phi^*\beta)$
- $\phi_1^*(\phi_2^*(\alpha)) = (\phi_1 \circ \phi_2)^*(\alpha)$ (naturality)

but I won't take the time to check these here (one can verify them all by expanding with a basis).

§28.3 Cells

Prototypical example for this section: A disk in $\mathbb{R}^2$ can be thought of as the cell $[0, R] \times [0, 2\pi] \rightarrow \mathbb{R}^2$ by $(r, \theta) \mapsto (r \cos \theta)\mathbf{e}_1 + (r \sin \theta)\mathbf{e}_2$.

Now that we have the notion of a pullback, we can define the notion of an integral for more general spaces. Specifically, to generalize the notion of integrals we had before:

Definition 28.3.1. A **k -cell** is a smooth function $c : [a_1, b_1] \times [a_2, b_2] \times \dots \times [a_k, b_k] \rightarrow V$.

Example 28.3.2 (Examples of cells)

Let $V = \mathbb{R}^2$ for convenience.

- (a) A 0-cell consists of a single point.
- (b) As we saw, a 1-cell is an arbitrary curve.
- (c) A 2-cell corresponds to a 2-dimensional surface. For example, the map $c : [0, R] \times [0, 2\pi]$ by

$$c : (r, \theta) \mapsto (r \cos \theta, r \sin \theta)$$

can be thought of as a disk of radius R .

Then, to define an integral

$$\int_c \alpha$$

for a differential k -form α and a k -cell $c : [0, 1]^k$, we simply take the pullback

$$\int_{[0,1]^k} c^* \alpha$$

Since $c^* \alpha$ is a k -form on the k -dimensional unit box, it can be written as $f(x_1, \dots, x_n) dx_1 \wedge \dots \wedge dx_n$, so the above integral can be written as

$$\int_0^1 \dots \int_0^1 f(x_1, \dots, x_n) dx_1 \wedge \dots \wedge dx_n$$

Example 28.3.3 (Area of a circle)

Consider $V = \mathbb{R}^2$ and let $c : (r, \theta) \mapsto (r \cos \theta) \mathbf{e}_1 + (r \sin \theta) \mathbf{e}_2$ on $[0, R] \times [0, 2\pi]$ as before. Take the 2-form α which gives $\alpha_p = \mathbf{e}_1^\vee \wedge \mathbf{e}_2^\vee$ at every point p . Then

$$\begin{aligned} c^* \alpha &= (\cos \theta dr - r \sin \theta d\theta) \wedge (\sin \theta dr + r \cos \theta d\theta) \\ &= r(\cos^2 \theta + \sin^2 \theta)(dr \wedge d\theta) \\ &= r dr \wedge d\theta \end{aligned}$$

Thus,

$$\int_c \alpha = \int_0^R \int_0^{2\pi} r dr \wedge d\theta = \pi R^2$$

which is the area of a circle.

Here's some geometric intuition for what's happening. Given a k -cell in V , a differential k -form α accepts a point p and some tangent vectors $v_1, \dots, v_k$ and spits out a number $\alpha_p(v_1, \dots, v_k)$, which as before we view as a signed hypervolume. Then the integral *adds up all these infinitesimals across the entire cell*. In particular, if $V = \mathbb{R}^k$ and we take the form $\alpha : p \mapsto \mathbf{e}_1^\vee \wedge \dots \wedge \mathbf{e}_k^\vee$, then what these α 's give is the k th hypervolume of the cell. For this reason, this α is called the **volume form** on $\mathbb{R}^k$.

You'll notice I'm starting to play loose with the term "cell": while the cell $c : [0, R] \times [0, 2\pi] \rightarrow \mathbb{R}^2$ is supposed to be a function I have been telling you to think of it as a unit disk (i.e. in terms of its image). In the same vein, a curve $[0, 1] \rightarrow V$ should be thought of as a curve in space, rather than a function on time.

This error turns out to be benign. Let α be a k -form on U and $c : [a_1, b_1] \times \cdots \times [a_k, b_k] \rightarrow U$ a k -cell. Suppose $\phi : [a'_1, b'_1] \times \cdots \times [a'_k, b'_k] \rightarrow [a_1, b_1] \times \cdots \times [a_k, b_k]$; it is a **reparametrization** if ϕ is bijective and $(D\phi)_p$ is always invertible (think “change of variables”); thus

$$c \circ \phi : [a'_1, b'_1] \times \cdots \times [a'_k, b'_k] \rightarrow U$$

is a k -cell as well. Then it is said to **preserve orientation** if $\det(D\phi)_p > 0$ for all p and **reverse orientation** if $\det(D\phi)_p < 0$ for all p .

Exercise 28.3.4. Why is it that exactly one of these cases must occur?

Theorem 28.3.5 (Changing variables doesn't affect integrals)

Let c be a k -cell, α a k -form, and ϕ a reparametrization. Then

$$\int_{c \circ \phi} \alpha = \begin{cases} \int_c \alpha & \phi \text{ preserves orientation} \\ -\int_c \alpha & \phi \text{ reverses orientation.} \end{cases}$$

Proof. Use naturality of the pullback to reduce it to the corresponding theorem in normal calculus. $\square$

So for example, if we had parametrized the unit circle as $[0, 1] \times [0, 1] \rightarrow \mathbb{R}^2$ by $(r, t) \mapsto R \cos(2\pi t) \mathbf{e}_1 + R \sin(2\pi t) \mathbf{e}_2$, we would have arrived at the same result. So we really can think of a k -cell just in terms of the points it specifies.

§28.4 Boundaries

Prototypical example for this section: The boundary of $[a, b]$ is $\{b\} - \{a\}$. The boundary of a square goes around its edge counterclockwise.

First, I introduce a technical term that lets us consider multiple cells at once.

Definition 28.4.1. A **k -chain** U is a formal linear combination of k -cells over U , i.e. a sum of the form

$$c = a_1 c_1 + \cdots + a_m c_m$$

where each $a_i \in \mathbb{R}$ and c_i is a k -cell. We define $\int_c \alpha = \sum_i a_i \int c_i$.

In particular, a 0-cell consists of several points, each with a given weight.

Now, how do we define the boundary? For a 1-cell $[a, b] \rightarrow U$, as I hinted earlier we want the answer to be the 0-chain $\{c(b)\} - \{c(a)\}$. Here's how we do it in general.

Definition 28.4.2. Suppose $c : [0, 1]^k \rightarrow U$ is a k -cell. Then the **boundary** of c , denoted $\partial c : [0, 1]^{k-1} \rightarrow U$, is the $(k-1)$ -chain defined as follows. For each $i = 1, \dots, k$ define

$$\begin{aligned} c_i^{\text{start}}(t_1, \dots, t_{k-1}) &= (t_1, \dots, t_{i-1}, 0, t_i, \dots, t_k) \\ c_i^{\text{stop}}(t_1, \dots, t_{k-1}) &= (t_1, \dots, t_{i-1}, 1, t_i, \dots, t_k). \end{aligned}$$

Then


$$\partial c \stackrel{\text{def}}{=} \sum_{i=1}^k (-1)^{i+1} (c_i^{\text{stop}} - c_i^{\text{start}}).$$

Finally, the boundary of a chain is the sum of the boundaries of each cell (with the appropriate weights). That is, $\partial(\sum a_i c_i) = \sum a_i \partial c_i$.

Question 28.4.3. Satisfy yourself that one can extend this definition to a k -cell c defined on $c : [a_1, b_1] \times \cdots \times [a_k, b_k] \rightarrow V$ (rather than from $[0, 1]^k \rightarrow V$).

Example 28.4.4 (Examples of boundaries)

Consider the 2-cell $c : [0, 1]^2 \rightarrow \mathbb{R}^2$ shown below.



Here p_1, p_2, p_3, p_4 are the images of $(0, 0), (0, 1), (1, 0), (1, 1)$, respectively. Then we can think of ∂c as

$$\partial c = [p_1, p_2] + [p_2, p_3] + [p_3, p_4] + [p_4, p_1]$$

where each “interval” represents the 1-cell shown by the reddish arrows on the right. We can take the boundary of this as well, and obtain an empty chain as


$$\partial(\partial c) = \sum_{i=1}^4 \partial([p_i, p_{i+1}]) = \sum_{i=1}^4 \{p_{i+1}\} - \{p_i\} = 0.$$

Example 28.4.5 (Boundary of a unit disk)

Consider the unit disk given by

$$c : [0, 1] \times [0, 2\pi] \rightarrow \mathbb{R}^2 \quad \text{by} \quad (r, \theta) \mapsto r \cos(2\pi t) \mathbf{e}_1 + r \sin(2\pi t) \mathbf{e}_2.$$

The four parts of the boundary are shown in the picture below:



Note that two of the arrows more or less cancel each other out when they are integrated. Moreover, we interestingly have a *degenerate* 1-cell at the center of the circle; it is a constant function $[0, 1] \rightarrow \mathbb{R}^2$ which always gives the origin.

Obligatory theorem, analogous to $d^2 = 0$ and left as a problem.

Theorem 28.4.6 (The boundary of the boundary is empty)

$\partial^2 = 0$, in the sense that for any k -chain c we have $\partial^2(c) = 0$.

§28.5 Stokes' theorem

Prototypical example for this section: $\int_{[a,b]} dg = g(b) - g(a)$.

We now have all the ingredients to state Stokes' theorem for cells.

Theorem 28.5.1 (Stokes' theorem for cells)

Take $U \subseteq V$ as usual, let $c : [0, 1]^k \rightarrow U$ be a k -cell and let $\alpha : U \rightarrow \Lambda^k(V^\vee)$ be a k -form. Then

$$\int_c d\alpha = \int_{\partial c} \alpha.$$

In particular, if $d\alpha = 0$ then the left-hand side vanishes.

For example, if c is the interval $[a, b]$ then $\partial c = \{b\} - \{a\}$, and thus we obtain the fundamental theorem of calculus.

§28.6 Problems to think about

Problem 28A[†] (Green's theorem). Let $f, g : \mathbb{R}^2 \rightarrow \mathbb{R}$ be smooth functions. Prove that

$$\int_c \left(\frac{\partial g}{\partial x} - \frac{\partial f}{\partial y} \right) dx \wedge dy = \int_{\partial c} (f dx + g dy).$$

Problem 28B. Show that $\partial^2 = 0$.

Problem 28C (Pullback and d commute). Let U and U' be open sets of vector spaces V and V' and let $\phi : U \rightarrow U'$ be a smooth map between them. Prove that for any differential form α on U' we have

$$\phi^*(d\alpha) = d(\phi^*\alpha).$$

Problem 28D (Arc length isn't a form). Show that there does *not* exist a 1-form α on $\mathbb{R}^2$ such that for a curve $c : [0, 1] \rightarrow \mathbb{R}^2$, the integral $\int_c \alpha$ gives the arc length of c .



Problem 28E. An **exact** k -form α is one satisfying $\alpha = d\beta$ for some β . Prove that

$$\int_{C_1} \alpha = \int_{C_2} \alpha$$

where C_1 and C_2 are any concentric circles in the plane and α is some exact 1-form.

29 A bit of manifolds

Last chapter, we stated Stokes' theorem for cells. It turns out there is a much larger class of spaces, the so-called *smooth manifolds*, for which this makes sense.

Unfortunately, the definition of a smooth manifold is *complete garbage*, and so by the time I am done defining differential forms and orientations, I will be too lazy to actually define what the integral on it is, and just wave my hands and state Stokes' theorem.

§29.1 Topological manifolds

Prototypical example for this section: S^2 : “the Earth looks flat”.

Long ago, people thought the Earth was flat, i.e. homeomorphic to a plane, and in particular they thought that $\pi_2(\text{Earth}) = 0$. But in fact, as most of us know, the Earth is actually a sphere, which is not contractible and in particular $\pi_2(\text{Earth}) \cong \mathbb{Z}$. This observation underlies the definition of a manifold:

An n -manifold is a space which locally looks like $\mathbb{R}^n$.

Actually there are two ways to think about a topological manifold M :

- “Locally”: at every point $p \in M$, some neighborhood of p looks like an open set of $\mathbb{R}^n$. For example, to someone standing on the surface of the Earth, the Earth looks much like $\mathbb{R}^2$.
- “Globally”: there exists an open cover of M by open sets $\{U_i\}_i$ (possibly infinite) such that each U_i is homeomorphic to some open subset of $\mathbb{R}^n$. For example, from outer space, the Earth can be covered by two hemispherical pancakes.

Question 29.1.1. Check that these are equivalent.

While the first one is the best motivation for examples, the second one is easier to use formally.

Definition 29.1.2. A **topological n -manifold** M is a Hausdorff space with an open cover $\{U_i\}$ of sets homeomorphic to subsets of $\mathbb{R}^n$, say by homeomorphisms

$$\phi_i : U_i \xrightarrow{\cong} E_i \subseteq \mathbb{R}^n$$


where each E_i is an open subset of $\mathbb{R}^n$. Each $\phi_i : U_i \rightarrow E_i$ is called a **chart**, and together they form a so-called **atlas**.

Remark 29.1.3. Here “ E ” stands for “Euclidean”. I think this notation is not standard; usually people just write $\phi_i(U_i)$ instead.

Remark 29.1.4. This definition is nice because it doesn't depend on embeddings: a manifold is an *intrinsic* space M , rather than a subset of $\mathbb{R}^N$ for some N . Analogy: an abstract group G is an intrinsic object rather than a subgroup of S_n .

Example 29.1.5 (An atlas on S^1)

Here is a picture of an atlas for S^1 , with two open sets.



Question 29.1.6. Where do you think the words “chart” and “atlas” come from?

Example 29.1.7 (Some examples of topological manifolds)

- (a) As discussed at length, the sphere S^2 is a 2-manifold: every point in the sphere has a small neighborhood that looks like D^2 . One can cover the Earth with just two hemispheres, and each hemisphere is homeomorphic to a disk.
- (b) The circle S^1 is a 1-manifold; every point has a neighborhood that looks like an open interval.
- (c) The torus, Klein bottle, $\mathbb{RP}^2$, $\mathbb{CP}^2$ are all 2-manifolds.
- (d) $\mathbb{R}^n$ is trivially a manifold, as are its open sets.

All these spaces are compact except $\mathbb{R}^n$.

A non-example of a manifold is D^n , because it has a *boundary*; points on the boundary do not have locally Euclidean neighborhoods.

§29.2 Smooth manifolds

Prototypical example for this section: All the topological manifolds.


Let M be a topological n -manifold with atlas $\{U_i \xrightarrow{\phi_i} E_i\}$.

Definition 29.2.1. For any i, j such that $U_i \cap U_j \neq \emptyset$, the map

$$\phi_{ij} : E_i \cap \phi_i^{\text{pre}}(U_i \cap U_j) \xrightarrow{\phi_i^{-1} \circ \phi_j} E_j \cap \phi_j^{\text{pre}}(U_i \cap U_j)$$

is called a **transition map**.

The intersection with the pre-image $\phi_i^{\text{pre}}(U_i \cap U_j)$ and the image $\phi_j^{\text{pre}}(U_i \cap U_j)$ is a notational annoyance to make the map well-defined and a homeomorphism. The transition map is just the natural way to go from $E_i \rightarrow E_j$, restricted to overlaps. Picture below, where the intersections are just the green portions of each E_1 and E_2 :



We want to add enough structure so that we can use differential forms.

Definition 29.2.2. We say M is a **smooth manifold** if all its transition maps are smooth.

This definition makes sense, because we know what it means for a map between two open sets of $\mathbb{R}^n$ to be differentiable.

With smooth manifolds we can try to port over definitions that we built for $\mathbb{R}^n$ onto our manifolds. So in general, all definitions involving smooth manifolds will reduce to something on each of the coordinate charts, with a compatibility condition.

Example 29.2.3 (Examples of manifold definitions)

- (a) Let M be a smooth manifold. A continuous function $f : M \rightarrow \mathbb{R}$ is called **smooth** if the composition

$$E_i \xrightarrow{\phi_i^{-1}} U_i \hookrightarrow M \xrightarrow{f} \mathbb{R}$$

is smooth as a function $E_i \rightarrow \mathbb{R}$.

- (b) Let M and N be smooth with atlases $\{U_i^M \xrightarrow{\phi_i} E_i^M\}_i$ and $\{U_j^N \xrightarrow{\phi_j} E_j^N\}_j$. A map $f : M \rightarrow N$ is **smooth** if for every i and j , the composed map

$$E_i \xrightarrow{\phi_i^{-1}} U_i \hookrightarrow M \xrightarrow{f} N \twoheadrightarrow U_j \xrightarrow{\phi_j} E_j$$

is smooth, as a function $E_i \rightarrow E_j$.

§29.3 Differential forms on manifolds

We already know what a differential form is on an open set $U \subseteq \mathbb{R}^n$. So, we naturally try to port over the definition of differentiable form on each subset, plus a compatibility condition.

Let M be a smooth manifold with atlas $\{U_i \xrightarrow{\phi_i} E_i\}_i$.

Definition 29.3.1. A **differential k -form** α on a smooth manifold M is a collection $\{\alpha_i\}_i$ of differential k -forms on each E_i , such that for any j and i we have that

$$\alpha_j = \phi_{ij}^*(\alpha_i).$$

Here the pullback. In English: we specify a differential form on each chart, which is compatible under pullbacks of the transition maps.

§29.4 Orientations

Prototypical example for this section: Left versus right, clockwise vs. counterclockwise.

This still isn't enough to integrate on manifolds. We need one more definition: that of an orientation.

The main issue is the observation from standard calculus that

$$\int_a^b f(x) dx = - \int_b^a f(x) dx.$$

Consider then a space M which is homeomorphic to an interval. If we have a 1-form α , how do we integrate it over M ? Since M is just a topological space (rather than a subset of $\mathbb{R}$), there is no default “left” or “right” that we can pick. As another example, if $M = S^1$ is a circle, there is no default “clockwise” or “counterclockwise” unless we decide to embed M into $\mathbb{R}^2$.

To work around this we have to actually have to make additional assumptions about our manifold.


Definition 29.4.1. A smooth n -manifold is **orientable** if there exists a differential n -form ω on M such that for every $p \in M$,

$$\omega_p \neq 0.$$

Recall here that ω_p is an element of $\Lambda^n(V^\vee)$. In that case we say ω is a **volume form** of M .

How do we picture this definition? If we recall that an differential form is supposed to take tangent vectors of M and return real numbers. To this end, we can think of each point $p \in M$ as having a **tangent plane** $T_p(M)$ which is n -dimensional. Now since the volume form ω is n -dimensional, it takes an entire basis of the $T_p(M)$ and gives a real number. So a manifold is orientable if there exists a consistent choice of sign for the basis of tangent vectors at every point of the manifold.

For “embedded manifolds”, this just amounts to being able to pick a nonzero field of normal vectors to each point $p \in M$. For example, S^1 is orientable in this way.



Similarly, one can orient a sphere S^2 by having a field of vectors pointing away (or towards) the center. This is all non-rigorous, because I haven't defined the tangent plane $T_p(M)$; since M is in general an intrinsic object one has to be quite roundabout to define $T_p(M)$. In any event, the point is that guesses about the orientability of spaces are likely to be correct.

Example 29.4.2 (Orientable surfaces)

- (a) Spheres S^n , planes, and the torus $S^1 \times S^1$ are orientable.
- (b) The Möbius strip and Klein bottle are *not* orientable: they are “one-sided”.
- (c) $\mathbb{CP}^n$ is orientable for any n .
- (d) $\mathbb{RP}^n$ is orientable only for odd n .

§29.5 Stokes' theorem for manifolds

Stokes' theorem in the general case is based on the idea of a **manifold with boundary** M , which I won't define, other than to say its boundary ∂M is an $n - 1$ dimensional manifold, and that it is oriented if M is oriented. An example is $M = D^2$, which has boundary $\partial M = S^1$.

Next,

Definition 29.5.1. The **support** of a differential form α on M is the closure of the set

$$\{p \in M \mid \alpha_p \neq 0\}.$$

If this support is compact as a topological space, we say α is **compactly supported**.

Remark 29.5.2. For example, volume forms are supported on all of M .

Now, one can define integration on oriented manifolds, but I won't define this because the definition is truly awful. Then Stokes' theorem says

Theorem 29.5.3 (Stokes' theorem for manifolds)

Let M be an smooth oriented n -manifold with boundary and let α be a compactly supported n -form. Then

$$\int_M d\alpha = \int_{\partial M} \alpha.$$

All the omitted details are developed in full in [Sj05].

§29.6 Problems to think about

Problem 29A. Show that a differential 0-form on a smooth manifold M is the same thing as a smooth function $M \rightarrow \mathbb{R}$.

IX

Algebraic Topology II: Homology

30 Singular homology	291
30.1 Simplices and boundaries	291
30.2 The singular homology groups	293
30.3 The homology functor and chain complexes	296
30.4 More examples of chain complexes	300
30.5 Problems to think about	301
31 The long exact sequence	302
31.1 Short exact sequences and four examples	302
31.2 The long exact sequence of homology groups	304
31.3 The Mayer-Vietoris sequence	306
31.4 Problems to think about	311
32 Excision and relative homology	312
32.1 The long exact sequences	312
32.2 The category of pairs	313
32.3 Excision	314
32.4 Some applications	315
32.5 Invariance of dimension	316
32.6 Problems to think about	317
33 Bonus: Cellular homology	318
33.1 Degrees	318
33.2 Cellular chain complex	319
33.3 The cellular boundary formula	322
33.4 Problems to think about	324
34 Singular cohomology	325
34.1 Cochain complexes	325
34.2 Cohomology of spaces	326
34.3 Cohomology of spaces is functorial	327
34.4 Universal coefficient theorem	328
34.5 Example computation of cohomology groups	329
34.6 Relative cohomology groups	330
34.7 Problems to think about	331
35 Application of cohomology	332
35.1 Poincaré duality	332
35.2 de Rham cohomology	332
35.3 Graded rings	333
35.4 Cup products	335
35.5 Relative cohomology pseudo-rings	337
35.6 Wedge sums	337
35.7 Künneth formula	338
35.8 Problems to think about	340

30 Singular homology

Now that we've defined $\pi_1(X)$, we turn our attention to a second way of capturing the same idea, $H_1(X)$. We'll then define $H_n(X)$ for $n \geq 2$. The good thing about the H_n groups is that, unlike the π_n groups, they are much easier to compute in practice. The downside is that their definition will require quite a bit of setup, and the "algebraic" part of "algebraic topology" will become a lot more technical.

§30.1 Simplices and boundaries

Prototypical example for this section: $\partial[v_0, v_1, v_2] = [v_0, v_1] - [v_0, v_2] + [v_1, v_2]$.

First things first:

Definition 30.1.1. The **standard n -simplex**, denoted Δ^n , is defined as

$$\{(x_0, x_1, \dots, x_n) \mid x_i \geq 0, x_0 + \dots + x_n = 1\}.$$


Hence it's the convex hull of some vertices $[v_0, \dots, v_n]$. Note that we keep track of the order $v_0, \dots, v_n$ of the vertices, for reasons that will soon become clear.

Given a topological space X , a **singular n -simplex** is a map $\sigma : \Delta^n \rightarrow X$.

Example 30.1.2 (Singular simplices)

- (a) Since $\Delta^0 = [v_0]$ is just a point, a singular 0-simplex X is just a point of X .
- (b) Since $\Delta^1 = [v_0, v_1]$ is an interval, a singular 1-simplex X is just a path in X .
- (c) Since $\Delta^2 = [v_0, v_1, v_2]$ is an equilateral triangle, a singular 2-simplex X looks a "disk" in X .

Here is a picture of all three in a space X :



The arrows aren't strictly necessary, but I've included them to help keep track of the "order" of the vertices; this will be useful in just a moment.

Now we're going to do something much like when we were talking about Stokes' theorem: we'll put a boundary ∂ operator on the singular n -simplices. This will give us a formal linear sums of n -simplices $\sum_k a_k \sigma_k$, which we call an **n -chain**.

In that case,

Definition 30.1.3. Given a singular n -simplex σ with vertices $[v_0, \dots, v_n]$, note that for every i we have an $(n-1)$ simplex $[v_0, \dots, v_{i-1}, v_{i+1}, \dots, v_n]$. The **boundary operator** ∂ is then defined by

$$\partial(\sigma) \stackrel{\text{def}}{=} \sum_i (-1)^i [v_0, \dots, v_{i-1}, v_{i+1}, \dots, v_n].$$

The boundary operator then extends linearly to n -chains:

$$\partial \left(\sum_k a_k \sigma_k \right) \stackrel{\text{def}}{=} \sum_k a_k \partial(\sigma_k).$$

By convention, a 0-chain has empty boundary.

Example 30.1.4 (Boundary operator)

Consider the chains depicted in Example 30.1.2. Then

- (a) $\partial\sigma^0 = 0$.
- (b) $\partial(\sigma^1) = [v_1] - [v_0]$: it's the "difference" of the 0-chain corresponding to point v_1 and the 0-chain corresponding to point v_0 .
- (c) $\partial(\sigma^2) = [v_0, v_1] - [v_0, v_2] + [v_1, v_2]$; i.e. one can think of it as the sum of the three oriented arrows which make up the "sides" of σ^2 .
- (d) Notice that if we take the boundary again, we get

$$\begin{aligned} \partial(\partial(\sigma^2)) &= \partial([v_0, v_1]) - \partial([v_0, v_2]) + \partial([v_1, v_2]) \\ &= ([v_1] - [v_0]) - ([v_2] - [v_0]) + ([v_2] - [v_1]) \\ &= 0. \end{aligned}$$

The fact that $\partial^2 = 0$ is of course not a coincidence.

Theorem 30.1.5 ($\partial^2 = 0$)

For any chain c , $\partial(\partial(c)) = 0$.

Proof. Essentially identical to Problem 28B: this is just a matter of writing down a bunch of $\sum$ signs. Diligent readers are welcome to try the computation. $\square$

Remark 30.1.6. The eerie similarity between the chains used to integrate differential forms and the chains in homology is not a coincidence. The de Rham cohomology, discussed much later, will make the relation explicit.

§30.2 The singular homology groups


Prototypical example for this section: Probably $H_n(S^m)$, especially the case $m = n = 1$.

Let X be a topological space, and let $C_n(X)$ be the free abelian group of n -chains of X that we defined earlier. Our work above gives us a boundary operator ∂ , so we have a sequence of maps

$$\dots \xrightarrow{\partial} C_3(X) \xrightarrow{\partial} C_2(X) \xrightarrow{\partial} C_1(X) \xrightarrow{\partial} C_0(X) \xrightarrow{\partial} 0$$

(here I'm using 0 for the trivial group, which is standard notation for abelian groups.) We'll call this the **singular chain complex**.

Now, how does this let us detect holes in the space? To see why, let's consider an annulus, with a 1-chain c drawn in red:



Notice that

$$\partial c = ([v_1] - [v_0]) - ([v_2] - [v_0]) + ([v_2] - [v_1]) = 0$$

and so we can say this 1-chain c is a “cycle”, because it has trivial boundary. However, c is not itself the boundary of any 2-chain, because of the hole in the center of the space — it's impossible to “fill in” the interior of c ! So, we have detected the hole by the algebraic fact that

$$c \in \ker \left(C_1(X) \xrightarrow{\partial} C_0(X) \right) \quad \text{but} \quad c \notin \text{im} \left(C_2(X) \xrightarrow{\partial} C_1(X) \right).$$

Indeed, if the hole was not present then this statement would be false.

We can capture this idea in any dimension, as follows.

Definition 30.2.1. Let

$$\dots \xrightarrow{\partial} C_2(X) \xrightarrow{\partial} C_1(X) \xrightarrow{\partial} C_0(X) \xrightarrow{\partial} 0$$

as above. We say that $c \in C_n(X)$ is:

- a **cycle** if $\partial c \in \ker \left(C_n(X) \xrightarrow{\partial} C_{n-1}(X) \right)$, and
- a **boundary** if $c \in \text{im} \left(C_{n+1}(X) \xrightarrow{\partial} C_n(X) \right)$.

Denote the cycles and boundaries by $Z_n(X), B_n(X) \subseteq C_n(X)$, respectively.

Question 30.2.2. Just to get you used to the notation: check that B_n and Z_n are themselves abelian groups, and that $B_n(X) \subseteq Z_n(X) \subseteq C_n(X)$.

The key point is that we can now define:

Definition 30.2.3. The n th homology group $H_n(X)$ is defined as

$$H_n(X) \stackrel{\text{def}}{=} Z_n(X)/B_n(X).$$

Example 30.2.4 (The zeroth homology group)

Let's compute $H_0(X)$ for a topological space X . We take $C_0(X)$, which is just formal linear sums of points of X .

First, we consider the kernel of $\partial : C_0(X) \rightarrow 0$, so the kernel of ∂ is the entire space $C_0(X)$: that is, every point is a “cycle”.

Now, what is the boundary? The main idea is that $[b] - [a] = 0$ if and only if there's a 1-chain which connects a to b , i.e. there is a path from a to b . In particular,

$$X \text{ path connected} \implies H_0(X) \cong \mathbb{Z}.$$

More generally, we have

Proposition 30.2.5 (Homology groups split into path-connected components)

If $X = \bigcup_{\alpha} X_{\alpha}$ is a decomposition into path-connected components, then we have

$$H_n(X) \cong \bigoplus_{\alpha} H_n(X_{\alpha}).$$

In particular, if X has r path-connected components, then $H_0(X) \cong \mathbb{Z}^{\oplus r}$.

(If it's surprising to see $\mathbb{Z}^{\oplus r}$, remember that an abelian group is the same thing as a $\mathbb{Z}$ -module, so the notation $G \oplus H$ is customary in place of $G \times H$ when G, H are abelian.)

Now let's investigate the first homology group.

Theorem 30.2.6 (Hurewicz theorem)

Let X be path-connected. Then $H_1(X)$ is the *abelianization* of $\pi_1(X, x_0)$.

We won't prove this but you can see it roughly from the example. The group $H_1(X)$ captures the same information as $\pi_1(X, x_0)$: a cycle (in $Z_1(X)$) corresponds to the same thing as the loops we studied in $\pi_1(X, x_0)$, and the boundaries (in $B_1(X)$, i.e. the things we mod out by) are exactly the nullhomotopic loops in $\pi_1(X, x_0)$. The difference is that $H_1(X)$ allows loops to commute, whereas $\pi_1(X, x_0)$ does not.

Example 30.2.7 (The first homology group of the annulus)

To give a concrete example, consider the annulus X above. We found a chain c that wrapped once around the hole of X . The point is that in fact,

$$H_1(X) = \langle c \rangle \cong \mathbb{Z}$$


which is to say the chains $c, 2c, \dots$ are all not the same in $H_1(X)$, but that any other 1-chain is equivalent to one of these. This captures the fact that X is really just S^1 .

Example 30.2.8 (An explicit boundary in S^1)

In $X = S^1$, let a be the uppermost point and b the lowermost point. Let c be the simplex from a to b along the left half of the circle, and d the simplex from a to b along the right half. Finally, let γ be the simplex which represents a loop γ from a to itself, wrapping once counterclockwise around S^1 . We claim that in $H^1(S^1)$ we have

$$\gamma = c - d$$

which geometrically means that $c - d$ represents wrapping once around the circle (which is of course what we expect).



Indeed this can be seen from the picture above, where we have drawn a 2-simplex whose boundary is exactly $\gamma - c + d$. The picture is somewhat metaphorical: in reality $v_0 = v_1 = a$, and the entire 2-simplex is embedded in S^1 . This is why singular homology is so-called: the images of the simplex can sometimes look quite “singular”.

Example 30.2.9 (The first homology group of the figure eight)

Consider X_8 (see Example 21.2.9). Both homology and homotopy see the two loops in X_8 , call them a and b . The difference is that in $\pi_1(X_8, x_0)$, these two loops are not allowed to commute: we don't have $ab \neq ba$, because the group operation in π_1 is “concatenate paths”. But in the homology group $H_1(X)$ the way we add a and b is to add them formally, to get the 1-chain $a + b$. So

$$H_1(X) \cong \mathbb{Z}^{\oplus 2} \quad \text{while} \quad \pi_1(X, x_0) = \langle a, b \rangle.$$

Example 30.2.10 (The homology groups of S^2)

Consider S^2 , the two-dimensional sphere. Since it's path connected, we have $H_0(S^2) = \mathbb{Z}$. We also have $H_1(S^2) = 0$, for the same reason that $\pi_1(S^2)$ is trivial as well. On the other hand we claim that

$$H_2(S^2) \cong \mathbb{Z}.$$

The elements of $H_2(S^2)$ correspond to wrapping S^2 in a tetrahedral bag (or two bags, or three bags, etc.). Thus, the second homology group lets us detect the spherical cavity of S^2 .

Actually, more generally it turns out that we will have

$$H_n(S^m) \cong \begin{cases} \mathbb{Z} & n = m \text{ or } n = 0 \\ 0 & \text{otherwise.} \end{cases}$$

Example 30.2.11 (Contractible spaces)

Given any contractible space X , it turns out that

$$H_n(X) \cong \begin{cases} \mathbb{Z} & n = 0 \\ 0 & \text{otherwise.} \end{cases}$$

The reason is that, like homotopy groups, it turns out that homology groups are homotopy invariant. (We'll prove this next section.) So the homology groups of contractible X are the same as those of a one-point space, which are those above.

Example 30.2.12 (Homology groups of the torus)

While we won't be able to prove it for a while, it turns out that

$$H_n(S^1 \times S^1) \cong \begin{cases} \mathbb{Z} & n = 0, 2 \\ \mathbb{Z}^{\oplus 2} & n = 1 \\ 0 & \text{otherwise.} \end{cases}$$

The homology group at 1 corresponds to our knowledge that $\pi_1(S^1 \times S^1) \cong \mathbb{Z}^2$ and the homology group at 2 detects the “cavity” of the torus.

This is fantastic and all, but how does one go about actually computing any homology groups? This will be a rather long story, and we'll have to do a significant amount of both algebra and geometry before we're really able to compute any homology groups. In what follows, it will often be helpful to keep track of which things are purely algebraic (work for any chain complex), and which parts are actually stating something which is geometrically true.

§30.3 The homology functor and chain complexes

As I mentioned before, the homology groups are homotopy invariant. This will be a similar song and dance as the work we did to create a functor $\pi_1 : \mathbf{hTop}_* \rightarrow \mathbf{Grp}$. Rather

than working slowly and pulling away the curtain to reveal the category theory at the end, we'll instead start with the category theory right from the start just to save some time.

Definition 30.3.1. The category $\mathbf{hTop}$ is defined as follows:

- Objects: topological spaces.
- Morphisms: *homotopy classes* of morphisms $X \rightarrow Y$.

In particular, X and Y are isomorphic in $\mathbf{hTop}$ if and only if they are homotopic.

You'll notice this is the same as $\mathbf{hTop}_*$, except without the basepoints.

Theorem 30.3.2 (Homology is a functor $\mathbf{hTop} \rightarrow \mathbf{Grp}$)

For any particular n , H_n is a functor $\mathbf{hTop} \rightarrow \mathbf{Grp}$. In particular,

- Given any map $f : X \rightarrow Y$, we get an induced map $f_* : H_n(X) \rightarrow H_n(Y)$.
- For two homotopic maps $f, g : X \rightarrow Y$, $f_* = g_*$.
- Two homotopic spaces X and Y have isomorphic homology groups: if $f : X \rightarrow Y$ is a homotopy then $f_* : H_n(X) \rightarrow H_n(Y)$ is an isomorphism.
- (Insert your favorite result about functors here.)

In order to do this, we have to describe how to take a map $f : X \rightarrow Y$ and obtain a map $H_n(f) : H_n(X) \rightarrow H_n(Y)$. Then we have to show that this map doesn't depend on the choice of homotopy. (This is the analog of the work we did with $f_{\sharp}$ before.) It turns out that this time around, proving this is much more tricky, and we will have to go back to the chain complex $C_{\bullet}(X)$ that we built at the beginning.

Algebra of chain complexes

Let's start with the algebra. First, I'll define the following abstraction of the complex to any sequence of abelian groups. Actually, though, it works in any category (not just $\mathbf{AbGrp}$). The strategy is as follows: we'll define everything that we need completely abstractly, then show that the geometry concepts we want correspond to this setting.

Definition 30.3.3. A **chain complex** is a sequence of groups A_n and maps

$$\dots \xrightarrow{\partial} A_{n+1} \xrightarrow{\partial} A_n \xrightarrow{\partial} A_{n-1} \xrightarrow{\partial} \dots$$

such that the composition of any two adjacent maps is the zero morphism. We usually denote this by $A_{\bullet}$.

The n th homology group $H_n(A_{\bullet})$ is defined as $\ker(A_n \rightarrow A_{n-1}) / \operatorname{im}(A_{n+1} \rightarrow A_n)$. Cycles and boundaries are defined in the same way as before.

Obviously, this is just an algebraic generalization of the structure we previously looked at, rid of all its original geometric context.

Definition 30.3.4. A **morphism of chain complexes** $f : A_\bullet \rightarrow B_\bullet$ is a sequence of maps f_n for every n such that the diagram

$$\begin{array}{ccccccc} \cdots & \xrightarrow{\partial_A} & A_{n+1} & \xrightarrow{\partial_A} & A_n & \xrightarrow{\partial_A} & A_{n-1} & \xrightarrow{\partial_A} & \cdots \\ & & \downarrow f_{n+1} & & \downarrow f_n & & \downarrow f_{n-1} & & \\ \cdots & \xrightarrow{\partial_B} & B_{n+1} & \xrightarrow{\partial_B} & B_n & \xrightarrow{\partial_B} & B_{n-1} & \xrightarrow{\partial_B} & \cdots \end{array}$$

commutes. Under this definition, the set of chain complexes becomes a category, which we denote $\mathbf{Cmplx}$.

Note that given a morphism of chain complexes $f : A_\bullet \rightarrow B_\bullet$, every cycle in A_n gets sent to a cycle in B_n , since the square

$$\begin{array}{ccc} A_n & \xrightarrow{\partial_A} & A_{n-1} \\ \downarrow f_n & & \downarrow \\ B_n & \xrightarrow{\partial_B} & B_{n-1} \end{array}$$

commutes. Similarly, every boundary in A_n gets sent to a boundary in B_n . Thus,

Every map of $f : A_\bullet \rightarrow B_\bullet$ gives a map $f_* : H_n(A) \rightarrow H_n(B)$ for every n .

Exercise 30.3.5. Interpret H_n as a functor $\mathbf{Cmplx} \rightarrow \mathbf{Grp}$.

Next, we want to define what it means for two maps f and g to be homotopic. Here's the answer:

Definition 30.3.6. Let $f, g : A_\bullet \rightarrow B_\bullet$. Suppose that one can find a map $P_n : A_n \rightarrow B_{n+1}$ for every n such that

$$g_n - f_n = \partial_B \circ P_n + P_{n-1} \circ \partial_A$$

Then P is a **chain homotopy** from f to g and f and g are **chain homotopic**.

We can draw a picture to illustrate this (warning: the diagonal dotted arrows do NOT commute with all the other arrows):

$$\begin{array}{ccccccc} \cdots & \xrightarrow{\partial_A} & A_{n+1} & \xrightarrow{\partial_A} & A_n & \xrightarrow{\partial_A} & A_{n-1} & \xrightarrow{\partial_A} & \cdots \\ & & \downarrow g-f & & \downarrow g-f & & \downarrow g-f & & \\ & & \swarrow P_n & & \swarrow P_{n-1} & & & & \\ \cdots & \xrightarrow{\partial_B} & B_{n+1} & \xrightarrow{\partial_B} & B_n & \xrightarrow{\partial_B} & B_{n-1} & \xrightarrow{\partial_B} & \cdots \end{array}$$

The definition is that in each slanted “parallelogram”, the $g - f$ arrow is the sum of the two compositions along the sides.

Remark 30.3.7. This equation should look terribly unmotivated right now, aside from the fact that we are about to show it does the right algebraic thing. Its derivation comes from the geometric context that we have deferred until the next section, where “homotopy” will naturally give “chain homotopy”.

Now, the point of this definition is that

Proposition 30.3.8 (Chain homotopic maps induce the same map on homology groups)

Let $f, g : A_\bullet \rightarrow B_\bullet$ be chain homotopic maps $A_\bullet \rightarrow B_\bullet$. Then the induced maps $f_*, g_* : H_n(A_\bullet) \rightarrow H_n(B_\bullet)$ coincide for each n .

Proof. It's equivalent to show $g - f$ gives the zero map on homology groups. In other words, we need to check that every cycle of A_n becomes a boundary of B_n under $g - f$.

Question 30.3.9. Verify that this is true. □

Geometry of chain complexes

Now let's fill in the geometric details of the picture above. First:

Lemma 30.3.10 (Map of space $\implies$ map of singular chain complexes)

Each $f : X \rightarrow Y$ induces a map $C_n(X) \rightarrow C_n(Y)$.


Proof. Take the composition

$$\Delta^n \xrightarrow{\sigma} X \xrightarrow{f} Y.$$

In other words, a path in X becomes a path in Y , et cetera. (It's not hard to see that the squares involving ∂ commute; check it if you like.) □

Now, what we need is to show that if $f, g : X \rightarrow Y$ are homotopic, then they are chain homotopic. To produce a chain homotopy, we need to take every n -simplex X to an $(n+1)$ -chain in Y , thus defining the map P_n .

Let's think about how we might do this. Let's take the n -simplex $\sigma : \Delta^n \rightarrow X$ and feed it through f and g ; pictured below is a 1-simplex σ (i.e. a path in X) which has been mapped into the space Y . Homotopy means the existence of a map $F : X \times [0, 1] \rightarrow Y$ such that $F(-, 0) = f$ and $F(-, 1) = g$, parts of which I've illustrated below with grey arrows in the image for Y .



This picture suggests how we might proceed: we want to create a 2-chain on Y given the 1-chains we've drawn. The homotopy F provides us with a "square" structure on Y , i.e. the square bounded by v_0, v_1, w_1, w_0 . We split this up into two triangles; and that's our 2-chain.

We can make this formal by taking $\Delta^1 \times [0, 1]$ (which is a square) and splitting it into two triangles. Then, if we apply $\sigma \times \text{id}$, we'll get an $(n+1)$ -cell in $X \times [0, 1]$, and

then finally applying F will map everything into our space Y . In our example, the final image is the 2-chain, consisting of two triangles, which in our picture can be written as $[v_0, v_1, w_1] - [v_0, w_0, w_1]$; the boundaries are given by the red, green, grey.

More generally, for an n -simplex $\phi = [x_0, \dots, x_n]$ we define the so-called *prism operator* P_n as follows. Set $v_i = f(x_i)$ and $w_i = g(x_i)$ for each i . Then, we let

$$P_n(\phi) \stackrel{\text{def}}{=} \sum_{i=0}^n (-1)^i (F \circ (\phi \times \text{id})) [v_0, \dots, v_i, w_i, \dots, w_n].$$

This is just the generalization of the construction above to dimensions $n > 1$; we split $\Delta^n \times [0, 1]$ into $n + 1$ simplices, map it into X by $\phi \times \text{id}$ and then push the whole thing into Y . The $(-1)^i$ makes sure that the “diagonal” faces all cancel off with each other.

We now claim that for every σ ,

$$\partial_Y(P_n(\sigma)) = g(\sigma) - f(\sigma) - P_{n-1}(\partial_X \sigma).$$

In the picture, $\partial_Y \circ P_n$ is the boundary of the entire prism (in the figure, this becomes the red, green, and grey lines, not including diagonal grey, which is cancelled out). The $g - f$ is the green minus the red, and the $P_{n-1} \circ \partial_X$ represents the grey edges of the prism (not including the diagonal line from v_1 to w_0). Indeed, one can check (just by writing down several $\sum$ signs) that the above identity holds.

So that gives the chain homotopy from f to g , completing the proof of Theorem 30.3.2.

§30.4 More examples of chain complexes

We now end this chapter by providing some more examples of chain complexes, which we’ll use in the next chapter to finally compute topological homology groups.

Example 30.4.1 (Reduced homology groups)

Suppose X is a (nonempty) topological space. One can augment the standard singular complex as follows: do the same thing as before, but augment the end by adding a $\mathbb{Z}$, as shown:

$$\cdots \rightarrow C_1(X) \rightarrow C_0(X) \xrightarrow{\varepsilon} \mathbb{Z} \rightarrow 0$$

Here ε is defined by $\varepsilon(\sum n_i p_i) = \sum n_i$ for points $p_i \in X$. (Recall that a 0-chain is just a formal sum of points!) We denote this **augmented singular chain complex** by $\tilde{C}_\bullet(X)$.

Question 30.4.2. What’s the homology of the above chain at $\mathbb{Z}$? (Hint: you need X nonempty.)

Definition 30.4.3. The homology groups of the augmented chain complex are called the **reduced homology groups** $\tilde{H}_n(X)$ of the space X .

Obviously $\tilde{H}_n(X) \cong H_n(X)$ for $n > 0$. But when $n = 0$, the map $H_0(X) \rightarrow \mathbb{Z}$ by ε has kernel $\tilde{H}_0(X)$, thus $H_0(X) \cong \tilde{H}_0(X) \oplus \mathbb{Z}$.

This is usually just an added convenience. For example, it means that if X is contractible, then all its reduced homology groups vanish, and thus we won’t have to keep fussing with the special $n = 0$ case.

Question 30.4.4. Given the claim earlier about $H_n(S^m)$, what should $\tilde{H}_n(S^m)$ be?

Example 30.4.5 (Relative chain groups)

Suppose X is a topological space, and $A \subseteq X$ a subspace. We can “mod out” by A by defining

$$C_n(X, A) \stackrel{\text{def}}{=} C_n(X)/C_n(A)$$

for every n . Thus chains contained entirely in A are trivial.

Then, the usual ∂ on $C_n(X)$ generates a new chain complex

$$\dots \xrightarrow{\partial} C_{n+1}(X, A) \xrightarrow{\partial} C_n(X, A) \xrightarrow{\partial} C_{n-1}(X, A) \xrightarrow{\partial} \dots$$

This is well-defined since ∂ takes $C_n(A)$ into $C_{n-1}(A)$.

Definition 30.4.6. The homology groups of the relative chain complex are the **relative homology groups** and denoted $H_n(X, A)$.

One naïve guess is that this might equal $H_n(X)/H_n(A)$. This is not true and in general doesn’t even make sense; if we take X to be $\mathbb{R}^2$ and $A = S^1$ a circle inside it, we have $H_1(X) = H_1(\mathbb{R}^2) = 0$ and $H_1(S^1) = \mathbb{Z}$.

Another guess is that $H_n(X, A)$ might just be $\tilde{H}_n(X/A)$. This will turn out to be true for most reasonable spaces X and A , and we will discuss this when we reach the excision theorem.

Example 30.4.7 (Mayer-Vietoris sequence)

Suppose a space X is covered by two open sets U and V . We can define $C_n(U + V)$ as follows: it consists of chains such that each simplex is either entirely contained in U , or entirely contained in V .

Of course, ∂ then defines another chain complex

$$\dots \xrightarrow{\partial} C_{n+1}(U + V) \xrightarrow{\partial} C_n(U + V) \xrightarrow{\partial} C_{n-1}(U + V) \xrightarrow{\partial} \dots$$

So once again, we can define homology groups for this complex; we denote them by $H_n(U + V)$. Miraculously, it will turn out that $H_n(U + V) \cong H_n(X)$.

§30.5 Problems to think about

Problem 30A. For $n \geq 1$ show that the composition

$$S^{n-1} \hookrightarrow D^n \xrightarrow{F} S^{n-1}$$

cannot be the identity map on S^{n-1} for any F .

Problem 30B (Brouwer fixed point theorem). Use the previous problem to prove that any continuous function $f : D^n \rightarrow D^n$ has a fixed point.

31 The long exact sequence

In this chapter we introduce the key fact about chain complexes that will allow us to compute the homology groups of any space: the so-called “long exact sequence”.

For those that haven’t read about abelian categories: a sequence of morphisms of abelian groups

$$\cdots \rightarrow G_{n+1} \rightarrow G_n \rightarrow G_{n-1} \rightarrow \cdots$$

is **exact** if the image of any arrow is equal to the kernel of the next arrow. In particular,

- The map $0 \rightarrow A \rightarrow B$ is exact if and only if $A \rightarrow B$ is injective.
- the map $A \rightarrow B \rightarrow 0$ is exact if and only if $A \rightarrow B$ is surjective.

(On that note: what do you call a chain complex whose homology groups are all trivial?)
A short exact sequence is one of the form $0 \rightarrow A \hookrightarrow B \twoheadrightarrow C \rightarrow 0$.

§31.1 Short exact sequences and four examples

Prototypical example for this section: Relative sequence and Mayer-Vietoris sequence.

Let $\mathcal{A} = \text{AbGrp}$. Recall that we defined a morphism of chain complexes in $\mathcal{A}$ already.

Definition 31.1.1. Suppose we have a map of chain complexes

$$0 \rightarrow A_\bullet \xrightarrow{f} B_\bullet \xrightarrow{g} C_\bullet \rightarrow 0$$

It is said to be **short exact** if *each row* of the diagram below is short exact.

$$\begin{array}{ccccccc}
 & \vdots & & \vdots & & \vdots & \\
 & \downarrow \partial_A & & \downarrow \partial_B & & \downarrow \partial_C & \\
 0 & \longrightarrow & A_{n+1} & \xrightarrow{\subset f_{n+1}} & B_{n+1} & \xrightarrow{g_{n+1}} & C_{n+1} \longrightarrow 0 \\
 & & \downarrow \partial_A & & \downarrow \partial_B & & \downarrow \partial_C \\
 0 & \longrightarrow & A_n & \xrightarrow{\subset f_n} & B_n & \xrightarrow{g_n} & C_n \longrightarrow 0 \\
 & & \downarrow \partial_A & & \downarrow \partial_B & & \downarrow \partial_C \\
 0 & \longrightarrow & A_{n-1} & \xrightarrow{\subset f_{n-1}} & B_{n-1} & \xrightarrow{g_{n-1}} & C_{n-1} \longrightarrow 0 \\
 & & \downarrow \partial_A & & \downarrow \partial_B & & \downarrow \partial_C \\
 & & \vdots & & \vdots & & \vdots
 \end{array}$$

Example 31.1.2 (Mayer-Vietoris short exact sequence and its augmentation)

Let $X = U \cup V$ be an open cover. For each n consider

$$C_n(U \cap V) \hookrightarrow C_n(U) \oplus C_n(V) \twoheadrightarrow C_n(U + V)$$

$$c \longmapsto (c, -c)$$

$$(c, d) \longmapsto c + d$$

One can easily see (by taking a suitable basis) that the kernel of the latter map is exactly the image of the first map. This generates a short exact sequence

$$0 \rightarrow C_\bullet(U \cap V) \hookrightarrow C_\bullet(U) \oplus C_\bullet(V) \twoheadrightarrow C_\bullet(U + V) \rightarrow 0.$$

Example 31.1.3 (Augmented Mayer-Vietoris sequence)

We can *augment* each of the chain complexes in the Mayer-Vietoris sequence as well, by appending

$$\begin{array}{ccccccc} 0 & \longrightarrow & C_0(U \cap V) & \hookrightarrow & C_0(U) \oplus C_0(V) & \twoheadrightarrow & C_0(U + V) \longrightarrow 0 \\ & & \downarrow \varepsilon & & \downarrow \varepsilon \oplus \varepsilon & & \downarrow \varepsilon \\ 0 & \longrightarrow & \mathbb{Z} & \longrightarrow & \mathbb{Z} \oplus \mathbb{Z} & \longrightarrow & \mathbb{Z} \longrightarrow 0 \end{array}$$

to the bottom of the diagram. In other words we modify the above into

$$0 \rightarrow \tilde{C}_\bullet(U \cap V) \hookrightarrow \tilde{C}_\bullet(U) \oplus \tilde{C}_\bullet(V) \twoheadrightarrow \tilde{C}_\bullet(U + V) \rightarrow 0$$

where $\tilde{C}_\bullet$ is the chain complex defined in Definition 30.4.3.

Example 31.1.4 (Relative chain short exact sequence)

Since $C_n(X, A) \stackrel{\text{def}}{=} C_n(X)/C_n(A)$, we have a short exact sequence

$$0 \rightarrow C_\bullet(A) \hookrightarrow C_\bullet(X) \twoheadrightarrow C_\bullet(X, A) \rightarrow 0$$

for every space X and subspace A . This can be augmented: we get

$$0 \rightarrow \tilde{C}_\bullet(A) \hookrightarrow \tilde{C}_\bullet(X) \twoheadrightarrow C_\bullet(X, A) \rightarrow 0$$

by adding the final row

$$\begin{array}{ccccccc} 0 & \longrightarrow & C_0(A) & \hookrightarrow & C_0(X) & \twoheadrightarrow & C_0(X, A) \longrightarrow 0 \\ & & \downarrow \varepsilon & & \downarrow \varepsilon & & \downarrow \\ 0 & \longrightarrow & \mathbb{Z} & \xrightarrow{\text{id}} & \mathbb{Z} & \longrightarrow & 0 \longrightarrow 0. \end{array}$$

§31.2 The long exact sequence of homology groups

Consider a short exact sequence $0 \rightarrow A_\bullet \xrightarrow{f} B_\bullet \xrightarrow{g} C_\bullet \rightarrow 0$. Now, we know that we get induced maps of homology groups, i.e. we have

$$\begin{array}{ccccc} \vdots & & \vdots & & \vdots \\ H_{n+1}(A_\bullet) & \xrightarrow{f_*} & H_{n+1}(B_\bullet) & \xrightarrow{g_*} & H_{n+1}(C_\bullet) \\ H_n(A_\bullet) & \xrightarrow{f_*} & H_n(B_\bullet) & \xrightarrow{g_*} & H_n(C_\bullet) \\ H_n(A_\bullet) & \xrightarrow{f_*} & H_n(B_\bullet) & \xrightarrow{g_*} & H_n(C_\bullet) \\ \vdots & & \vdots & & \vdots \end{array}$$

But the theorem is that we can string these all together, taking each $H_{n+1}(C_\bullet)$ to $H_n(A_\bullet)$.

Theorem 31.2.1 (Short exact $\implies$ long exact)

Let $0 \rightarrow A_\bullet \xrightarrow{f} B_\bullet \xrightarrow{g} C_\bullet \rightarrow 0$ be *any* short exact sequence of chain complexes we like. Then there is an *exact* sequence

$$\begin{array}{ccccccc} & & & & \dots & \longrightarrow & H_{n+2}(C_\bullet) \\ & & & & \swarrow & \partial & \nearrow \\ H_{n+1}(A_\bullet) & \xrightarrow{f_*} & H_{n+1}(B_\bullet) & \xrightarrow{g_*} & H_{n+1}(C_\bullet) & & \\ & & \swarrow & \partial & \nearrow & & \\ H_n(A_\bullet) & \xrightarrow{f_*} & H_n(B_\bullet) & \xrightarrow{g_*} & H_n(C_\bullet) & & \\ & & \swarrow & \partial & \nearrow & & \\ H_{n-1}(A_\bullet) & \xrightarrow{f_*} & H_{n-1}(B_\bullet) & \xrightarrow{g_*} & H_{n-1}(C_\bullet) & & \\ & & \swarrow & \partial & \nearrow & & \\ H_{n-2}(A_\bullet) & \longrightarrow & \dots & & & & \end{array}$$

This is called a **long exact sequence** of homology groups.

Proof. A very long diagram chase, valid over any abelian category. (Alternatively, it's actually possible to use the snake lemma twice.) $\square$

Remark 31.2.2. The map $\partial : H_n(C_\bullet) \rightarrow H_{n-1}(A_\bullet)$ can be written explicitly as follows. Recall that H_n is “cycles modulo boundaries”, and consider the sub-diagram

$$\begin{array}{ccccc} B_n & \xrightarrow{g_n} & C_n & & \\ \downarrow \partial_B & & \downarrow \partial_C & & \\ A_{n-1} & \xrightarrow{f_{n-1}} & B_{n-1} & \xrightarrow{g_{n-1}} & C_{n-1} \end{array}$$

We need to take every cycle in C_n to a cycle in A_{n-1} . (Then we need to check a ton of “well-defined” issues, but let's put that aside for now.)

Suppose $c \in C_n$ is a cycle (so $\partial_C(c) = 0$). By surjectivity, there is a $b \in B_n$ with $g_n(b) = c$, which maps down to $\partial_B(b)$. Now, the image of $\partial_B(b)$ under g_{n-1} is zero by commutativity of the square, and so we can pull back under f_{n-1} to get a unique element of A_{n-1} (by exactness at B_{n-1}).

In summary: we go “left, down, left” to go from c to a :

$$\begin{array}{ccccc}
 & b & \xrightarrow{g_n} & \boxed{c} & \\
 & \downarrow \partial_B & & \downarrow \partial_C & \\
 \boxed{a} & \xrightarrow{f_{n-1}} & \partial_B(b) & \xrightarrow{g_{n-1}} & 0
 \end{array}$$

Exercise 31.2.3. Check quickly that the recovered a is actually a cycle, meaning $\partial_A(a) = 0$. (You’ll need another row, and the fact that $\partial_B^2 = 0$.)

The final word is that:

Short exact sequences of chain complexes give long exact sequences of homology groups.

In particular, let us take the four examples given earlier.

Example 31.2.4 (Mayer-Vietoris long exact sequence, provisional version)

The Mayer-Vietoris ones give, for $X = U \cup V$ an open cover,

$$\cdots \rightarrow H_n(U \cap V) \rightarrow H_n(U) \oplus H_n(V) \rightarrow H_n(U + V) \rightarrow H_{n-1}(U \cap V) \rightarrow \cdots$$

and its reduced version

$$\cdots \rightarrow \tilde{H}_n(U \cap V) \rightarrow \tilde{H}_n(U) \oplus \tilde{H}_n(V) \rightarrow \tilde{H}_n(U + V) \rightarrow \tilde{H}_{n-1}(U \cap V) \rightarrow \cdots$$

This version is “provisional” because in the next section we will replace $H_n(U + V)$ and $\tilde{H}_n(U + V)$ with something better. As for the relative homology sequences, we have:

Theorem 31.2.5 (Long exact sequence for relative homology)

Let X be a space, and let $A \subseteq X$ be a subspace. There are long exact sequences

$$\cdots \rightarrow H_n(A) \rightarrow H_n(X) \rightarrow H_n(X, A) \rightarrow H_{n-1}(A) \rightarrow \cdots$$

and

$$\cdots \rightarrow \tilde{H}_n(A) \rightarrow \tilde{H}_n(X) \rightarrow H_n(X, A) \rightarrow \tilde{H}_{n-1}(A) \rightarrow \cdots$$

The exactness of these sequences will give **tons of information** about $H_n(X)$ if only we knew something about what $H_n(U + V)$ or $H_n(X, A)$ looked like. This is the purpose of the next chapter.

§31.3 The Mayer-Vietoris sequence

Prototypical example for this section: The computation of $H_n(S^m)$ by splitting S^m into two hemispheres.

Now that we have done so much algebra, we need to invoke some geometry. There are two major geometric results in the Napkin. One is the excision theorem, which we discuss next chapter. The other we present here, which will let us take advantage of the Mayer-Vietoris sequence. The proofs are somewhat involved and are thus omitted; see [Ha02] for details.

The first theorem is that the notation $H_n(U + V)$ that we have kept until now is redundant, and can be replaced with just $H_n(X)$:

Theorem 31.3.1 (Open cover homology theorem)

Consider the inclusion $\iota : C_\bullet(U + V) \hookrightarrow C_\bullet(X)$. Then ι induces an isomorphism

$$H_n(U + V) \cong H_n(X).$$

Remark 31.3.2. In fact, this is true for any open cover (even uncountable), not just those with two covers $U \cup V$. But we only state the special case with two open sets, because this is what is needed for Example 31.1.2.

So, Example 31.1.2 together with the above theorem implies, after replacing all the $H_n(U + V)$'s with $H_n(X)$'s:

Theorem 31.3.3 (Mayer-Vietoris long exact sequence)

If $X = U \cup V$ is an open cover, then we have long exact sequences

$$\cdots \rightarrow H_n(U \cap V) \rightarrow H_n(U) \oplus H_n(V) \rightarrow H_n(X) \rightarrow H_{n-1}(U \cap V) \rightarrow \cdots$$

and

$$\cdots \rightarrow \tilde{H}_n(U \cap V) \rightarrow \tilde{H}_n(U) \oplus \tilde{H}_n(V) \rightarrow \tilde{H}_n(X) \rightarrow \tilde{H}_{n-1}(U \cap V) \rightarrow \cdots$$

At long last, we can compute the homology groups of the spheres.


Theorem 31.3.4 (The homology groups of S^m)

For integers m and n ,

$$\tilde{H}_n(S^m) \cong \begin{cases} \mathbb{Z} & n = m \\ 0 & \text{otherwise.} \end{cases}$$

The generator $\tilde{H}_n(S^n)$ is an n -cell which covers S^n exactly once (for example, the generator for $\tilde{H}_1(S^1)$ is a loop which wraps around S^1 once).

Proof. This one's fun, so I'll only spoil the case $m = 1$, and leave the rest to you. Decompose the circle S^1 into two arcs U and V , as shown:



Each of U and V is contractible, so all their reduced homology groups vanish. Moreover, $U \cap V$ is homotopy equivalent to two points, hence

$$\tilde{H}_n(U \cap V) \cong \begin{cases} \mathbb{Z} & n = 0 \\ 0 & \text{otherwise.} \end{cases}$$

Now consider again the segment of the short exact sequence


$$\cdots \rightarrow \underbrace{\tilde{H}_n(U) \oplus \tilde{H}_n(V)}_{=0} \rightarrow \tilde{H}_n(S^1) \xrightarrow{\partial} \tilde{H}_{n-1}(U \cap V) \rightarrow \underbrace{\tilde{H}_{n-1}(U) \oplus \tilde{H}_{n-1}(V)}_{=0} \rightarrow \cdots$$

From this we derive that $\tilde{H}_n(S^1)$ is $\mathbb{Z}$ for $n = 1$ and 0 elsewhere.

It remains to analyze the generators of $\tilde{H}_1(S^1)$. Note that the isomorphism was given by the connecting homomorphism ∂ , which is given by a “left, down, left” procedure (Remark 31.2.2) in the diagram

$$\begin{array}{ccc} C_1(U) \oplus C_1(V) & \longrightarrow & C_1(U + V) \\ \downarrow \partial \oplus \partial & & \\ C_0(U \cap V) & \longrightarrow & C_0(U) \oplus C_0(V). \end{array}$$

Mark the points a and b as shown in the two disjoint paths of $U \cap V$.



Then $a - b$ is a cycle which represents a generator of $H_0(U \cap V)$. We can find the pre-image of ∂ as follows: letting c and d be the chains joining a and b , with c contained in U , and d contained in V , the diagram completes as

$$\begin{array}{ccc} (c, d) & \longmapsto & c - d \\ \downarrow & & \\ a - b & \longmapsto & (a - b, a - b) \end{array}$$

In other words $\partial(c - d) = a - b$, so $c - d$ is a generator for $\tilde{H}^1(S^1)$.

Thus we wish to show that $c - d$ is (in $H^1(S^1)$) equivalent to the loop γ wrapping around S^1 once, counterclockwise. This was illustrated in Example 30.2.8. $\square$

Thus, the key idea in Mayer-Vietoris is that

Mayer-Vietoris lets us compute $H_n(X)$ by splitting X into two open sets.

Here are some more examples.


Proposition 31.3.5 (The homology groups of the figure eight)

Let $X = S^1 \wedge S^1$ be the figure eight. Then

$$\tilde{H}_n(X) \cong \begin{cases} \mathbb{Z}^{\oplus 2} & n = 1 \\ 0 & \text{otherwise.} \end{cases}$$

The generators for $\tilde{H}_1(X)$ are the two loops of the figure eight.

Proof. Again, for simplicity we work with reduced homology groups. Let U be the “left” half of the figure eight plus a little bit of the right, as shown below.



The set V is defined symmetrically. In this case $U \cap V$ is contractible, while each of U and V is homotopic to S^1 .

Thus, we can read a segment of the long exact sequence as

$$\cdots \rightarrow \underbrace{\tilde{H}_n(U \cap V)}_{=0} \rightarrow \tilde{H}_n(U) \oplus \tilde{H}_n(V) \rightarrow \tilde{H}_n(X) \rightarrow \underbrace{\tilde{H}_{n-1}(U \cap V)}_{=0} \rightarrow \cdots$$

So we get that $\tilde{H}_n(X) \cong \tilde{H}_n(S^1) \oplus \tilde{H}_n(S^1)$. The claim about the generators follows from the fact that, according to the isomorphism above, the generators of $\tilde{H}_n(X)$ are the generators of $\tilde{H}_n(U)$ and $\tilde{H}_n(V)$, which we described geometrically in the last theorem. $\square$


Up until now, we have been very fortunate that we have always been able to make certain parts of the space contractible. This is not always the case, and in the next example we will have to actually understand the maps in question to complete the solution.

Proposition 31.3.6 (Homology groups of the torus)

Let $X = S^1 \times S^1$ be the torus. Then

$$\tilde{H}_n(X) = \begin{cases} \mathbb{Z}^{\oplus 2} & n = 1 \\ \mathbb{Z} & n = 2 \\ 0 & \text{otherwise.} \end{cases}$$

Proof. To make our diagram look good on 2D paper, we'll represent the torus as a square with its edges identified, though three-dimensionally the picture makes sense as well. Consider U (shaded light orange) and V (shaded green) as shown. (Note that V is connected due to the identification of the left and right (blue) edges, even if it doesn't look connected in the picture).



In the three dimensional picture, U and V are two cylinders which together give the torus. This time, U and V are each homotopic to S^1 , and the intersection $U \cap V$ is the disjoint union of two circles: thus $\tilde{H}_1(U \cap V) \cong \mathbb{Z} \oplus \mathbb{Z}$, and $H_0(U \cap V) \cong \mathbb{Z}^{\oplus 2} \implies \tilde{H}_0(U \cap V) \cong \mathbb{Z}$.

For $n \geq 3$, we have

$$\cdots \rightarrow \underbrace{\tilde{H}_n(U \cap V)}_{=0} \rightarrow \tilde{H}_n(U) \oplus \tilde{H}_n(V) \rightarrow \tilde{H}_n(X) \rightarrow \underbrace{\tilde{H}_{n-1}(U \cap V)}_{=0} \rightarrow \cdots$$

and so $H_n(X) \cong 0$ for $n \geq 3$. Also, we have $H_0(X) \cong \mathbb{Z}$ since X is path-connected. So it remains to compute $H_2(X)$ and $H_1(X)$.

Let's find $H_2(X)$ first. We first consider the segment

$$\cdots \rightarrow \underbrace{\tilde{H}_2(U) \oplus \tilde{H}_2(V)}_{=0} \rightarrow \tilde{H}_2(X) \xrightarrow{\delta} \underbrace{\tilde{H}_1(U \cap V)}_{\cong \mathbb{Z} \oplus \mathbb{Z}} \xrightarrow{\phi} \underbrace{\tilde{H}_1(U) \oplus \tilde{H}_1(V)}_{\cong \mathbb{Z} \oplus \mathbb{Z}} \rightarrow \cdots$$

Unfortunately, this time it's not immediately clear what $\tilde{H}_2(X)$ because we only have one zero at the left. In order to do this, we have to actually figure out what the maps δ and ϕ look like. Note that, as we'll see, ϕ isn't an isomorphism even though the groups are isomorphic.

The presence of the zero term has allowed us to make the connecting map δ injective. First, $\tilde{H}_2(X)$ is isomorphic to the image of δ , which is exactly the kernel of the arrow ϕ inserted. To figure out what $\ker \phi$ is, we have to think back to how the map $C_\bullet(U \cap V) \rightarrow C_\bullet(U) \oplus C_\bullet(V)$ was constructed: it was $c \mapsto (c, -c)$. So the induced maps of homology groups is actually what you would guess: a 1-cycle z in $\tilde{H}_1(U \cap V)$ gets sent $(z, -z)$ in $\tilde{H}_1(U) \oplus \tilde{H}_1(V)$.

In particular, consider the two generators z_1 and z_2 of $\tilde{H}_1(U \cap V) = \mathbb{Z} \oplus \mathbb{Z}$, i.e. one cycle in each connected component of $U \cap V$. (To clarify: $U \cap V$ consists of two "wristbands"; z_i wraps around the i th one once.) Moreover, let α_U denote a generator of $\tilde{H}_1(U) \cong \mathbb{Z}$, and α_V a generator of $\tilde{H}_1(V) \cong \mathbb{Z}$. Then we have that

$$z_1 \mapsto (\alpha_U, -\alpha_V) \quad \text{and} \quad z_2 \mapsto (\alpha_U, -\alpha_V).$$

(The signs may differ on which direction you pick for the generators; note that $\mathbb{Z}$ has two possible generators.) We can even format this as a matrix:

$$\phi = \begin{bmatrix} 1 & 1 \\ -1 & -1 \end{bmatrix}.$$

And we observe $\phi(z_1 - z_2) = 0$, meaning this map has nontrivial kernel! That is,

$$\ker \phi = \langle z_1 - z_2 \rangle \cong \mathbb{Z}.$$

Thus, $\tilde{H}_2(X) \cong \text{im } \delta \cong \ker \phi \cong \mathbb{Z}$. We'll also note that $\text{im } \phi$ is the set generated by $(\alpha_U, -\alpha_V)$; (in particular $\text{im } \phi \cong \mathbb{Z}$ and the quotient by $\text{im } \phi$ is $\mathbb{Z}$ too).

The situation is similar with $\tilde{H}_1(X)$: this time, we have

$$\dots \xrightarrow{\phi} \underbrace{\tilde{H}_1(U) \oplus \tilde{H}_1(V)}_{\cong \mathbb{Z} \oplus \mathbb{Z}} \xrightarrow{\psi} \tilde{H}_1(X) \xrightarrow{\partial} \underbrace{\tilde{H}_0(U \cap V)}_{\cong \mathbb{Z}} \rightarrow \underbrace{\tilde{H}_0(U) \oplus \tilde{H}_0(V)}_{=0} \rightarrow \dots$$

and so we know that the connecting map ∂ is surjective, hence $\text{im } \partial \cong \mathbb{Z}$. Now, we also have

$$\begin{aligned} \ker \partial &\cong \text{im } \psi \cong \left(\tilde{H}_1(U) \oplus \tilde{H}_1(V) \right) / \ker \psi \\ &\cong \left(\tilde{H}_1(U) \oplus \tilde{H}_1(V) \right) / \text{im } \phi \\ &\cong \mathbb{Z} \end{aligned}$$

by what we knew about $\text{im } \phi$ already. To finish off we need some algebraic tricks. The first is Proposition 25.5.1, which gives us a short exact sequence

$$0 \rightarrow \underbrace{\ker \partial}_{\cong \text{im } \psi \cong \mathbb{Z}} \hookrightarrow \tilde{H}_1(X) \twoheadrightarrow \underbrace{\text{im } \partial}_{\cong \mathbb{Z}} \rightarrow 0.$$

You should satisfy yourself that $\tilde{H}_1(X) \cong \mathbb{Z} \oplus \mathbb{Z}$ is the only possibility, but we'll prove this rigorously with Lemma 31.3.7. $\square$

Note that the previous example is of a different attitude than the previous ones, because we had to figure out what the maps in the long exact sequence actually were to even compute the groups. In principle, you could also figure out all the isomorphisms in the previous proof and explicitly compute the generators of $\tilde{H}_1(S^1 \times S^1)$, but to avoid getting bogged down in detail I won't do so here.

Finally, to fully justify the last step, we present:

Lemma 31.3.7 (Splitting lemma)

For a short exact sequence $0 \rightarrow A \xrightarrow{f} B \xrightarrow{g} C \rightarrow 0$ of abelian groups, the following are equivalent:

- (a) There exists $p : B \rightarrow A$ such that $A \xrightarrow{f} B \xrightarrow{p} A$ is the identity.
- (b) There exists $s : C \rightarrow B$ such that $C \xrightarrow{s} B \xrightarrow{g} C$ is the identity.
- (c) There is an isomorphism from B to $A \oplus C$ such that the diagram

$$\begin{array}{ccccccc} & & & B & & & \\ & & \nearrow f & & \searrow g & & \\ 0 & \longrightarrow & A & & & C & \longrightarrow 0 \\ & & \searrow & \cong \downarrow & \nearrow & & \\ & & & A \oplus C & & & \end{array}$$

commutes. (The maps attached to $A \oplus C$ are the obvious ones.)

In particular, (b) holds anytime C is free.

In these cases we say the short exact sequence **splits**. The point is that

An exact sequence which splits let us obtain B given A and C .

In particular, for $C = \mathbb{Z}$ or any free abelian group, condition (b) is necessarily true. So, once we obtained the short exact sequence $0 \rightarrow \mathbb{Z} \rightarrow \tilde{H}_1(X) \rightarrow \mathbb{Z} \rightarrow 0$, we were done.

Remark 31.3.8. Unfortunately, not all exact sequences split: An example of a short exact sequence which doesn't split is

$$0 \rightarrow \mathbb{Z}_2 \xrightarrow{\times 2} \mathbb{Z}_4 \rightarrow \mathbb{Z}_2 \rightarrow 0$$

since it is not true that $\mathbb{Z}_4 \cong \mathbb{Z}_2 \oplus \mathbb{Z}_2$.

Remark 31.3.9. The splitting lemma is true in any abelian category. The “direct sum” is the colimit of the two objects A and C .

§31.4 Problems to think about

Problem 31A. Complete the proof of Theorem 31.3.4, i.e. compute $H_n(S^m)$ for all m and n . (Try doing $m = 2$ first, and you'll see how to proceed.)

Problem 31B. Compute the reduced homology groups of $\mathbb{R}^n$ with $p \geq 1$ points removed.

Problem 31C*. Let $n \geq 1$ and $k \geq 0$ be integers. Compute $H_k(\mathbb{R}^n, \mathbb{R}^n \setminus \{0\})$.

Problem 31D* (Triple long exact sequence). Let $B \subseteq A \subseteq X$ be subspaces. Show that there is a long exact sequence

$$\cdots \rightarrow H_n(B, A) \rightarrow H_n(X, A) \rightarrow H_n(X, B) \rightarrow H_{n-1}(B, A) \rightarrow \cdots$$



Problem 31E* (Klein bottle). Show that the reduced homology groups of the Klein bottle K are given by

$$\tilde{H}_n(K) = \begin{cases} \mathbb{Z} \oplus \mathbb{Z}_2 & n = 1 \\ 0 & \text{otherwise.} \end{cases}$$

32 Excision and relative homology

We have already seen how to use the Mayer-Vietoris sequence: we started with a sequence

$$\cdots \rightarrow H_n(U \cap V) \rightarrow H_n(U) \oplus H_n(V) \rightarrow H_n(U + V) \rightarrow H_{n-1}(U \cap V) \rightarrow \cdots$$

and its reduced version, then appealed to the geometric fact that $H_n(U + V) \cong H_n(X)$. This allowed us to algebraically make computations on $H_n(X)$.

In this chapter, we turn our attention to the long exact sequence associated to the chain complex

$$0 \rightarrow C_n(A) \hookrightarrow C_n(X) \twoheadrightarrow C_n(X, A) \rightarrow 0.$$

The setup will look a lot like the previous two chapters, except in addition to $H_n : \mathbf{hTop} \rightarrow \mathbf{Grp}$ we will have a functor $H_n : \mathbf{hPairTop} \rightarrow \mathbf{Grp}$ which takes a pair (X, A) to $H_n(X, A)$. Then, we state (again without proof) the key geometric result, and use this to make deductions.

§32.1 The long exact sequences

Recall Theorem 31.2.5, which says that the sequences

$$\cdots \rightarrow H_n(A) \rightarrow H_n(X) \rightarrow H_n(X, A) \rightarrow H_{n-1}(A) \rightarrow \cdots$$

and

$$\cdots \rightarrow \tilde{H}_n(A) \rightarrow \tilde{H}_n(X) \rightarrow H_n(X, A) \rightarrow \tilde{H}_{n-1}(A) \rightarrow \cdots$$

are long exact. By Problem 31D* we even have a long exact sequence

$$\cdots \rightarrow H_n(B, A) \rightarrow H_n(X, A) \rightarrow H_n(X, B) \rightarrow H_{n-1}(B, A) \rightarrow \cdots$$

for $A \subseteq B \subseteq X$. An application of the second long exact sequence above gives:

Lemma 32.1.1 (Homology relative to contractible spaces)

Let X be a topological space, and let $A \subseteq X$ be contractible. For all n ,

$$H_n(X, A) \cong \tilde{H}_n(X).$$

Proof. Since A is contractible, we have $\tilde{H}_n(A) = 0$ for every n . For each n there's a segment of the long exact sequence given by

$$\cdots \rightarrow \underbrace{\tilde{H}_n(A)}_{=0} \rightarrow \tilde{H}_n(X) \rightarrow H_n(X, A) \rightarrow \underbrace{\tilde{H}_{n-1}(A)}_{=0} \rightarrow \cdots$$

So since $0 \rightarrow \tilde{H}_n(X) \rightarrow H_n(X, A) \rightarrow 0$ is exact, this means $H_n(X, A) \cong \tilde{H}_n(X)$. $\square$

In particular, the theorem applies if A is a single point. The case $A = \emptyset$ is also worth noting. We compile these results into a lemma:

Lemma 32.1.2 (Relative homology generalizes absolute homology)

Let X be any space, and $* \in X$ a point. Then for all n ,

$$H_n(X, \{*\}) \cong \tilde{H}_n(X) \quad \text{and} \quad H_n(X, \emptyset) = H_n(X).$$

§32.2 The category of pairs

Since we now have an $H_n(X, A)$ instead of just $H_n(X)$, a natural next step is to create a suitable category of *pairs* and give ourselves the same functorial setup as before.

Definition 32.2.1. Let $\emptyset \neq A \subseteq X$ and $\emptyset \neq B \subseteq X$ be subspaces, and consider a map $f : X \rightarrow Y$. If $f(A) \subseteq B$ we write

$$f : (X, A) \rightarrow (Y, B).$$

We say f is a **map of pairs**, between the pairs (X, A) and (Y, B) .

Definition 32.2.2. We say that $f, g : (X, A) \rightarrow (Y, B)$ are **pair-homotopic** if they are “homotopic through maps of pairs”.

More formally, a **pair-homotopy** $f, g : (X, A) \rightarrow (Y, B)$ is a map $F : [0, 1] \times X \rightarrow Y$, which we’ll write as $F_t(X)$, such that F is a homotopy of the maps $f, g : X \rightarrow Y$ and each F_t is itself a map of pairs.

Thus, we naturally arrive at two categories:

- **PairTop**, the category of *pairs* of topological spaces, and
- **hPairTop**, the same category except with maps only equivalent up to homotopy.

Definition 32.2.3. As before, we say pairs (X, A) and (Y, B) are **pair-homotopy equivalent** if they are isomorphic in **hPairTop**. An isomorphism of **hPairTop** is a **pair-homotopy equivalence**.

We can do the same song and dance as before with the prism operator to obtain:

Lemma 32.2.4 (Induced maps of relative homology)

We have a functor

$$H_n : \mathbf{hPairTop} \rightarrow \mathbf{Grp}.$$

That is, if $f : (X, A) \rightarrow (Y, B)$ then we obtain an induced map

$$f_* : H_n(X, A) \rightarrow H_n(Y, B).$$

and if two such f and g are pair-homotopic then $f_* = g_*$.

Now, we want an analog of contractible spaces for our pairs: i.e. pairs of spaces (X, A) such that $H_n(X, A) = 0$. The correct definition is:

Definition 32.2.5. Let $A \subseteq X$. We say that A is a **deformation retract** of X if there is a map of pairs $r : (X, A) \rightarrow (A, A)$ which is a pair homotopy equivalence.

Example 32.2.6 (Examples of deformation retracts)

- If a single point p is a deformation retract of a space X , then X is contractible, since the retraction $r : X \rightarrow \{*\}$ (when viewed as a map $X \rightarrow X$) is homotopic to the identity map $\text{id}_X : X \rightarrow X$.
- The punctured disk $D^2 \setminus \{0\}$ deformation retracts onto its boundary S^1 .
- More generally, $D^n \setminus \{0\}$ deformation retracts onto its boundary S^{n-1} .
- Similarly, $\mathbb{R}^n \setminus \{0\}$ deformation retracts onto a sphere S^{n-1} .

Of course in this situation we have that

$$H_n(X, A) \cong H_n(A, A) = 0.$$

Exercise 32.2.7. Show that if $A \subseteq V \subseteq X$, and A is a deformation retract of V , then $H_n(X, A) \cong H_n(X, V)$ for all n . (Use Problem 31D*. Solution in next section.)

§32.3 Excision

Now for the key geometric result, which is the analog of Theorem 31.3.1 for our relative homology groups.

Theorem 32.3.1 (Excision)

Let $Z \subseteq A \subseteq X$ be subspaces such that the closure of Z is contained in the interior of A . Then the inclusion $\iota(X \setminus Z, A \setminus Z) \hookrightarrow (X, A)$ (viewed as a map of pairs) induces an isomorphism of relative homology groups

$$H_n(X \setminus Z, A \setminus Z) \cong H_n(X, A).$$

This means we can *excise* (delete) a subset Z of A in computing the relative homology groups $H_n(X, A)$. This should intuitively make sense: since we are “modding out by points in A ”, the internals of the point A should not matter so much.

The main application of excision is to decide when $H_n(X, A) \cong \tilde{H}_n(X/A)$. Answer:

Theorem 32.3.2 (Relative homology $\implies$ quotient space)

Let X be a space and A be a subspace such that A is a deformation retract of some neighborhood $V \subseteq X$. Then the quotient map $q : X \rightarrow X/A$ induces an isomorphism

$$H_n(X, A) \cong H_n(X/A, A/A) \cong \tilde{H}_n(X/A).$$

Proof. By hypothesis, we can consider the following maps of pairs:

$$\begin{aligned} r &: (V, A) \rightarrow (A, A) \\ q &: (X, A) \rightarrow (X/A, A/A) \\ \hat{q} &: (X - A, V - A) \rightarrow (X/A - A/A, V/A - A/A). \end{aligned}$$

Moreover, r is a pair-homotopy equivalence. Considering the long exact sequence of of a triple (which was Problem 31D*) we have a diagram

$$\begin{array}{ccccccc} H_n(V, A) & \longrightarrow & H_n(X, A) & \xrightarrow{f} & H_n(X, V) & \longrightarrow & H_{n-1}(V, A) \\ \cong \downarrow r & & & & & & \cong \downarrow r \\ \underbrace{H_n(A, A)}_{=0} & & & & & & \underbrace{H_{n-1}(A, A)}_{=0} \end{array}$$

where the isomorphisms arise since r is a pair-homotopy equivalence. So f is an isomorphism. Similarly the map

$$g : H_n(X/A, A/A) \rightarrow H_n(X/A, V/A)$$

is an isomorphism.

Now, consider the commutative diagram

$$\begin{array}{ccccc} H_n(X, A) & \xrightarrow{f} & H_n(X, V) & \xleftarrow{\text{Excise}} & H_n(X - A, V - A) \\ \downarrow q_* & & & & \downarrow \hat{q}_* \cong \\ H_n(X/A, A/A) & \xrightarrow{g} & H_n(X/A, V/A) & \xleftarrow{\text{Excise}} & H_n(X/A - A/A, V/A - A/A) \end{array}$$

and observe that the rightmost arrow $\hat{q}_*$ is an isomorphism, because outside of A the map $\hat{q}$ is the identity. We know f and g are isomorphisms, as are the two arrows marked with “Excise” (by excision). From this we conclude that q_* is an isomorphism. Of course we already know that homology relative to a point is just the relative homology groups (this is the important case of Lemma 32.1.1). $\square$

§32.4 Some applications

One nice application of excision is to compute $\tilde{H}_n(X \vee Y)$.

Theorem 32.4.1 (Homology of wedge sums)

Let X and Y be spaces with basepoints $x_0 \in X$ and $y_0 \in Y$, and assuming each point is a deformation retract of some neighborhood. Then for every n we have

$$\tilde{H}_n(X \vee Y) = \tilde{H}_n(X) \oplus \tilde{H}_n(Y).$$

Proof. Apply Theorem 32.3.2 with the subset $\{x_0, y_0\}$ of $X \amalg Y$,

$$\begin{aligned} \tilde{H}_n(X \vee Y) &\cong \tilde{H}_n((X \amalg Y)/\{x_0, y_0\}) \cong H_n(X \amalg Y, \{x_0, y_0\}) \\ &\cong H_n(X, \{x_0\}) \oplus H_n(Y, \{y_0\}) \\ &\cong \tilde{H}_n(X) \oplus \tilde{H}_n(Y). \end{aligned} \quad \square$$

Another application is to give a second method of computing $H_n(S^m)$. To do this, we will prove that

$$\tilde{H}_n(S^m) \cong \tilde{H}_{n-1}(S^{m-1})$$

for any $n, m > 1$. However,

- $\tilde{H}_0(S^n)$ is $\mathbb{Z}$ for $n = 0$ and 0 otherwise.
- $\tilde{H}_n(S^0)$ is $\mathbb{Z}$ for $m = 0$ and 0 otherwise.

So by induction on $\min\{m, n\}$ we directly obtain that

$$\tilde{H}_n(S^m) \cong \begin{cases} \mathbb{Z} & m = n \\ 0 & \text{otherwise} \end{cases}$$

which is what we wanted.

To prove the claim, let's consider the exact sequence formed by the pair $X = D^2$ and $A = S^1$.

Example 32.4.2 (The long exact sequence for $(X, A) = (D^2, S^1)$)

Consider D^2 (which is contractible) with boundary S^1 . Clearly S^1 is a deformation retraction of $D^2 \setminus \{0\}$, and if we fuse all points on the boundary together we get $D^2/S^1 \cong S^2$. So we have a long exact sequence

$$\begin{array}{ccccc}
 \tilde{H}_2(S^1) & \longrightarrow & \underbrace{\tilde{H}_2(D^2)}_{=0} & \longrightarrow & \tilde{H}_2(S^2) \\
 & \searrow & & & \uparrow \\
 \tilde{H}_1(S^1) & \longrightarrow & \underbrace{\tilde{H}_1(D^2)}_{=0} & \longrightarrow & \tilde{H}_1(S^2) \\
 & \searrow & & & \uparrow \\
 \tilde{H}_0(S^1) & \longrightarrow & \underbrace{\tilde{H}_0(D^2)}_{=0} & \longrightarrow & \underbrace{\tilde{H}_0(S^2)}_{=0}
 \end{array}$$

From this diagram we read that

$$\dots, \quad \tilde{H}_3(S^2) = \tilde{H}_2(S^1), \quad \tilde{H}_2(S^2) = \tilde{H}_1(S^1), \quad \tilde{H}_1(S^2) = \tilde{H}_0(S^1).$$

More generally, the exact sequence for the pair $(X, A) = (D^m, S^{m-1})$ shows that $\tilde{H}_n(S^m) \cong \tilde{H}_{n-1}(S^{m-1})$, which is the desired conclusion.

§32.5 Invariance of dimension

Here is one last example of an application of excision.

Definition 32.5.1. Let X be a space and $p \in X$ a point. The k th **local homology group** of p at X is defined as

$$H_k(X, X \setminus \{p\}).$$

Note that for any neighborhood U of p , we have by excision that

$$H_k(X, X \setminus \{p\}) \cong H_k(U, U \setminus \{p\}).$$

Thus this local homology group only depends on the space near p .

Theorem 32.5.2 (Invariance of dimension, Brouwer 1910)

Let $U \subseteq \mathbb{R}^n$ and $V \subseteq \mathbb{R}^m$ be nonempty open sets. If U and V are homeomorphic, then $m = n$.

Proof. Consider a point $x \in U$ and its local homology groups. By excision,

$$H_k(\mathbb{R}^n, \mathbb{R}^n \setminus \{x\}) \cong H_k(U, U \setminus \{x\}).$$

But since $\mathbb{R}^n \setminus \{x\}$ is homotopic to S^{n-1} , the long exact sequence of Theorem 31.2.5 tells us that

$$H_k(\mathbb{R}^n, \mathbb{R}^n \setminus \{x\}) \cong \begin{cases} \mathbb{Z} & k = n \\ 0 & \text{otherwise.} \end{cases}$$

Analogously, given $y \in V$ we have

$$H_k(\mathbb{R}^m, \mathbb{R}^m \setminus \{y\}) \cong H_k(V, V \setminus \{y\}).$$

If $U \cong V$, we thus deduce that

$$H_k(\mathbb{R}^n, \mathbb{R}^n \setminus \{x\}) \cong H_k(\mathbb{R}^m, \mathbb{R}^m \setminus \{y\})$$

for all k . This of course can only happen if $m = n$. □

§32.6 Problems to think about

Problem 32A. Let $X = S^1 \times S^1$ and $Y = S^1 \vee S^1 \vee S^2$. Show that

$$H_n(X) \cong H_n(Y)$$

for every integer n .

Problem 32B (Hatcher §2.1 exercise 18). Consider $\mathbb{Q} \subset \mathbb{R}$. Compute $\tilde{H}_1(\mathbb{R}, \mathbb{Q})$.

Problem 32C*. What are the local homology groups of a topological n -manifold?

Problem 32D. Let

$$X = \{(x, y) \mid x \geq 0\} \subseteq \mathbb{R}^2$$

denote the half-plane. What are the local homology groups of points in X ?

33 Bonus: Cellular homology

We now introduce cellular homology, which essentially lets us compute the homology groups of any CW complex we like.

§33.1 Degrees

Prototypical example for this section: $z \mapsto z^d$ has degree d .

For any $n > 0$ and map $f : S^n \rightarrow S^n$, consider

$$f_* : \underbrace{H_n(S^n)}_{\cong \mathbb{Z}} \rightarrow \underbrace{H_n(S^n)}_{\cong \mathbb{Z}}$$

which must be multiplication by some constant d . This d is called the **degree** of f , denoted $\deg f$.

Question 33.1.1. Show that $\deg(f \circ g) = \deg(f) \deg(g)$.

Example 33.1.2 (Degree)

- (a) For $n = 1$, the map $z \mapsto z^k$ (viewing $S^1 \subseteq \mathbb{C}$) has degree k .
- (b) A reflection map $(x_0, x_1, \dots, x_n) \mapsto (-x_0, x_1, \dots, x_n)$ has degree -1 ; we won't prove this, but geometrically this should be clear.
- (c) The antipodal map $x \mapsto -x$ has degree $(-1)^{n+1}$ since it's the composition of $n + 1$ reflections as above. We denote this map by $-\text{id}$.

Obviously, if f and g are homotopic, then $\deg f = \deg g$. In fact, a theorem of Hopf says that this is a classifying invariant: anytime $\deg f = \deg g$, we have that f and g are homotopic.

One nice application of this:

Theorem 33.1.3 (Hairy ball theorem)


If $n > 0$ is even, then S^n doesn't have a continuous field of nonzero tangent vectors.

Proof. If the vectors are nonzero then WLOG they have norm 1; that is for every x we have an orthogonal unit vector $v(x)$. Then we can construct a homotopy map $F : S^n \times [0, 1] \rightarrow S^n$ by

$$(x, t) \mapsto (\cos \pi t)x + (\sin \pi t)v(x).$$

which gives a homotopy from id to $-\text{id}$. So $\deg(\text{id}) = \deg(-\text{id})$, which means $1 = (-1)^{n+1}$ so n must be odd. $\square$

Of course, the one can construct such a vector field whenever n is odd. For example, when $n = 1$ such a vector field is drawn below.



§33.2 Cellular chain complex

Before starting, we state:

Lemma 33.2.1 (CW homology groups)

Let X be a CW complex. Then

$$H_k(X^n, X^{n-1}) \cong \begin{cases} \mathbb{Z}^{\oplus \#n\text{-cells of } X} & k = n \\ 0 & \text{otherwise.} \end{cases}$$

and

$$H_k(X^n) \cong \begin{cases} H_k(X) & k \leq n-1 \\ 0 & k \geq n+1. \end{cases}$$

Proof. The first part is immediate by noting that (X^n, X^{n-1}) is a good pair and X^n/X^{n-1} is a wedge sum of two spheres. For the second part, fix k and note that, as long as $n \leq k-1$ or $n \geq k+2$,

$$\underbrace{H_{k+1}(X^n, X^{n-1})}_{=0} \rightarrow H_k(X^{n-1}) \rightarrow H_k(X^n) \rightarrow \underbrace{H_k(X^n, X^{n-1})}_{=0}.$$

So we have isomorphisms

$$H_k(X^{k-1}) \cong H_k(X^{k-2}) \cong \dots \cong H_k(X^0) = 0$$

and

$$H_k(X^{k+1}) \cong H_k(X^{k+2}) \cong \dots \cong H_k(X). \quad \square$$


So, we know that the groups $H_k(X^k, X^{k-1})$ are super nice: they are free abelian with basis given by the cells of X . So, we give them a name:

Definition 33.2.2. For a CW complex X , we define

$$\text{Cells}_k(X) = H_k(X^k, X^{k-1})$$

where $\text{Cells}_0(X) = H_0(X^0, \emptyset) = H_0(X^0)$ by convention. So $\text{Cells}_k(X)$ is an abelian group with basis given by the k -cells of X .

Now, using $\text{Cells}_k = H_k(X^k, X^{k-1})$ let's use our long exact sequence and try to string together maps between these. Consider the following diagram.



The idea is that we have taken all the exact sequences generated by adjacent skeletons, and strung them together at the groups $H_k(X^k)$, with half the exact sequences being laid out vertically and the other half horizontally.

In that case, composition generates a sequence of dotted maps between the $H_k(X^k, X^{k-1})$ as shown.

Question 33.2.3. Show that the composition of two adjacent dotted arrows is zero.

So from the diagram above, we can read off a sequence of arrows

$$\dots \xrightarrow{d_5} \text{Cells}_4(X) \xrightarrow{d_4} \text{Cells}_3(X) \xrightarrow{d_3} \text{Cells}_2(X) \xrightarrow{d_2} \text{Cells}_1(X) \xrightarrow{d_1} \text{Cells}_0(X) \xrightarrow{d_0} 0.$$

This is a chain complex, called the **cellular chain complex**; as mentioned before before all the homology groups are free, but these ones are especially nice because for most

reasonable CW complexes, they are also finitely generated (unlike the massive $C_\bullet(X)$ that we had earlier). In other words, the $H_k(X^k, X^{k-1})$ are especially nice “concrete” free groups that one can actually work with.

The other reason we care is that in fact:

Theorem 33.2.4 (Cellular chain complex gives $H_n(X)$)

The k th homology group of the cellular chain complex is isomorphic to $H_k(X)$.

Proof. Follows from the diagram; Problem 33D. $\square$

A nice application of this is to define the **Euler characteristic** of a finite CW complex X . Of course we can write

$$\chi(X) = \sum_n (-1)^n \cdot \#(n\text{-cells of } X)$$

which generalizes the familiar $V - E + F$ formula. However, this definition is unsatisfactory because it depends on the choice of CW complex, while we actually want $\chi(X)$ to only depend on the space X itself (and not how it was built). In light of this, we prove that:

Theorem 33.2.5 (Euler characteristic via Betti numbers)

For any finite CW complex X we have

$$\chi(X) = \sum_n (-1)^n \text{rank } H_n(X).$$

Thus $\chi(X)$ does not depend on the choice of CW decomposition. The numbers

$$b_n = \text{rank } H_n(X)$$

are called the **Betti numbers** of X . In fact, we can use this to define $\chi(X)$ for any reasonable space; we are happy because in the (frequent) case that X is a CW complex,

Proof. We quote the fact that if $0 \rightarrow A \rightarrow B \rightarrow C \rightarrow D \rightarrow 0$ is exact then $\text{rank } B + \text{rank } D = \text{rank } A + \text{rank } C$. Then for example the row

$$\underbrace{H_2(X^1)}_{=0} \rightarrow H_2(X^2) \hookrightarrow H_2(X^2, X^1) \xrightarrow{\partial_2} H_1(X^1) \twoheadrightarrow \underbrace{H_1(X^2)}_{\cong H_1(X)} \rightarrow \underbrace{H_1(X^2, X^1)}_{=0}$$

from the cellular diagram gives

$$\#(2\text{-cells}) + \text{rank } H_1(X) = \text{rank } H_2(X^2) + \text{rank } H_1(X^1).$$

More generally,

$$\#(k\text{-cells}) + \text{rank } H_{k-1}(X) = \text{rank } H_k(X^k) + \text{rank } H_{k-1}(X^{k-1})$$

which holds also for $k = 0$ if we drop the H_{-1} terms (since $\#0\text{-cells} = \text{rank } H_0(X^0)$ is obvious). Multiplying this by $(-1)^k$ and summing across $k \geq 0$ gives the conclusion. $\square$

Example 33.2.6 (Examples of Betti numbers)

- (a) The Betti numbers of S^n are $b_0 = b_n = 1$, and zero elsewhere. The Euler characteristic is $1 + (-1)^n$.
- (b) The Betti numbers of a torus $S^1 \times S^1$ are $b_0 = 1$, $b_1 = 2$, $b_2 = 1$, and zero elsewhere. Thus the Euler characteristic is 0.
- (c) The Betti numbers of $\mathbb{CP}^n$ are $b_0 = b_2 = \cdots = b_{2n} = 1$, and zero elsewhere. Thus the Euler characteristic is $n + 1$.
- (d) The Betti numbers of the Klein bottle are $b_0 = 1$, $b_1 = 1$ and zero elsewhere. Thus the Euler characteristic is 0, the same as the sphere (also since their CW structures use the same number of cells).

One notices that in the “nice” spaces S^n , $S^1 \times S^1$ and $\mathbb{CP}^n$ there is a nice symmetry in the Betti numbers, namely $b_k = b_{n-k}$. This is true more generally; see Poincaré duality and Problem 35A[†].

§33.3 The cellular boundary formula

In fact, one can describe explicitly what the maps d_n are. Recalling that $H_k(X^k, X^{k-1})$ has a basis the k -cells of X , we obtain:

Theorem 33.3.1 (Cellular boundary formula for $k = 1$)

For $k = 1$,

$$d_1 : \text{Cells}_1(X) \rightarrow \text{Cells}_0(X)$$

is just the boundary map.

Theorem 33.3.2 (Cellular boundary for $k > 1$)

Let $k > 1$ be a positive integer. Let e^k be an k -cell, and let $\{e_\beta^{k-1}\}_\beta$ denote all $(k-1)$ -cells of X . Then

$$d_k : \text{Cells}_k(X) \rightarrow \text{Cells}_{k-1}(X)$$

is given on basis elements by

$$d_k(e^k) = \sum_{\beta} d_{\beta} e_{\beta}^{k-1}$$

where d_{β} is the degree of the composed map

$$S^{k-1} = \partial D_{\beta}^k \xrightarrow{\text{attach}} X^{k-1} \twoheadrightarrow S_{\beta}^{k-1}.$$

Here the first arrow is the attaching map for e^k and the second arrow is the quotient of collapsing $X^{k-1} \setminus e_{\beta}^{k-1}$ to a point.


This gives us an algorithm for computing homology groups of a CW complex:

- Construct the cellular chain complex, where $\text{Cells}_k(X)$ is $\mathbb{Z}^{\oplus \#k\text{-cells}}$.
- $d_1 : \text{Cells}_1(X) \rightarrow \text{Cells}_0(X)$ is just the boundary map (so $d_1(e^1)$ is the difference of the two endpoints).
- For any $k > 1$, we compute $d_k : \text{Cells}_k(X) \rightarrow \text{Cells}_{k-1}(X)$ on basis elements as follows. Repeat the following for each k -cell e^k :
 - For every $k-1$ cell e_β^{k-1} , compute the degree of the boundary of e^k welded onto the boundary of e_β^{k-1} , say d_β .
 - Then $d_k(e^k) = \sum_\beta d_\beta e_\beta^{k-1}$.
- Now we have the maps of the cellular chain complex, so we can compute the homologies directly (by taking the quotient of the kernel by the image).

We can use this for example to compute the homology groups of the torus again, as well as the Klein bottle and other spaces.

Example 33.3.3 (Cellular homology of a torus)

Consider the torus built from e^0 , e_a^1 , e_b^1 and e^2 as before, where e^2 is attached via the word $aba^{-1}b^{-1}$. For example, X^1 is



The cellular chain complex is

$$0 \longrightarrow \mathbb{Z}e^2 \xrightarrow{d_2} \mathbb{Z}e_a^1 \oplus \mathbb{Z}e_b^1 \xrightarrow{d_1} \mathbb{Z}e^0 \xrightarrow{d_0} 0$$

Now apply the cellular boundary formulas:

- Recall that d_1 was the boundary formula. We have $d_1(e_a^1) = e_0 - e_0 = 0$ and similarly $d_1(e_b^1) = 0$. So $d_1 = 0$.
- For d_2 , consider the image of the boundary e^2 on e_a^1 . Around X^1 , it wraps once around e_a^1 , once around e_b^1 , again around e_a^1 (in the opposite direction), and again around e_b^1 . Once we collapse the entire e_b^1 to a point, we see that the degree of the map is 0. So $d_2(e^2)$ has no e_a^1 coefficient. Similarly, it has no e_b^1 coefficient, hence $d_2 = 0$.

Thus

$$d_1 = d_2 = 0.$$

So at every map in the complex, the kernel of the map is the whole space while the image is $\{0\}$. So the homology groups are $\mathbb{Z}$, $\mathbb{Z}^{\oplus 2}$, $\mathbb{Z}$.

Example 33.3.4 (Cellular homology of the Klein bottle)

Let X be a Klein bottle. Consider cells e^0 , e_a^1 , e_b^1 and e^2 as before, but this time e^2 is attached via the word $abab^{-1}$. So d_1 is still zero, but this time we have $d_2(e^2) = 2e_a^1$ instead (why?). So our diagram looks like

$$\begin{array}{ccccccc} 0 & \longrightarrow & \mathbb{Z}e^2 & \xrightarrow{d_2} & \mathbb{Z}e_a^1 \oplus \mathbb{Z}e_b^1 & \xrightarrow{d_1} & \mathbb{Z}e^0 \xrightarrow{d_0} 0 \\ & & & & e^2 \longmapsto & 2e_a^1 & \\ & & & & e_a^1 \longmapsto & 0 & \\ & & & & e_b^1 \longmapsto & 0 & \end{array}$$

So we get that $H_0(X) \cong \mathbb{Z}$, but

$$H_1(X) \cong \mathbb{Z} \oplus \mathbb{Z}_2$$

this time (it is $\mathbb{Z}^{\oplus 2}$ modulo a copy of $2\mathbb{Z}$). Also, $\ker d_2 = 0$, and so now $H_2(X) = 0$.

§33.4 Problems to think about

Problem 33A[†]. Let n be a positive integer. Show that

$$H_k(\mathbb{CP}^n) \cong \begin{cases} \mathbb{Z} & k = 0, 2, 4, \dots, 2n \\ 0 & \text{otherwise.} \end{cases}$$

Problem 33B. Show that a non-surjective map $f : S^n \rightarrow S^n$ has degree zero.



Problem 33C (Moore spaces). Let $G_1, G_2, \dots, G_N$ be a sequence of finitely generated abelian groups. Construct a space X such that

$$\tilde{H}_n(X) \cong \begin{cases} G_n & 1 \leq n \leq N \\ 0 & \text{otherwise.} \end{cases}$$

Problem 33D. Prove Theorem 33.2.4, showing that the homology groups of X coincide with the homology groups of the cellular chain complex.



Problem 33E[†]. Let n be a positive integer. Show that

$$H_k(\mathbb{RP}^n) \cong \begin{cases} \mathbb{Z} & \text{if } k = 0 \text{ or } k = n \equiv 1 \pmod{2} \\ \mathbb{Z}_2 & \text{if } k \text{ is odd and } 0 < k < n \\ 0 & \text{otherwise.} \end{cases}$$

34 Singular cohomology

Here's one way to motivate this chapter. It turns out that:

- $H_n(\mathbb{CP}^2) \cong H_n(S^2 \vee S^4)$ for every n .
- $H_n(\mathbb{CP}^3) \cong H_n(S^2 \times S^4)$ for every n .

This is unfortunate, because if possible we would like to be able to tell these spaces apart (as they are in fact not homotopy equivalent), but the homology groups cannot tell the difference between them.

In this chapter, we'll define a *cohomology group* $H^n(X)$ and $H^n(Y)$. In fact, the H^n 's are completely determined by the H_n 's by the so-called *universal coefficient theorem*. However, it turns out that one can take all the cohomology groups and put them together to form a *cohomology ring* $H^\bullet$. We will then see that $H^\bullet(X) \not\cong H^\bullet(Y)$ as rings.

§34.1 Cochain complexes

Definition 34.1.1. A **cochain complex** $A^\bullet$ is algebraically the same as a chain complex, except that the indices increase. So it is a sequence of abelian groups

$$\dots \xrightarrow{\delta} A^{n-1} \xrightarrow{\delta} A^n \xrightarrow{\delta} A^{n+1} \xrightarrow{\delta} \dots$$

such that $\delta^2 = 0$. Notation-wise, we're now using subscripts, and use δ rather than ∂ . We define the **cohomology groups** by

$$H^n(A^\bullet) = \ker \left(A^n \xrightarrow{\delta} A^{n+1} \right) / \operatorname{im} \left(A^{n-1} \xrightarrow{\delta} A^n \right).$$

Example 34.1.2 (de Rham cohomology)

We have already met one example of a cochain complex: let M be a smooth manifold and $\Omega^k(M)$ be the additive group of k -forms on M . Then we have a cochain complex

$$0 \xrightarrow{d} \Omega^0(M) \xrightarrow{d} \Omega^1(M) \xrightarrow{d} \Omega^2(M) \xrightarrow{d} \dots$$

The resulting cohomology is called **de Rham cohomology**, described later.

Aside from de Rham's cochain complex, **the most common way to get a cochain complex is to dualize a chain complex**. Specifically, pick an abelian group G ; note that $\operatorname{Hom}(-, G)$ is a contravariant functor, and thus takes every chain complex

$$\dots \xrightarrow{\partial} A_{n+1} \xrightarrow{\partial} A_n \xrightarrow{\partial} A_{n-1} \xrightarrow{\partial} \dots$$

into a cochain complex: letting $A^n = \operatorname{Hom}(A_n, G)$ we obtain

$$\dots \xrightarrow{\delta} A^{n-1} \xrightarrow{\delta} A^n \xrightarrow{\delta} A^{n+1} \xrightarrow{\delta} \dots$$

where $\delta(A_n \xrightarrow{f} G) = A_{n+1} \xrightarrow{\partial} A \xrightarrow{f} G$.

These are the cohomology groups we study most in algebraic topology, so we give a special notation to them.

Definition 34.1.3. Given a chain complex $A_\bullet$ of abelian groups and another group G , we let

$$H^n(A_\bullet; G)$$

denote the cohomology groups of the dual cochain complex $A^\bullet$ obtained by applying $\text{Hom}(-, G)$. In other words, $H^n(A_\bullet; G) = H^n(A^\bullet)$.

§34.2 Cohomology of spaces

Prototypical example for this section: $C^0(X; G)$ all functions $X \rightarrow G$ while $H^0(X)$ are those functions $X \rightarrow G$ constant on path components.

The case of interest is our usual geometric situation, with $C_\bullet(X)$.

Definition 34.2.1. For a space X and abelian group G , we define $C^\bullet(X; G)$ to be the dual to the singular chain complex $C_\bullet(X)$, called the **singular cochain complex** of X ; its elements are called **cochains**.

Then we define the **cohomology groups** of the space X as

$$H^n(X; G) \stackrel{\text{def}}{=} H^n(C_\bullet(X); G) = H_n(C^\bullet(X; G)).$$

Remark 34.2.2. Note that if G is also a ring (like $\mathbb{Z}$ or $\mathbb{R}$), then $H^n(X; G)$ is not only an abelian group but actually a G -module.

Example 34.2.3 ($C^0(X; G)$, $C^1(X; G)$, and $H^0(X; G)$)

Let X be a topological space and consider $C^\bullet(X)$.

- $C_0(X)$ is the free abelian group on X , and $C^0(X) = \text{Hom}(C_0(X), G)$. So a 0-cochain is a function that takes every point of X to an element of G .
- $C_1(X)$ is the free abelian group on 1-simplices in X . So $C^1(X)$ needs to take every 1-simplex to an element of G .

Let's now try to understand $\delta : C^0(X) \rightarrow C^1(X)$. Given a 0-cochain $\phi \in C^0(X)$, i.e. a homomorphism $\phi : C_0(X) \rightarrow G$, what is $\delta\phi : C_1(X) \rightarrow G$? Answer:

$$\delta\phi : [v_0, v_1] \mapsto \phi([v_0]) - \phi([v_1]).$$

Hence, elements of $\ker(C^0 \xrightarrow{\delta} C^1) \cong H^0(X; G)$ are those cochains that are *constant on path-connected components*.

In particular, much like $H_0(X)$, we have

$$H^0(X) \cong G^{\oplus r}$$

if X has r path-connected components (where r is finite¹).

¹Something funny happens if X has *infinitely* many path-connected components: say $X = \coprod_\alpha X_\alpha$ over an infinite indexing set. In this case we have $H_0(X) = \bigoplus_\alpha G$ while $H^0(X) = \prod_\alpha G$. For homology we get a *direct sum* while for cohomology we get a *direct product*.

These are actually different for infinite indexing sets. For general modules $\bigoplus_\alpha M_\alpha$ is *defined* to only allow to have *finitely many* zero terms. (This was never mentioned earlier in the Napkin, since I only ever defined $M \oplus N$ and extended it to finite direct sums.) No such restriction holds for $\prod_\alpha G_\alpha$ a product of groups. This corresponds to the fact that $C_0(X)$ is formal linear sums of 0-chains (which, like all formal sums, are finite) from the path-connected components of G . But a cochain of $C^0(X)$ is a *function* from each path-connected component of X to G , where there is no restriction.

To the best of my knowledge, the higher cohomology groups $H^n(X; G)$ (or even the cochain groups $C^n(X; G) = \text{Hom}(C_n(X), G)$) are harder to describe concretely.

Abuse of Notation 34.2.4. In this chapter the only cochain complexes we will consider are dual complexes as above. So, any time we write a chain complex $A^\bullet$ it is implicitly given by applying $\text{Hom}(-, G)$ to $A_\bullet$.

§34.3 Cohomology of spaces is functorial

We now check that the cohomology groups still exhibit the same nice functorial behavior. First, let's categorize the previous results we had:

Question 34.3.1. Define CoCmplx the category of cochain complexes.

Exercise 34.3.2. Interpret $\text{Hom}(-, G)$ as a contravariant functor from

$$\text{Hom}(-, G) : \text{Cmplx}^{\text{op}} \rightarrow \text{CoCmplx}.$$

This means in particular that given a chain map $f : A_\bullet \rightarrow B_\bullet$, we naturally obtain a dual map $f^\vee : B^\bullet \rightarrow A^\bullet$.

Question 34.3.3. Interpret $H^n : \text{CoCmplx} \rightarrow \text{Grp}$ as a functor. Compose these to get a contravariant functor $H^n(-; G) : \text{Cmplx}^{\text{op}} \rightarrow \text{Grp}$.

Then in exact analog to our result that $H_n : \text{hTop} \rightarrow \text{Grp}$ we have:

Theorem 34.3.4 ($H^n(-; G) : \text{hTop}^{\text{op}} \rightarrow \text{Grp}$)

For every n , $H^n(-; G)$ is a contravariant functor from hTop^{op} to Grp .

Proof. The idea is to leverage the work we already did in constructing the prism operator earlier. First, we construct the entire sequence of functors from $\text{Top}^{\text{op}} \rightarrow \text{Grp}$:

$$\begin{array}{ccccccc} \text{Top}^{\text{op}} & \xrightarrow{C_\bullet} & \text{Cmplx}^{\text{op}} & \xrightarrow{\text{Hom}(-; G)} & \text{CoCmplx} & \xrightarrow{H^n} & \text{Grp} \\ \\ \begin{array}{c} X \\ \downarrow f \\ Y \end{array} & \xrightarrow{\quad} & \begin{array}{c} C_\bullet(X) \\ \downarrow f_\# \\ C_\bullet(Y) \end{array} & \xrightarrow{\quad} & \begin{array}{c} C^\bullet(X; G) \\ \uparrow f^\sharp \\ C^\bullet(Y; G) \end{array} & \xrightarrow{\quad} & \begin{array}{c} H^n(X; G) \\ \uparrow f^* \\ H^n(Y; G) \end{array} \end{array}$$

Here $f^\sharp = (f_\#)^\vee$, and f^* is the resulting induced map on homology groups of the cochain complex.

So as before all we have to show is that $f \simeq g$, then $f^* = g^*$. Recall now that there is a prism operator such that $f_\# - g_\# = P\partial + \partial P$. If we apply the entire functor $\text{Hom}(-; G)$ we get that $f^\sharp - g^\sharp = \delta P^\vee + P^\vee \delta$ where $P^\vee : C^{n+1}(Y; G) \rightarrow C^n(X; G)$. So $f^\sharp$ and $g^\sharp$ are chain homotopic thus $f^* = g^*$. $\square$

§34.4 Universal coefficient theorem

We now wish to show that the cohomology groups are determined up to isomorphism by the homology groups: given $H_n(A_\bullet)$, we can extract $H^n(A_\bullet; G)$. This is achieved by the *universal coefficient theorem*.

Theorem 34.4.1 (Universal coefficient theorem)

Let $A_\bullet$ be a chain complex of *free* abelian groups, and let G be another abelian group. Then there is a natural short exact sequence

$$0 \rightarrow \text{Ext}(H_{n-1}(A_\bullet), G) \rightarrow H^n(A_\bullet; G) \xrightarrow{h} \text{Hom}(H_n(A_\bullet), G) \rightarrow 0.$$

In addition, this exact sequence is *split* so in particular

$$H^n(C_\bullet; G) \cong \text{Ext}(H_{n-1}(A_\bullet), G) \oplus \text{Hom}(H_n(A_\bullet), G).$$

Fortunately, in our case of interest, $A_\bullet$ is $C_\bullet(X)$ which is by definition free.

There are two things we need to explain, what the map h is and the map Ext is.

It's not too hard to guess how

$$h : H^n(A_\bullet; G) \rightarrow \text{Hom}(H_n(A_\bullet), G)$$

is defined. An element of $H^n(A_\bullet; G)$ is represented by a function which sends a cycle in A_n to an element of G . The content of the theorem is to show that h is surjective with kernel $\text{Ext}(H_{n-1}(A_\bullet), G)$.

What about Ext ? It turns out that $\text{Ext}(-, G)$ is the so-called **Ext functor**, defined as follows. Let H be an abelian group, and consider a **free resolution** of H , by which we mean an exact sequence

$$\dots \xrightarrow{f_2} F_1 \xrightarrow{f_1} F_0 \xrightarrow{f_0} H \rightarrow 0$$

with each F_i free. Then we can apply $\text{Hom}(-, G)$ to get a cochain complex

$$\dots \xleftarrow{f_2^\vee} \text{Hom}(F_1, G) \xleftarrow{f_1^\vee} \text{Hom}(F_0, G) \xleftarrow{f_0^\vee} \text{Hom}(H, G) \leftarrow 0.$$

but *this cochain complex need not be exact* (in categorical terms, $\text{Hom}(-, G)$ does not preserve exactness). We define

$$\text{Ext}(H, G) \stackrel{\text{def}}{=} \ker(f_2^\vee) / \text{im}(f_1^\vee)$$

and it's a theorem that this doesn't depend on the choice of the free resolution. There's a lot of homological algebra that goes into this, which I won't take the time to discuss; but the upshot of the little bit that I did include is that the Ext functor is very easy to compute in practice, since you can pick any free resolution you want and compute the above.

Lemma 34.4.2 (Computing the Ext functor)

For any abelian groups G, H, H' we have

- (a) $\text{Ext}(H \oplus H', G) = \text{Ext}(H, G) \oplus \text{Ext}(H', G)$.
- (b) $\text{Ext}(H, G) = 0$ for H free, and
- (c) $\text{Ext}(\mathbb{Z}_n, G) = G/nG$.

Proof. For (a), note that if $\cdots \rightarrow F_1 \rightarrow F_0 \rightarrow H \rightarrow 0$ and $\cdots \rightarrow F'_1 \rightarrow F'_0 \rightarrow F'_0 \rightarrow H' \rightarrow 0$ are free resolutions, then so is $F_1 \oplus F'_1 \rightarrow F_0 \oplus F'_0 \rightarrow H \oplus H' \rightarrow 0$.

For (b), note that $0 \rightarrow H \rightarrow H \rightarrow 0$ is a free resolution.

Part (c) follows by taking the free resolution

$$0 \rightarrow \mathbb{Z} \xrightarrow{\times n} \mathbb{Z} \rightarrow \mathbb{Z}_n \rightarrow 0$$

and applying $\text{Hom}(-, G)$ to it.

Question 34.4.3. Finish the proof of (c) from here. □

Question 34.4.4. Some Ext practice: compute $\text{Ext}(\mathbb{Z}^{\oplus 2015}, G)$ and $\text{Ext}(\mathbb{Z}_{30}, \mathbb{Z}_4)$.

§34.5 Example computation of cohomology groups

Prototypical example for this section: Possibly $H^n(S^m)$.

The universal coefficient theorem gives us a direct way to compute any cohomology groups, provided we know the homology ones.

Example 34.5.1 (Cohomology groups of S^m)

It is straightforward to compute $H^n(S^m)$ now: all the Ext terms vanish since $H_n(S^m)$ is always free, and hence we obtain that

$$H^n(S^m) \cong \text{Hom}(H_n(S^m), G) \cong \begin{cases} G & n = m, n = 0 \\ 0 & \text{otherwise.} \end{cases}$$

Example 34.5.2 (Cohomology groups of torus)

This example has no nonzero Ext terms either, since this time $H^n(S^1 \times S^1)$ is always free. So we obtain

$$H^n(S^1 \times S^1) \cong \text{Hom}(H_n(S^1 \times S^1), G).$$

Since $H_n(S^1 \times S^1)$ is $\mathbb{Z}, \mathbb{Z}^{\oplus 2}, \mathbb{Z}$ in dimensions $n = 1, 2, 1$ we derive that

$$H^n(S^1 \times S^1) \cong \begin{cases} G & n = 0, 2 \\ G^{\oplus 2} & n = 1. \end{cases}$$

From these examples one might notice that:

Lemma 34.5.3 (0th homology groups are just duals)

For $n = 0$ and $n = 1$, we have

$$H^n(X; G) \cong \text{Hom}(H_n(X), G).$$

Proof. It's already been shown for $n = 0$. For $n = 1$, notice that $H_0(X)$ is free, so the Ext term vanishes. □

Example 34.5.4 (Cohomology groups of Klein bottle)

This example will actually have Ext term. Recall that if K is a Klein Bottle then its homology groups are $\mathbb{Z}$ in dimension $n = 0$ and $\mathbb{Z} \oplus \mathbb{Z}_2$ in $n = 1$, and 0 elsewhere.

For $n = 0$, we again just have $H^0(K; G) \cong \text{Hom}(\mathbb{Z}, G) \cong G$. For $n = 1$, the Ext term is $\text{Ext}(H_0(K), G) \cong \text{Ext}(\mathbb{Z}, G) = 0$ so

$$H^1(K; G) \cong \text{Hom}(\mathbb{Z} \oplus \mathbb{Z}_2, G) \cong G \oplus \text{Hom}(\mathbb{Z}_2, G).$$

We have that $\text{Hom}(\mathbb{Z}_2, G)$ is the subgroup of elements of order 2 in G (and $0 \in G$).

But for $n = 2$, we have our first interesting Ext group: the exact sequence is

$$0 \rightarrow \text{Ext}(\mathbb{Z} \oplus \mathbb{Z}_2, G) \rightarrow H^2(X; G) \rightarrow \underbrace{H_2(X)}_{=0} \rightarrow 0.$$

Thus, we have

$$H^2(X; G) \cong (\text{Ext}(\mathbb{Z}, G) \oplus \text{Ext}(\mathbb{Z}_2, G)) \oplus 0 \cong G/2G.$$

All the higher groups vanish. In summary:

$$H^n(X; G) \cong \begin{cases} G & n = 0 \\ G \oplus \text{Hom}(\mathbb{Z}_2, G) & n = 1 \\ G/2G & n = 2 \\ 0 & n \geq 3. \end{cases}$$

§34.6 Relative cohomology groups

One can also define relative cohomology groups in the obvious way: dualize the chain complex

$$\dots \xrightarrow{\partial} C_1(X, A) \xrightarrow{\partial} C_0(X, A) \rightarrow 0$$

to obtain a cochain complex

$$\dots \xleftarrow{\delta} C^1(X, A; G) \xleftarrow{\delta} C^0(X, A; G) \leftarrow 0.$$

We can take the cohomology groups of this.

Definition 34.6.1. The groups thus obtained are the **relative cohomology groups** are denoted $H^n(X, A; G)$.

In addition, we can define reduced cohomology groups as well. One way to do it is to take the augmented singular chain complex

$$\dots \xrightarrow{\partial} C_1(X) \xrightarrow{\partial} C_0(X) \xrightarrow{\epsilon} \mathbb{Z} \rightarrow 0$$

and dualize it to obtain

$$\dots \xleftarrow{\delta} C^1(X; G) \xleftarrow{\delta} C^0(X; G) \xleftarrow{\epsilon^\vee} \underbrace{\text{Hom}(\mathbb{Z}, G)}_{\cong G} \leftarrow 0.$$

Since the $\mathbb{Z}$ we add is also free, the universal coefficient theorem still applies. So this will give us reduced cohomology groups.

However, since we already defined the relative cohomology groups, it is easiest to simply define:

Definition 34.6.2. The **reduced cohomology groups** of a nonempty space X , denoted $\tilde{H}^n(X; G)$, are defined to be $H^n(X, \{*\}; G)$ for some point $*$ in X .

§34.7 Problems to think about

Problem 34A* (Wedge product cohomology). For any G and n we have

$$\tilde{H}^n(X \vee Y; G) \cong \tilde{H}^n(X; G) \oplus \tilde{H}^n(Y; G).$$

Problem 34B[†]. Prove that for a field F of characteristic zero and a space X with finitely generated homology groups:

$$H^k(X, F) \cong (H_k(X))^{\vee}.$$

Thus over fields cohomology is the dual of homology.

Problem 34C ($\mathbb{Z}_2$ -cohomology of $\mathbb{RP}^n$). Prove that

$$H^m(\mathbb{RP}^n, \mathbb{Z}_2) \cong \begin{cases} \mathbb{Z} & m = 0, \text{ or } m \text{ is odd and } m = n \\ \mathbb{Z}_2 & 0 < m < n \text{ and } m \text{ is odd} \\ 0 & \text{otherwise.} \end{cases}$$

35 Application of cohomology

In this final chapter on topology, I'll state (mostly without proof) some nice properties of cohomology groups, and in particular introduce the so-called cup product. For an actual treatise on the cup product, see [Ha02] or [Ma13a].

§35.1 Poincaré duality

First cool result: you may have noticed symmetry in the (co)homology groups of “nice” spaces like the torus or S^n . In fact this is predicted by:

Theorem 35.1.1 (Poincaré duality)

If M is a smooth oriented compact n -manifold, then we have a natural isomorphism

$$H^k(M; \mathbb{Z}) \cong H_{n-k}(M)$$

for every k . In particular, $H^k(M) = 0$ for $k > n$.

So for smooth oriented compact manifolds, cohomology and homology groups are not so different.

From this follows the symmetry that we mentioned when we first defined the Betti numbers:

Corollary 35.1.2 (Symmetry of Betti numbers)

Let M be a smooth oriented compact n -manifold, and let b_k denote its Betti number. Then

$$b_k = b_{n-k}.$$

Proof. Problem 35A[†]. □

§35.2 de Rham cohomology

We now reveal the connection between differential forms and singular cohomology.

Let M be a smooth manifold. We are interested in the homology and cohomology groups of M . We specialize to the case $G = \mathbb{R}$, the additive group of real numbers.

Question 35.2.1. Check that $\text{Ext}(H, \mathbb{R}) = 0$ for any finitely generated abelian group H .

Thus, with real coefficients the universal coefficient theorem says that

$$H^k(M; \mathbb{R}) \cong \text{Hom}(H_k(M), \mathbb{R}) = (H_k(M))^\vee$$

where we view $H_k(X)$ as a real vector space. So, we'd like to get a handle on either $H_k(M)$ or $H^k(M; \mathbb{R})$.

Consider the cochain complex

$$0 \rightarrow \Omega^0(M) \xrightarrow{d} \Omega^1(M) \xrightarrow{d} \Omega^2(M) \xrightarrow{d} \Omega^3(M) \xrightarrow{d} \dots$$

and let $H_{\text{dR}}^k(M)$ denote its cohomology groups. Thus the de Rham cohomology is the closed forms modulo the exact forms.

$$\text{Cochain} : \text{Cocycle} : \text{Coboundary} = k\text{-form} : \text{Closed form} : \text{Exact form}.$$

The whole punch line is:

Theorem 35.2.2 (de Rham's theorem)

For any smooth manifold M , we have a natural isomorphism

$$H^k(M; \mathbb{R}) \cong H_{\text{dR}}^k(M).$$

So the theorem is that the real cohomology groups of manifolds M are actually just given by the behavior of differential forms. Thus,

One can metaphorically think of elements of cohomology groups as G -valued differential forms on the space.

Why does this happen? In fact, we observed already behavior of differential forms which reflects holes in the space. For example, let $M = S^1$ be a circle and consider the **angle form** α (see Example 27.5.4). The form α is closed, but not exact, because it is possible to run a full circle around S^1 . So the failure of α to be exact is signaling that $H_1(S^1) \cong \mathbb{Z}$.

§35.3 Graded rings

Prototypical example for this section: Polynomial rings are commutative graded rings, while $\Lambda^\bullet(V)$ is anticommutative.

In the de Rham cohomology, the differential forms can interact in another way: given a k -form α and an ℓ -form β , we can consider a $(k + \ell)$ -form

$$\alpha \wedge \beta.$$

So we can equip the set of forms with a “product”, satisfying $\beta \wedge \alpha = (-1)^{k\ell} \alpha \wedge \beta$. This is a special case of a more general structure:

Definition 35.3.1. A **graded pseudo-ring** R is an abelian group

$$R = \bigoplus_{d \geq 0} R^d$$

where $R^0, R^1, \dots$, are abelian groups, with an additional associative binary operation $\times : R \rightarrow R$. We require that if $r \in R^d$ and $s \in R^e$, we have $rs \in R^{d+e}$. Elements of an R^d are called **homogeneous elements**; if $r \in R^d$ and $r \neq 0$, we write $|r| = d$.

Note that we do *not* assume commutativity. In fact, these “rings” may not even have an identity 1. We use other words if there are additional properties:

Definition 35.3.2. A **graded ring** is a graded pseudo-ring with 1. If it is commutative we say it is a **commutative graded ring**.

Definition 35.3.3. A graded (pseudo-)ring R is **anticommutative** if for any homogeneous r and s we have

$$rs = (-1)^{|r||s|}sr.$$

To summarize:

Flavors of graded rings	Need not have 1	Must have a 1
No Assumption	graded pseudo-ring	graded ring
Anticommutative	anticommutative pseudo-ring	anticommutative ring
Commutative		commutative graded ring

Example 35.3.4 (Examples of graded rings)

- (a) The ring $R = \mathbb{Z}[x]$ is a **commutative graded ring**, with the d th component being the multiples of x^d .
- (b) The ring $R = \mathbb{Z}[x, y, z]$ is a **commutative graded ring**, with the d th component being the abelian group of homogeneous degree d polynomials (and 0).
- (c) Let V be a vector space, and consider the abelian group

$$\Lambda^\bullet(V) = \bigoplus_{d \geq 0} \Lambda^d(V).$$

For example, $e_1 + (e_2 \wedge e_3) \in \Lambda^\bullet(V)$, say. We endow $\Lambda^\bullet(V)$ with the product $\wedge$, which makes it into an **anticommutative ring**.

- (d) Consider the set of differential forms of a manifold M , say

$$\Omega^\bullet(M) = \bigoplus_{d \geq 0} \Omega^d(M)$$

endowed with the product $\wedge$. This is an **anticommutative ring**.

All four examples have a multiplicative identity.

Let's return to the situation of $\Omega^\bullet(M)$. Consider again the de Rham cohomology groups $H_{\text{dR}}^k(M)$, whose elements are closed forms modulo exact forms. We claim that:

Lemma 35.3.5 (Wedge product respects de Rham cohomology)

The wedge product induces a map

$$\wedge : H_{\text{dR}}^k(M) \times H_{\text{dR}}^\ell(M) \rightarrow H_{\text{dR}}^{k+\ell}(M).$$

Proof. First, we recall that the operator d satisfies

$$d(\alpha \wedge \beta) = (d\alpha) \wedge \beta + \alpha \wedge (d\beta).$$

Now suppose α and β are closed forms. Then from the above, $\alpha \wedge \beta$ is clearly closed. Also if α is closed and $\beta = d\omega$ is exact, then $\alpha \wedge \beta$ is exact, from the identity

$$d(\alpha \wedge \omega) = d\alpha \wedge \omega + \alpha \wedge d\omega = \alpha \wedge \beta.$$

Similarly if α is exact and β is closed then $\alpha \wedge \beta$ is exact. Thus it makes sense to take the product modulo exact forms, giving the theorem above. $\square$

Therefore, we can obtain a *anticommutative ring*

$$H_{\text{dR}}^\bullet(M) = \bigoplus_{k \geq 0} H_{\text{dR}}^k(M)$$

with $\wedge$ as a product, and $1 \in \Lambda^0(\mathbb{R}) = \mathbb{R}$ as the identity

§35.4 Cup products

Inspired by this, we want to see if we can construct a similar product on $\bigoplus_{k \geq 0} H^k(X; R)$ for any topological space X and ring R (where R is commutative with 1 as always). The way to do this is via the *cup product*.

Then this gives us a way to multiply two cochains, as follows.

Definition 35.4.1. Suppose $\phi \in C^k(X; R)$ and $\psi \in C^\ell(X; R)$. Then we can define their **cup product** $\phi \smile \psi \in C^{k+\ell}(X; R)$ to be

$$(\phi \smile \psi)([v_0, \dots, v_{k+\ell}]) = \phi([v_0, \dots, v_k]) \cdot \psi([v_k, \dots, v_{k+\ell}])$$

where the multiplication is in R .

Question 35.4.2. Assuming R has a 1, which 0-cochain is the identity for $\smile$?

First, we prove an analogous result as before:

Lemma 35.4.3 (δ with cup products)

We have $\delta(\phi \smile \psi) = \delta\phi \smile \psi + (-1)^k \phi \smile \delta\psi$.

Proof. Direct $\sum$ computations. $\square$

Thus, by the same routine we used for de Rham cohomology, we get an induced map

$$\smile: H^k(X; R) \times H^\ell(X; R) \rightarrow H^{k+\ell}(X; R).$$

We then define the **singular cohomology ring** whose elements are finite sums in

$$H^\bullet(X; R) = \bigoplus_{k \geq 0} H^k(X; R)$$

and with multiplication given by $\smile$. Thus it is a graded ring (with $1_R \in R$ the identity) and is in fact anticommutative:

Proposition 35.4.4 (Cohomology is anticommutative)

$H^\bullet(X; R)$ is an anticommutative ring, meaning $\phi \smile \psi = (-1)^{k\ell} \psi \smile \phi$.

For a proof, see [Ha02, Theorem 3.11, pages 210-212]. Moreover, we have the de Rham isomorphism

Theorem 35.4.5 (de Rham extends to ring isomorphism)

For any smooth manifold M , the isomorphism of de Rham cohomology groups to singular cohomology groups in fact gives an isomorphism

$$H^\bullet(M; \mathbb{R}) \cong H_{\text{dR}}^\bullet(M)$$

of anticommutative rings.

Therefore, if “differential forms” are the way to visualize the elements of a cohomology group, the wedge product is the correct way to visualize the cup product.

We now present (mostly without proof) the cohomology rings of some common spaces.

Example 35.4.6 (Cohomology of torus)

The cohomology ring $H^\bullet(S^1 \times S^1; \mathbb{Z})$ of the torus is generated by elements $|\alpha| = |\beta| = 1$ which satisfy the relations $\alpha \smile \alpha = \beta \smile \beta = 0$, and $\alpha \smile \beta = -\beta \smile \alpha$. (It also includes an identity 1.) Thus as a $\mathbb{Z}$ -module it is

$$H^\bullet(S^1 \times S^1; \mathbb{Z}) \cong \mathbb{Z} \oplus [\alpha\mathbb{Z} \oplus \beta\mathbb{Z}] \oplus (\alpha \smile \beta)\mathbb{Z}.$$

This gives the expected dimensions $1 + 2 + 1 = 4$. It is anti-commutative.

Example 35.4.7 (Cohomology ring of S^n)

Consider S^n for $n \geq 1$. The nontrivial cohomology groups are given by $H^0(S^n; \mathbb{Z}) \cong H^n(S^n; \mathbb{Z}) \cong \mathbb{Z}$. So as an abelian group

$$H^\bullet(S^n; \mathbb{Z}) \cong \mathbb{Z} \oplus \alpha\mathbb{Z}$$

where α is the generator of $H^n(S^n, \mathbb{Z})$.

Now, observe that $|\alpha \smile \alpha| = 2n$, but since $H^{2n}(S^n; \mathbb{Z}) = 0$ we must have $\alpha \smile \alpha = 0$. So even more succinctly,

$$H^\bullet(S^n; \mathbb{Z}) \cong \mathbb{Z}[\alpha]/(\alpha^2).$$

Confusingly enough, this graded ring is both commutative *and* anti-commutative. The reason is that $\alpha \smile \alpha = 0 = -(\alpha \smile \alpha)$.

Example 35.4.8 (Cohomology ring of real and complex projective space)

It turns out that

$$\begin{aligned} H^\bullet(\mathbb{RP}^n; \mathbb{Z}_2) &\cong \mathbb{Z}_2[\alpha]/(\alpha^{n+1}) \\ H^\bullet(\mathbb{CP}^n; \mathbb{Z}) &\cong \mathbb{Z}[\beta]/(\beta^{n+1}) \end{aligned}$$

where $|\alpha| = 1$ is a generator of $H^1(\mathbb{RP}^n; \mathbb{Z}_2)$ and $|\beta| = 2$ is a generator of $H^2(\mathbb{CP}^n; \mathbb{Z})$.

Confusingly enough, both graded rings are commutative *and* anti-commutative. In the first case it is because we work in $\mathbb{Z}_2$, for which $1 = -1$, so anticommutative is actually equivalent to commutative. In the second case, all nonzero homogeneous elements have degree 2.

§35.5 Relative cohomology pseudo-rings

For $A \subseteq X$, one can also define a relative cup product

$$H^k(X, A; R) \times H^\ell(X, A; R) \rightarrow H^{k+\ell}(X, A; R).$$

After all, if either cochain vanishes on chains in A , then so does their cup product. This lets us define **relative cohomology pseudo-ring** and **reduced cohomology pseudo-ring** (by $A = \{*\}$), say

$$H^\bullet(X, A; R) = \bigoplus_{k \geq 0} H^k(X, A; R)$$

$$\tilde{H}^\bullet(X; R) = \bigoplus_{k \geq 0} \tilde{H}^k(X; R).$$

These are both **anticommutative pseudo-rings**. Indeed, often we have $\tilde{H}^0(X; R) = 0$ and thus there is no identity at all.

Once again we have functoriality:

Theorem 35.5.1 (Cohomology (pseudo-)rings are functorial)

Fix a ring R (commutative with 1). Then we have functors

$$H^\bullet(-; R) : \mathbf{hTop}^{\text{op}} \rightarrow \mathbf{GradedRings}$$

$$H^\bullet(-, -; R) : \mathbf{hPairTop}^{\text{op}} \rightarrow \mathbf{GradedPseudoRings}.$$

Unfortunately, unlike with (co)homology groups, it is a nontrivial task to determine the cup product for even nice spaces like CW complexes. So we will not do much in the way of computation. However, there is a little progress we can make.

§35.6 Wedge sums

Our goal is to now compute $\tilde{H}^\bullet(X \wedge Y)$. To do this, we need to define the product of two graded pseudo-rings:

Definition 35.6.1. Let R and S be two graded pseudo-rings. The **product pseudo-ring** $R \times S$ is the graded pseudo-ring defined by taking the underlying abelian group as

$$R \oplus S = \bigoplus_{d \geq 0} (R^d \oplus S^d).$$

Multiplication comes from R and S , followed by declaring $r \cdot s = 0$ for $r \in R, s \in S$.

Note that this is just graded version of the product ring defined in Example 12.2.10.

Exercise 35.6.2. Show that if R and S are graded rings (meaning they have 1_R and 1_S), then so is $R \times S$.

Now, the theorem is that:

Theorem 35.6.3 (Cohomology pseudo-rings of wedge sums)

We have

$$\tilde{H}^\bullet(X \wedge Y; R) \cong \tilde{H}^\bullet(X; R) \times \tilde{H}^\bullet(Y; R)$$

as graded pseudo-rings.

This allows us to resolve the first question posed at the beginning. Let $X = \mathbb{CP}^2$ and $Y = S^2 \vee S^4$. We have that

$$H^\bullet(\mathbb{CP}^2; \mathbb{Z}) \cong \mathbb{Z}[\alpha]/(\alpha^3).$$

Hence this is a graded ring generated by three elements:

- 1, in dimension 0.
- α , in dimension 2.
- α^2 , in dimension 4.

Next, consider the reduced cohomology pseudo-ring

$$\tilde{H}^\bullet(S^2 \vee S^4; \mathbb{Z}) \cong \tilde{H}^\bullet(S^2; \mathbb{Z}) \oplus \tilde{H}^\bullet(S^4; \mathbb{Z}).$$

Thus the absolute cohomology ring $H^\bullet(S^2 \vee S^4; \mathbb{Z})$ is a graded ring also generated by three elements.

- 1, in dimension 0 (once we add back in the 0th dimension).
- a_2 , in dimension 2 (from $H^\bullet(S^2; \mathbb{Z})$).
- a_4 , in dimension 4 (from $H^\bullet(S^4; \mathbb{Z})$).

Each graded component is isomorphic, like we expected. However, in the former, the product of two degree 2 generators is

$$\alpha \cdot \alpha = \alpha^2.$$

In the latter, the product of two degree 2 generators is

$$a_2 \cdot a_2 = a_2^2 = 0$$

since $a_2 \smile a_2 = 0 \in H^\bullet(S^2; \mathbb{Z})$.

Thus $S^2 \vee S^4$ and $\mathbb{CP}^2$ are not homotopy equivalent.

§35.7 Künneth formula

We now wish to tell apart the spaces $S^2 \times S^4$ and $\mathbb{CP}^3$. In order to do this, we will need a formula for $H^n(X \times Y; R)$ in terms of $H^n(X; R)$ and $H^n(Y; R)$. Thus formulas are called **Künneth formulas**. In this section we will only use a very special case, which involves the tensor product of two graded rings.

Definition 35.7.1. Let A and B be two graded rings which are also R -modules (where R is a commutative ring with 1). We define the **tensor product** $A \otimes_R B$ as follows. As an abelian group, it is

$$A \otimes_R B = \bigoplus_{d \geq 0} \left(\bigoplus_{k=0}^d A^k \otimes_R B^{d-k} \right).$$

The multiplication is given on basis elements by

$$(a_1 \otimes b_1)(a_2 \otimes b_2) = (a_1 a_2) \otimes (b_1 b_2).$$

Of course the multiplicative identity is $1_A \otimes 1_B$.

Now let X and Y be topological spaces, and take the product: we have a diagram

$$\begin{array}{ccc} & X \times Y & \\ \pi_X \swarrow & & \searrow \pi_Y \\ X & & Y \end{array}$$

where π_X and π_Y are projections. As $H^k(-; R)$ is functorial, this gives induced maps

$$\begin{aligned} \pi_X^* : H^k(X \times Y; R) &\rightarrow H^k(X; R) \\ \pi_Y^* : H^k(X \times Y; R) &\rightarrow H^k(Y; R) \end{aligned}$$

for every k .

By using this, we can define a so-called cross product.

Definition 35.7.2. Let R be a ring, and X and Y spaces. Let π_X and π_Y be the projections of $X \times Y$ onto X and Y . Then the **cross product** is the map

$$H^\bullet(X; R) \otimes_R H^\bullet(Y; R) \xrightarrow{\times} H^\bullet(X \times Y; R)$$

acting on cocycles as follows: $\phi \times \psi = \pi_X^*(\phi) \smile \pi_Y^*(\psi)$.

This is just the most natural way to take a k -cycle on X and an ℓ -cycle on Y , and create a $(k + \ell)$ -cycle on the product space $X \times Y$.

Theorem 35.7.3 (Künneth formula)

Let X and Y be CW complexes such that $H^k(Y; R)$ is a finitely generated free R -module for every k . Then the cross product is an isomorphism of anticommutative rings

$$H^\bullet(X; R) \otimes_R H^\bullet(Y; R) \rightarrow H^\bullet(X \times Y; R).$$

In any case, this finally lets us resolve the question set out at the beginning. We saw that $H_n(\mathbb{CP}^3) \cong H_n(S^2 \times S^4)$ for every n , and thus it follows that $H^n(\mathbb{CP}^3; \mathbb{Z}) \cong H^n(S^2 \times S^4; \mathbb{Z})$ too.

But now let us look at the cohomology rings. First, we have

$$H^\bullet(\mathbb{CP}^3; \mathbb{Z}) \cong \mathbb{Z}[\alpha]/(\alpha^3) \cong \mathbb{Z} \oplus \alpha\mathbb{Z} \oplus \alpha^2\mathbb{Z} \oplus \alpha^3\mathbb{Z}$$

where $|\alpha| = 2$; hence this is a graded ring generated by

- 1, in degree 0.

- α , in degree 2.
- α^2 , in degree 4.
- α^3 , in degree 6.

Now let's analyze

$$H^\bullet(S^2 \times S^4; \mathbb{Z}) \cong \mathbb{Z}[\beta]/(\beta^2) \otimes \mathbb{Z}[\gamma]/(\gamma^2).$$

It is thus generated thus by the following elements:

- $1 \otimes 1$, in degree 0.
- $\beta \otimes 1$, in degree 2.
- $1 \otimes \gamma$, in degree 4.
- $\beta \otimes \gamma$, in degree 6.

Again in each dimension we have the same abelian group. But notice that if we square $\beta \otimes 1$ we get

$$(\beta \otimes 1)(\beta \otimes 1) = \beta^2 \otimes 1 = 0.$$

Yet the degree 2 generator of $H^\bullet(\mathbb{CP}^3; \mathbb{Z})$ does not have this property. Hence these two graded rings are not isomorphic.

So it follows that $\mathbb{CP}^3$ and $S^2 \times S^4$ are not homotopy equivalent.

§35.8 Problems to think about

Problem 35A[†] (Symmetry of Betti numbers by Poincaré duality). Let M be a smooth oriented compact n -manifold, and let b_k denote its Betti number. Prove that $b_k = b_{n-k}$.

Problem 35B. Show that $\mathbb{RP}^n$ is not orientable for even n .

Problem 35C. Show that $\mathbb{RP}^3$ is not homotopy equivalent to $\mathbb{RP}^2 \wedge S^3$.



Problem 35D. Show that $S^m \wedge S^n$ is not a deformation retract of $S^m \times S^n$ for any $m, n \geq 1$.

X

Algebraic NT I: Rings of Integers

36 Algebraic integers	342
36.1 Motivation from high school algebra	342
36.2 Algebraic numbers and algebraic integers	343
36.3 Number fields	344
36.4 Norms and traces	345
36.5 The ring of integers	348
36.6 Primitive element theorem, and monogenic extensions	351
36.7 Problems to think about	352
37 Unique factorization (finally!)	354
37.1 Motivation	354
37.2 Ideal arithmetic	355
37.3 Field of fractions	356
37.4 Dedekind domains	356
37.5 Unique factorization works	357
37.6 The factoring algorithm	359
37.7 Fractional ideals	361
37.8 The ideal norm	363
37.9 Problems to think about	364
38 Minkowski bound and class groups	365
38.1 The class group	365
38.2 The discriminant of a number field	365
38.3 The signature of a number field	368
38.4 Minkowski's theorem	370
38.5 The trap box	371
38.6 The Minkowski bound	372
38.7 The class group is finite	373
38.8 Computation of class numbers	374
38.9 Problems to think about	377
39 More properties of the discriminant	378
39.1 Problems to think about	378
40 Bonus: Let's solve Pell's equation!	379
40.1 Units	379
40.2 Dirichlet's unit theorem	380
40.3 Finding fundamental units	381
40.4 Pell's equation	382
40.5 Problems to think about	383

36 Algebraic integers

Here's a first taste of algebraic number theory.

This is really close to the border between olympiads and higher math. You've always known that $a + \sqrt{2}b$ had a "norm" $a^2 - 2b^2$, and that somehow this norm was multiplicative. You've also always known that roots come in conjugate pairs. You might have heard of minimal polynomials but not know much about them.

This chapter will make all these vague notions precise. It's drawn largely from the first chapter of [Og10].

§36.1 Motivation from high school algebra

This is adapted from my blog, *Power Overwhelming*¹.

In high school precalculus, you'll often be asked to find the roots of some polynomial with integer coefficients. For instance,

$$x^3 - x^2 - x - 15 = (x - 3)(x^2 + 2x + 5)$$

has roots 3, $-1 + 2i$, $-1 - 2i$. Or as another example,

$$x^3 - 3x^2 - 2x + 2 = (x + 1)(x^2 - 4x + 2)$$

has roots -1 , $2 + \sqrt{2}$, $2 - \sqrt{2}$. You'll notice that the irrational roots, like $-1 \pm 2i$ and $2 \pm \sqrt{2}$, are coming up in pairs. In fact, I think precalculus explicitly tells you that the imaginary roots come in conjugate pairs. More generally, it seems like all the roots of the form $a + b\sqrt{c}$ come in "conjugate pairs". And you can see why.

But a polynomial like

$$x^3 - 8x + 4$$

has no rational roots. (The roots of this are approximately -3.0514 , 0.51730 , 2.5341 .) Or even simpler,

$$x^3 - 2$$

has only one real root, $\sqrt[3]{2}$. These roots, even though they are irrational, have no "conjugate" pairs. Or do they?

Let's try and figure out exactly what's happening. Let α be any complex number. We define the **minimal polynomial** of α over $\mathbb{Q}$ to be the polynomial such that

- $P(x)$ has rational coefficients, and leading coefficient 1,
- $P(\alpha) = 0$.
- The degree of P is as small as possible. We call $\deg P$ the **degree** of α .

¹URL: <https://usamo.wordpress.com/2014/10/19/why-do-roots-come-in-conjugate-pairs/>

Example 36.1.1 (Examples of minimal polynomials)

- (a) $\sqrt{2}$ has minimal polynomial $x^2 - 2$.
- (b) The imaginary unit $i = \sqrt{-1}$ has minimal polynomial $x^2 + 1$.
- (c) A primitive p th root of unity, $\zeta_p = e^{\frac{2\pi i}{p}}$, has minimal polynomial $x^{p-1} + x^{p-2} + \cdots + 1$, where p is a prime.

Note that $100x^2 - 200$ is also a polynomial of the same degree which has $\sqrt{2}$ as a root; that's why we want to require the polynomial to be monic. That's also why we choose to work in the rational numbers; that way, we can divide by leading coefficients without worrying if we get non-integers.

Why do we care? The point is as follows: suppose we have another polynomial $A(x)$ such that $A(\alpha) = 0$. Then we claim that $P(x)$ actually divides $A(x)$! That means that all the other roots of P will also be roots of A .

The proof is by contradiction: if not, by polynomial long division we can find a quotient and remainder $Q(x), R(x)$ such that

$$A(x) = Q(x)P(x) + R(x)$$

and $R(x) \neq 0$. Notice that by plugging in $x = \alpha$, we find that $R(\alpha) = 0$. But $\deg R < \deg P$, and $P(x)$ was supposed to be the minimal polynomial. That's impossible!

Let's look at a more concrete example. Consider $A(x) = x^3 - 3x^2 - 2x + 2$ from the beginning. The minimal polynomial of $2 + \sqrt{2}$ is $P(x) = x^2 - 4x + 2$ (why?). Now we know that if $2 + \sqrt{2}$ is a root, then $A(x)$ is divisible by $P(x)$. And that's how we know that if $2 + \sqrt{2}$ is a root of A , then $2 - \sqrt{2}$ must be a root too.

As another example, the minimal polynomial of $\sqrt[3]{2}$ is $x^3 - 2$. So $\sqrt[3]{2}$ actually has **two** conjugates, namely, $\alpha = \sqrt[3]{2}(\cos 120^\circ + i \sin 120^\circ)$ and $\beta = \sqrt[3]{2}(\cos 240^\circ + i \sin 240^\circ)$. Thus any polynomial which vanishes at $\sqrt[3]{2}$ also has α and β as roots!

Question 36.1.2. Let α be a root of the polynomial $P(x)$. Show that $P(x)$ is the minimal polynomial if and only if it is irreducible. (This is tautological: the point is just to realize that “minimal polynomials” and “irreducible polynomials” are the same beasts.)

§36.2 Algebraic numbers and algebraic integers

Prototypical example for this section: $\sqrt{2}$ is an algebraic integer (root of $x^2 - 2$), $\frac{1}{2}$ is an algebraic number but not an algebraic integer (root of $x - \frac{1}{2}$).

Let's now work in much vaster generality. First, let's give names to the new numbers we've discussed above.

Definition 36.2.1. An **algebraic number** is any $\alpha \in \mathbb{C}$ which is the root of *some* polynomial with coefficients in $\mathbb{Q}$. The set of algebraic numbers is denoted $\overline{\mathbb{Q}}$.

Remark 36.2.2. One can equally well say algebraic numbers are those of which are roots of some polynomial with coefficients in $\mathbb{Z}$ (rather than $\mathbb{Q}$), since any polynomial in $\mathbb{Q}[x]$ can be scaled to one in $\mathbb{Z}[x]$.

Definition 36.2.3. Consider an algebraic number α and its minimal polynomial P (which is monic and has rational coefficients). If it turns out the coefficients of P are integers, then we say α is an **algebraic integer**.

The set of algebraic integers is denoted $\overline{\mathbb{Z}}$.

Remark 36.2.4. One can show, using *Gauss's Lemma*, that if α is the root of *any* monic polynomial with integer coefficients, then α is an algebraic integer. So in practice, if I want to prove that $\sqrt{2} + \sqrt{3}$ is an algebraic integer, then I only have to say “the polynomial $(x^2 - 5)^2 - 24$ works” without checking that it's minimal.

Sometimes for clarity, we refer to elements of $\mathbb{Z}$ as **rational integers**.

Example 36.2.5 (Examples of algebraic integers)

The numbers

$$4, i = \sqrt{-1}, \sqrt[3]{2}, \sqrt{2} + \sqrt{3}$$

are all algebraic integers, since they are the roots of the monic polynomials $x - 4$, $x^2 + 1$, $x^3 - 2$ and $(x^2 - 5)^2 - 24$.

The number $\frac{1}{2}$ has minimal polynomial $x - \frac{1}{2}$, so it's an algebraic number but not an algebraic integer. (In fact, the rational root theorem also directly implies that any monic integer polynomial does not have $\frac{1}{2}$ as a root!)

Proposition 36.2.6 (Rational algebraic integers are rational integers)

An algebraic integer is rational if and only if it is a rational integer. In symbols,

$$\overline{\mathbb{Z}} \cap \mathbb{Q} = \mathbb{Z}.$$

Question 36.2.7. (Important) Prove this.

§36.3 Number fields

Prototypical example for this section: $\mathbb{Q}(\sqrt{2})$ is a typical number field.

Given any algebraic number α , we're able to consider fields of the form $\mathbb{Q}(\alpha)$. Let us write down the more full version.

Definition 36.3.1. A **number field** K is a field containing $\mathbb{Q}$ as a subfield which is a *finite-dimensional* $\mathbb{Q}$ -vector space. The **degree** of K is its dimension.

Example 36.3.2 (Prototypical example)

Consider the field

$$K = \mathbb{Q}(\sqrt{2}) = \{a + b\sqrt{2} \mid a, b \in \mathbb{Q}\}.$$

This is a field extension of $\mathbb{Q}$, and has degree 2 (the basis being 1 and $\sqrt{2}$).

Example 36.3.3 (Adjoining an algebraic number)

Let α be the root of some irreducible polynomial $P(x)$ in $\mathbb{Q}$. The field $\mathbb{Q}(\alpha)$ is a field extension as well, and the basis is $1, \alpha, \alpha^2, \dots, \alpha^{m-1}$, where m is the degree of α . In particular, the degree of $\mathbb{Q}(\alpha)$ is just the degree of α .

Example 36.3.4 (Non-examples of number fields)

$\mathbb{R}$ and $\mathbb{C}$ are not number fields since there is no *finite* $\mathbb{Q}$ -basis of them.

§36.4 Norms and traces

Prototypical example for this section: $a + b\sqrt{2}$ as an element of $\mathbb{Q}(\sqrt{2})$ has norm $a^2 - 2b^2$ and trace $2a$.

Remember when you did olympiads and we had like $a^2 + b^2$ was the “norm” of $a + bi$? Cool, let me tell you what’s actually happening.

First, let me make precise the notion of a conjugate.

Definition 36.4.1. Let α be an algebraic number, and let $P(x)$ be its minimal polynomial, of degree m . Then the m roots of P are the (Galois) **conjugates** of α .

It’s worth showing at the moment that there are no repeated conjugates.

Lemma 36.4.2 (Irreducible polynomials have distinct roots)

An irreducible polynomial cannot have a complex double root.

Proof. Let $f(x) \in \mathbb{Q}[x]$ be the irreducible polynomial and assume it has a double root α . **Take the derivative** $f'(x)$. This derivative has three interesting properties.

- The degree of f' is one less than the degree of f .
- The polynomials f and f' are not relatively prime because they share a factor $x - \alpha$.
- The coefficients of f' are also in $\mathbb{Q}$.

Consider $P = \gcd(f, f')$. We must have $P \in \mathbb{Q}[x]$ by Euclidean algorithm. But the first two facts about f' ensure that P is nonconstant and $\deg P < \deg f$. Yet P divides f , contradiction. $\square$

Hence α has exactly as many conjugates as the degree of α .

Now, we would *like* to define the *norm* of an element $N(\alpha)$ as the product of its conjugates. For example, we want $2 + i$ to have norm $(2 + i)(2 - i) = 5$, and in general for $a + bi$ to have norm $a^2 + b^2$. It would be *really cool* if the norm was multiplicative; we already know this is true for complex numbers!

Unfortunately, this doesn’t quite work: consider

$$N(2 + i) = 5 \text{ and } N(2 - i) = 5.$$

But $(2 + i)(2 - i) = 5$, which doesn’t have norm 25 like we want, since 5 is degree 1 and has no conjugates at all. The reason this “bad” thing is happening is that we’re trying to define the norm of an *element*, when we really ought to be defining the norm of an element *with respect to a particular K* .

What I’m driving at is that the norm should have different meanings depending on which field you’re in. If we think of 5 as an element of $\mathbb{Q}$, then its norm is 5. But thought of as an element of $\mathbb{Q}(i)$, its norm really ought to be 25. Let’s make this happen: for K a number field, we will now define $N_{K/\mathbb{Q}}(\alpha)$ to be the norm of α *with respect to K* as follows.

Definition 36.4.3. Let $\alpha \in K$ have degree n , so $\mathbb{Q}(\alpha) \subseteq K$, and set $k = (\deg K)/n$. The **norm** of α is defined as

$$N_{K/\mathbb{Q}}(\alpha) \stackrel{\text{def}}{=} \left(\prod \text{Galois conj of } \alpha \right)^k.$$

The **trace** is defined as

$$\text{Tr}_{K/\mathbb{Q}}(\alpha) \stackrel{\text{def}}{=} k \cdot \left(\sum \text{Galois conj of } \alpha \right).$$

The exponent of k is a “correction factor” that makes the norm of 5 into $5^2 = 25$ when we view 5 as an element of $\mathbb{Q}(i)$ rather than an element of $\mathbb{Q}$. For a “generic” element of K , we expect $k = 1$.

Exercise 36.4.4. Use what you know about nested vector spaces to convince yourself that k is actually an integer.

Example 36.4.5 (Norm of $a + b\sqrt{2}$)

Let $\alpha = a + b\sqrt{2} \in \mathbb{Q}(\sqrt{2})$. If $b \neq 0$, then α and K have the degree 2. Thus the only conjugates of α are $a \pm b\sqrt{2}$, which gives the norm

$$(a + b\sqrt{2})(a - b\sqrt{2}) = a^2 - 2b^2,$$

The trace is $(a - b\sqrt{2}) + (a + b\sqrt{2}) = 2a$.

Nicely, the formula $a^2 - 2b^2$ and $2a$ also works when $b = 0$.

Of importance is:

Proposition 36.4.6 (Norms and traces are rational integers)

If α is an algebraic integer, its norm and trace are rational integers.

Question 36.4.7. Prove it. (Vieta.)

That’s great, but it leaves a question unanswered: why is the norm multiplicative? To do this, I have to give a new definition of norm and trace.

Theorem 36.4.8 (Morally correct definition of norm and trace)

Let K be a number field of degree n , and let $\alpha \in K$. Let $\mu_\alpha : K \rightarrow K$ denote the map

$$x \mapsto \alpha x$$

viewed as a linear map of $\mathbb{Q}$ -vector spaces. Then,

- the norm of α equals the determinant $\det \mu_\alpha$, and
- the trace of α equals the trace $\text{Tr } \mu_\alpha$.

Since the trace and determinant don’t depend on the choice of basis, you can pick whatever basis you want and use whatever definition you got in high school. Fantastic, right?

Example 36.4.9 (Explicit computation of matrices for $a + b\sqrt{2}$)

Let $K = \mathbb{Q}(\sqrt{2})$, and let $1, \sqrt{2}$ be the basis of K . Let

$$\alpha = a + b\sqrt{2}$$

(possibly even $b = 0$), and notice that

$$(a + b\sqrt{2})(x + y\sqrt{2}) = (ax + 2yb) + (bx + ay)\sqrt{2}.$$

We can rewrite this in matrix form as

$$\begin{bmatrix} a & 2b \\ b & a \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} ax + 2yb \\ bx + ay \end{bmatrix}.$$

Consequently, we can interpret μ_α as the matrix

$$\mu_\alpha = \begin{bmatrix} a & 2b \\ b & a \end{bmatrix}.$$

Of course, the matrix will change if we pick a different basis, but the determinant and trace do not: they are always given by

$$\det \mu_\alpha = a^2 - 2b^2 \text{ and } \operatorname{Tr} \mu_\alpha = 2a.$$

This interpretation explains why the same formula should work for $a + b\sqrt{2}$ even in the case $b = 0$.

Proof. I'll prove the result for just the norm; the trace falls out similarly. Set

$$n = \deg \alpha \quad kn = \deg K.$$

The proof is split into two parts, depending on whether or not $k = 1$.

Proof if $k = 1$. Set $n = \deg \alpha = \deg K$. Thus the norm actually *is* the product of the Galois conjugates. Also,

$$\{1, \alpha, \dots, \alpha^{n-1}\}$$

is linearly independent in K , and hence a basis (as $\dim K = n$). Let's use this as the basis for μ_α .

Let

$$x^n + c_{n-1}x^{n-1} + \dots + c_0$$

be the minimal polynomial of α . Thus $\mu_\alpha(1) = \alpha$, $\mu_\alpha(\alpha) = \alpha^2$, and so on, but $\mu_\alpha(\alpha^{n-1}) = -c_{n-1}\alpha^{n-1} - \dots - c_0$. Therefore, μ_α is given by the matrix

$$M = \begin{bmatrix} 0 & 0 & 0 & \dots & 0 & -c_0 \\ 1 & 0 & 0 & \dots & 0 & -c_1 \\ 0 & 1 & 0 & \dots & 0 & -c_2 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & 1 & -c_{n-1} \end{bmatrix}.$$

Thus

$$\det M = (-1)^n c_0$$

and we're done by Vieta's formulas. ■

Proof if $k > 1$. We have nested vector spaces

$$\mathbb{Q} \subseteq \mathbb{Q}(\alpha) \subseteq K.$$

Let $e_1, \dots, e_k$ be a $\mathbb{Q}(\alpha)$ -basis for K (meaning: interpret K as a vector space over $\mathbb{Q}(\alpha)$, and pick that basis). Since $\{1, \alpha, \dots, \alpha^{n-1}\}$ is a $\mathbb{Q}$ basis for $\mathbb{Q}(\alpha)$, the elements

$$\begin{array}{cccc} e_1, & e_1\alpha, & \dots, & e_1\alpha^{n-1} \\ e_2, & e_2\alpha, & \dots, & e_2\alpha^{n-1} \\ \vdots & \vdots & \ddots & \vdots \\ e_k, & e_k\alpha, & \dots, & e_k\alpha^{n-1} \end{array}$$

constitute a $\mathbb{Q}$ -basis of K . Using *this* basis, the map μ_α looks like

$$\underbrace{\begin{bmatrix} M & & & \\ & M & & \\ & & \ddots & \\ & & & M \end{bmatrix}}_{k \text{ times}}$$

where M is the same matrix as above: we just end up with one copy of our old matrix for each e_i . Thus $\det \mu_\alpha = (\det M)^k$, as needed. $\blacksquare$

Question 36.4.10. Verify the result for traces as well. $\square$

From this it follows immediately that

$$N_{K/\mathbb{Q}}(\alpha\beta) = N_{K/\mathbb{Q}}(\alpha) N_{K/\mathbb{Q}}(\beta)$$

because by definition we have

$$\mu_{\alpha\beta} = \mu_\alpha \circ \mu_\beta,$$

and that the determinant is multiplicative. In the same way, the trace is additive.

§36.5 The ring of integers

Prototypical example for this section: If $K = \mathbb{Q}(\sqrt{2})$, then $\mathcal{O}_K = \mathbb{Z}[\sqrt{2}]$. But if $K = \mathbb{Q}(\sqrt{5})$, then $\mathcal{O}_K = \mathbb{Z}[\frac{1+\sqrt{5}}{2}]$.

$\mathbb{Z}$ makes for better number theory than $\mathbb{Q}$. In the same way, focusing on the *algebraic integers* of K gives us some really nice structure, and we'll do that here.

Definition 36.5.1. Given a number field K , we define

$$\mathcal{O}_K \stackrel{\text{def}}{=} K \cap \overline{\mathbb{Z}}$$

to be the **ring of integers** of K ; in other words $\mathcal{O}_K$ consists of the algebraic integers of K .

We do the classical example of a quadratic field now. Before proceeding, I need to write a silly number theory fact.

Exercise 36.5.2. Let a and b be rational numbers, and d a squarefree positive integer.

- If $d \equiv 2, 3 \pmod{4}$, prove that $2a, a^2 - db^2 \in \mathbb{Z}$ if and only if $a, b \in \mathbb{Z}$.
- For $d \equiv 1 \pmod{4}$, prove that $2a, a^2 - db^2 \in \mathbb{Z}$ if and only if $a, b \in \mathbb{Z}$ OR if $a - \frac{1}{2}, b - \frac{1}{2} \in \mathbb{Z}$.

This is annoying but not hard; you'll need to take mod 4.

Example 36.5.3 (Ring of integers of $K = \mathbb{Q}(\sqrt{3})$)

Let K be as above. We claim that

$$\mathcal{O}_K = \mathbb{Z}[\sqrt{3}] = \left\{ m + n\sqrt{3} \mid m, n \in \mathbb{Z} \right\}.$$

We set $\alpha = a + b\sqrt{3}$. Then $\alpha \in \mathcal{O}_K$ when the minimal polynomial has integer coefficients.

If $b = 0$, then the minimal polynomial is $x - \alpha = x - a$, and thus α works if and only if it's an integer. If $b \neq 0$, then the minimal polynomial is

$$(x - a)^2 - 3b^2 = x^2 - 2a \cdot x + (a^2 - 3b^2).$$

From the exercise, this occurs exactly for $a, b \in \mathbb{Z}$.

Example 36.5.4 (Ring of integers of $K = \mathbb{Q}(\sqrt{5})$)

We claim that in this case

$$\mathcal{O}_K = \mathbb{Z} \left[\frac{1 + \sqrt{5}}{2} \right] = \left\{ m + n \cdot \frac{1 + \sqrt{5}}{2} \mid m, n \in \mathbb{Z} \right\}.$$

The proof is exactly the same, except the exercise tells us instead that for $b \neq 0$, we have both the possibility that $a, b \in \mathbb{Z}$ or that $a, b \in \mathbb{Z} - \frac{1}{2}$. This reflects the fact that $\frac{1+\sqrt{5}}{2}$ is the root of $x^2 - x - 1 = 0$; no such thing is possible with $\sqrt{3}$.

In general, the ring of integers of $K = \mathbb{Q}(\sqrt{d})$ is

$$\mathcal{O}_K = \begin{cases} \mathbb{Z}[\sqrt{d}] & d \equiv 2, 3 \pmod{4} \\ \mathbb{Z} \left[\frac{1+\sqrt{d}}{2} \right] & d \equiv 1 \pmod{4}. \end{cases}$$

What we're going to show is that $\mathcal{O}_K$ behaves in K a lot like the integers do in $\mathbb{Q}$. First we show K consists of quotients of numbers in $\mathcal{O}_K$. In fact, we can do better:

Example 36.5.5 (Rationalizing the denominator)

For example, consider $K = \mathbb{Q}(\sqrt{3})$. The number $x = \frac{1}{4+\sqrt{3}}$ is an element of K , but by "rationalizing the denominator" we can write

$$\frac{1}{4 + \sqrt{3}} = \frac{4 - \sqrt{3}}{13}.$$

So we see that in fact, x is $\frac{1}{13}$ of an integer in $\mathcal{O}_K$.

The theorem holds true more generally.

Theorem 36.5.6 ($K = \mathbb{Q} \cdot \mathcal{O}_K$)

Let K be a number field, and let $x \in K$ be any element. Then there exists an integer n such that $nx \in \mathcal{O}_K$; in other words,

$$x = \frac{1}{n}\alpha$$

for some $\alpha \in \mathcal{O}_K$.

Exercise 36.5.7. Prove this yourself. (Start by using the fact that x has a minimal polynomial with integer coefficients. Alternatively, take the norm.)

Now we are going to show $\mathcal{O}_K$ is a ring; we'll check it is closed under addition and multiplication. To do so, the easiest route is:

Lemma 36.5.8 ($\alpha \in \overline{\mathbb{Z}} \iff \mathbb{Z}[\alpha]$ finitely generated)

Let $\alpha \in \overline{\mathbb{Q}}$. Then α is an algebraic integer if and only if the abelian group $\mathbb{Z}[\alpha]$ is finitely generated.

Proof. Note that α is an algebraic integer if and only if it's the root of some nonzero, monic polynomial with integer coefficients. Suppose first that

$$\alpha^N = c_{N-1}\alpha^{N-1} + c_{N-2}\alpha^{N-2} + \cdots + c_0.$$

Then the set $1, \alpha, \dots, \alpha^{N-1}$ generates $\mathbb{Z}[\alpha]$, since we can repeatedly replace α^N until all powers of α are less than N .

Conversely, suppose that $\mathbb{Z}[\alpha]$ is finitely generated by some $b_1, \dots, b_m$. Viewing the b_i as polynomials in α , we can select a large integer N (say $N = \deg b_1 + \cdots + \deg b_m + 2015$) and express α^N in the b_i 's to get

$$\alpha^N = c_1 b_1(\alpha) + \cdots + c_m b_m(\alpha).$$

The above gives us a monic polynomial in α , and the choice of N guarantees it is not zero. So α is an algebraic integer. $\square$

Example 36.5.9 ($\frac{1}{2}$ isn't an algebraic integer)

We already know $\frac{1}{2}$ isn't an algebraic integer. So we expect

$$\mathbb{Z}\left[\frac{1}{2}\right] = \left\{ \frac{a}{2^m} \mid a, m \in \mathbb{Z} \text{ and } m \geq 0 \right\}.$$

to not be finitely generated, and this is the case.

Question 36.5.10. To make the last example concrete: name all the elements of $\mathbb{Z}[\frac{1}{2}]$ that cannot be written as an integer combination of

$$\left\{ \frac{1}{2}, \frac{7}{8}, \frac{13}{64}, \frac{2015}{4096}, \frac{1}{1048576} \right\}$$

Now we can state the theorem.

Theorem 36.5.11 (Algebraic integers are closed under $+$ and $\times$)

The set $\overline{\mathbb{Z}}$ is closed under addition and multiplication; i.e. it is a ring. In particular, $\mathcal{O}_K$ is also a ring for any number field K .

Proof. Let $\alpha, \beta \in \overline{\mathbb{Z}}$. Then $\mathbb{Z}[\alpha]$ and $\mathbb{Z}[\beta]$ are finitely generated. Hence so is $\mathbb{Z}[\alpha, \beta]$. (Details: if $\mathbb{Z}[\alpha]$ has $\mathbb{Z}$ -basis $a_1, \dots, a_m$ and $\mathbb{Z}[\beta]$ has $\mathbb{Z}$ -basis $b_1, \dots, b_n$, then take the mn elements $a_i b_j$.)

Now $\mathbb{Z}[\alpha \pm \beta]$ and $\mathbb{Z}[\alpha\beta]$ are subsets of $\mathbb{Z}[\alpha, \beta]$ and so they are also finitely generated. Hence $\alpha \pm \beta$ and $\alpha\beta$ are algebraic integers. $\square$

In fact, something even better is true. As you saw, for $\mathbb{Q}(\sqrt{3})$ we had $\mathcal{O}_K = \mathbb{Z}[\sqrt{3}]$; in other words, $\mathcal{O}_K$ was generated by 1 and $\sqrt{3}$. Something similar was true for $\mathbb{Q}(\sqrt{5})$. We claim that in fact, the general picture looks exactly like this.

Theorem 36.5.12 ($\mathcal{O}_K$ is a free $\mathbb{Z}$ -module of rank n)

Let K be a number field of degree n . Then $\mathcal{O}_K$ is a free $\mathbb{Z}$ -module of rank n , i.e. $\mathcal{O}_K \cong \mathbb{Z}^{\oplus n}$ as an abelian group. In other words, $\mathcal{O}_K$ has a $\mathbb{Z}$ -basis of n elements as

$$\mathcal{O}_K = \{c_1\alpha_1 + \dots + c_{n-1}\alpha_{n-1} + c_n\alpha_n \mid c_i \in \mathbb{Z}\}$$

where α_i are algebraic integers in $\mathcal{O}_K$.

Proof. I find this kind of fun: Problem 36F. $\square$

This last theorem shows that in many ways $\mathcal{O}_K$ is a “lattice” in K . That is, for a number field K we can find $\alpha_1, \dots, \alpha_n$ in $\mathcal{O}_K$ such that

$$\begin{aligned} \mathcal{O}_K &\cong \alpha_1\mathbb{Z} \oplus \alpha_2\mathbb{Z} \oplus \dots \oplus \alpha_n\mathbb{Z} \\ K &\cong \alpha_1\mathbb{Q} \oplus \alpha_2\mathbb{Q} \oplus \dots \oplus \alpha_n\mathbb{Q} \end{aligned}$$

as abelian groups.

§36.6 Primitive element theorem, and monogenic extensions

Prototypical example for this section: $\mathbb{Q}(\sqrt{3}, \sqrt{5}) \cong \mathbb{Q}(\sqrt{3} + \sqrt{5})$. Can you see why?

I’m only putting this theorem here because I was upset that no one told me it was true (it’s a very natural conjecture), and I hope to not do the same to the reader. However, I’m not going to use it in anything that follows.

Theorem 36.6.1 (Artin's primitive element theorem)

Every number field K is isomorphic to $\mathbb{Q}(\alpha)$ for some algebraic number α .

The proof is left as Problem 41E, since to prove it I need to talk about field extensions first.

The prototypical example

$$\mathbb{Q}(\sqrt{3}, \sqrt{5}) \cong \mathbb{Q}(\sqrt{3} + \sqrt{5})$$

makes it clear why this theorem should not be too surprising.

We might hope that something similar is true of the ring of integers: that we can write

$$\mathcal{O}_K = \mathbb{Z}[\theta]$$

in which case $\{1, \theta, \dots, \theta^{n-1}\}$ serves both as a basis of K and as the $\mathbb{Z}$ -basis for $\mathcal{O}_K$ (here $n = [K : \mathbb{Q}]$). In other words, we hope that the basis of $\mathcal{O}_K$ is actually a “power basis”.

This is true for the most common examples we use:

- the quadratic field, and
- the cyclotomic field in Problem 36G[†].

Unfortunately, it is not true in general: the first counterexample is $\mathbb{Q}(\alpha)$ for α a root of $X^3 - X^2 - 2X - 8$.

We call an extension with this nice property **monogenic**. As we'll later see, monogenic extensions have a really nice factoring algorithm, Theorem 37.6.4.

§36.7 Problems to think about

Problem 36A^{*}. Show that α is a unit of $\mathcal{O}_K$ (meaning $\alpha^{-1} \in \mathcal{O}_K$) if and only if $N_{K/\mathbb{Q}}(\alpha) = \pm 1$.

Problem 36B (Russia 1984). Find all integers m and n such that

$$(5 + 3\sqrt{2})^m = (3 + 5\sqrt{2})^n.$$

Problem 36C (Black MOP 2010). There are $n > 2$ lamps arranged in a circle; initially one is on and the others are off. We may select any regular polygon whose vertices are among the lamps and toggle the states of all the lamps simultaneously. Show it is impossible to turn all lamps off.

Problem 36D (USA TST 2012). Decide whether there exist $a, b, c > 2010$ satisfying

$$a^3 + 2b^3 + 4c^3 = 6abc + 1.$$

Problem 36E^{*}. Consider n roots of unity $\varepsilon_1, \dots, \varepsilon_n$. Assume the average $\frac{1}{n}(\varepsilon_1 + \dots + \varepsilon_n)$ is an algebraic integer. Prove that either the average is zero or $\varepsilon_1 = \dots = \varepsilon_n$. (Used in Lemma 49.2.2.)



Problem 36F. Let K be a number field of degree n . Using the norm, show that $\mathcal{O}_K$ is a free $\mathbb{Z}$ -module of rank n (thus proving Theorem 36.5.12).



Problem 36G[†] (Cyclotomic Field). Let p be an odd rational prime and ζ_p a primitive p th root of unity. Let $K = \mathbb{Q}(\zeta_p)$. Prove that $\mathcal{O}_K = \mathbb{Z}[\zeta_p]$. (In fact, the result is true even if p is not a prime.)



Problem 36H (Kronecker's theorem). Let α be an algebraic integer. Suppose all its Galois conjugates have absolute value one. Prove that $\alpha^N = 1$ for some positive integer N .

DRAFT (Evan Chen)
Updated February 15, 2017

37 Unique factorization (finally!)

Took long enough.

§37.1 Motivation

Suppose we're interested in solutions to the Diophantine equation $n = x^2 + 5y^2$ for a given n . The idea is to try and “factor” n in $\mathbb{Z}[\sqrt{-5}]$, for example

$$6 = (1 + \sqrt{-5})(1 - \sqrt{-5}).$$

Unfortunately, this is not so simple, because as I've said before we don't have unique factorization of elements:

$$6 = 2 \cdot 3 = (1 + \sqrt{-5})(1 - \sqrt{-5}).$$

One reason this doesn't work is that we don't have a notion of a *greatest common divisor*. We can write $(35, 77) = 7$, but what do we make of $(3, 1 + \sqrt{-5})$?

The trick is to use ideals as a “generalized GCD”. Recall that by (a, b) I mean the ideal $\{ax + by \mid x, y \in \mathbb{Z}[\sqrt{-5}]\}$. You can see that $(35, 77) = (7)$, but $(3, 1 + \sqrt{-5})$ will be left “unsimplified” because it doesn't represent an actual value in the ring. Using these *sets* (ideals) as elements, it turns out that we can develop a full theory of prime factorization, and we do so in this chapter.

In other words, we use the ideal $(a_1, \dots, a_m)$ to interpret a “generalized GCD” of $a_1, \dots, a_m$. In particular, if we have a number x we want to represent, we encode it as just (x) .

Going back to our example of 6,

$$(6) = (2) \cdot (3) = (1 + \sqrt{-5}) \cdot (1 - \sqrt{-5}).$$

Please take my word for it that in fact, the complete prime factorization of (6) into prime ideals is

$$(6) = (2, 1 - \sqrt{-5})^2 (3, 1 + \sqrt{-5}) (3, 1 - \sqrt{-5}) = \mathfrak{p}^2 \mathfrak{q}_1 \mathfrak{q}_2.$$

In fact, $(2) = \mathfrak{p}^2$, $(3) = \mathfrak{q}_1 \mathfrak{q}_2$, $(1 + \sqrt{-5}) = \mathfrak{p} \mathfrak{q}_1$, $(1 - \sqrt{-5}) = \mathfrak{p} \mathfrak{q}_2$. So 6 indeed factorizes uniquely into ideals, even though it doesn't factor into elements.

As one can see above, ideal factorization is more refined than element factorization. Once you have the factorization into *ideals*, you can from there recover all the factorizations into *elements*. The upshot of this is that if we want to write n as $x^2 + 5y^2$, we just have to factor n into ideals, and from there we can recover all factorizations into elements, and finally all ways to write n as $x^2 + 5y^2$. Since we can already break n into rational prime factors (for example $6 = 2 \cdot 3$ above) we just have to figure out how each rational prime $p \mid n$ breaks down. There's a recipe for this, Theorem 37.6.4! In fact, I'll even tell you what it says in this special case:

- If $t^2 + 5$ factors as $(t + c)(t - c) \pmod{p}$, then $(p) = (p, c + \sqrt{-5})(p, c - \sqrt{-5})$.
- Otherwise, (p) is a prime ideal.

In this chapter we'll develop this theory of unique factorization in full generality.

Remark 37.1.1. In this chapter, I'll be using the letters $\mathfrak{a}$, $\mathfrak{b}$, $\mathfrak{p}$, $\mathfrak{q}$ for ideals of $\mathcal{O}_K$. When fractional ideals arise, I'll use I and J for them.

§37.2 Ideal arithmetic

Prototypical example for this section: $(x)(y) = (xy)$. In any case, think in terms of generators.

First, I have to tell you how to add and multiply two ideals $\mathfrak{a}$ and $\mathfrak{b}$.

Definition 37.2.1. Given two ideals $\mathfrak{a}$ and $\mathfrak{b}$ of a ring R , we define

$$\begin{aligned}\mathfrak{a} + \mathfrak{b} &\stackrel{\text{def}}{=} \{a + b \mid a \in \mathfrak{a}, b \in \mathfrak{b}\} \\ \mathfrak{a} \cdot \mathfrak{b} &\stackrel{\text{def}}{=} \{a_1 b_1 + \cdots + a_n b_n \mid a_i \in \mathfrak{a}, b_i \in \mathfrak{b}\}.\end{aligned}$$

(Note that infinite sums don't make sense in general rings, which is why in $\mathfrak{a} \cdot \mathfrak{b}$ we cut off the sum after some finite number of terms.) You can readily check these are actually ideals. This definition is more natural if you think about it in terms of the generators of $\mathfrak{a}$ and $\mathfrak{b}$.

Proposition 37.2.2 (Ideal arithmetic via generators)

Suppose $\mathfrak{a} = (a_1, a_2, \dots, a_n)$ and $\mathfrak{b} = (b_1, \dots, b_m)$ are ideals in a ring R . Then

- (a) $\mathfrak{a} + \mathfrak{b}$ is the ideal generated by $a_1, \dots, a_n, b_1, \dots, b_m$.
- (b) $\mathfrak{a} \cdot \mathfrak{b}$ is the ideal generated by $a_i b_j$, for $1 \leq i \leq n$ and $1 \leq j \leq m$.

Proof. Pretty straightforward; just convince yourself that this result is correct. $\square$

In other words, for sums you append the two sets of generators together, and for products you take products of the generators. Note that for principal ideals, this coincides with “normal” multiplication, for example

$$(3) \cdot (5) = (15)$$

in $\mathbb{Z}$.

Remark 37.2.3. Note that for an ideal $\mathfrak{a}$ and an element c , the set

$$c\mathfrak{a} = \{ca \mid a \in \mathfrak{a}\}$$

is equal to $(c) \cdot \mathfrak{a}$. So “scaling” and “multiplying by principal ideals” are the same thing. This is important, since we'll be using the two notions interchangeably.

Finally, since we want to do factorization we better have some notion of divisibility. So we define:

Definition 37.2.4. We say $\mathfrak{a}$ divides $\mathfrak{b}$ and write $\mathfrak{a} \mid \mathfrak{b}$ if $\mathfrak{a} \supseteq \mathfrak{b}$.

Note the reversal of inclusions! So (3) divides (15) , because (15) is contained in (3) ; every multiple of 15 is a multiple of 3.

Finally, the **prime ideals** are defined as in Definition 12.7.2: $\mathfrak{p}$ is prime if $xy \in \mathfrak{p}$ implies $x \in \mathfrak{p}$ or $y \in \mathfrak{p}$. This is compatible with the definition of divisibility:

Exercise 37.2.5. A nonzero proper ideal $\mathfrak{p}$ is prime if and only if whenever $\mathfrak{p}$ divides $\mathfrak{a}\mathfrak{b}$, $\mathfrak{p}$ divides one of $\mathfrak{a}$ or $\mathfrak{b}$.

As mentioned in Remark 12.7.4, this also lets us ignore multiplication by units: $(-3) = (3)$.

§37.3 Field of fractions

Prototypical example for this section: The field of fractions of $\mathbb{Z}$ is $\mathbb{Q}$

As long as we're trying to get an analog of factoring things in $\mathbb{Z}$, we may as well try to extend it to $\mathbb{Q}$: every rational number $\frac{a}{b}$ can be factored by just factoring its numerator and denominator. So we want some analog of $\mathbb{Q}$.

Given an integral domain R , we define its **field of fractions** K as follows: it consists of elements a/b , where $a, b \in R$ and $b \neq 0$. We set $a/b \sim c/d$ if and only if $bc = ad$.

For example, the field of fractions of $\mathbb{Z}$ is $\mathbb{Q}$, *by definition*. In fact everything you know about $\mathbb{Q}$ basically carries over by analogy. In particular, we define

$$\frac{a}{b} + \frac{c}{d} \stackrel{\text{def}}{=} \frac{ad + bc}{bd}$$

and

$$\frac{a}{b} \cdot \frac{c}{d} \stackrel{\text{def}}{=} \frac{ac}{bd}.$$

You can prove if you want that this indeed a field, but considering how comfortable we are that $\mathbb{Q}$ is well-defined, I wouldn't worry about it...

Question 37.3.1. Show that the ring of integers $\mathcal{O}_K$ has field of fractions K .

§37.4 Dedekind domains

Prototypical example for this section: Any $\mathcal{O}_K$ is a Dedekind domain.

We now define a Dedekind domain as follows.

Definition 37.4.1. An integral domain $\mathcal{A}$ is a **Dedekind domain** if it is Noetherian, integrally closed, and *every nonzero prime ideal of $\mathcal{A}$ is in fact maximal*. (The last condition is the important one.)

Here there's one new word I have to define for you, but we won't make much use of it.

Definition 37.4.2. Let R be an integral domain and let K be its field of fractions. We say R is **integrally closed** if the only elements $a \in K$ which are roots of *monic* polynomials in R are the elements of R (which are roots of the trivial $x - r$ polynomial).

The *interesting* condition in the definition of a Dedekind domain is the last one: prime ideals and maximal ideals are the same thing. The other conditions are just technicalities, but "primes are maximal" has real substance.

Example 37.4.3 ($\mathbb{Z}$ is a dedekind domain)

The ring $\mathbb{Z}$ is a Dedekind domain. Note that

- $\mathbb{Z}$ is Noetherian (for obvious reasons).
- $\mathbb{Z}$ has field of fractions $\mathbb{Q}$. If $f(x) \in \mathbb{Z}[x]$ is monic, then by the rational root theorem any rational roots are integers (this is the same as the proof that $\mathbb{Z} \cap \mathbb{Q} = \mathbb{Z}$). Hence $\mathbb{Z}$ is integrally closed.
- The nonzero prime ideals of $\mathbb{Z}$ are (p) , which also happen to be maximal.

The case of interest is a ring $\mathcal{O}_K$ in which we wish to do factorizing. We're now going to show that for any number field K , the ring $\mathcal{O}_K$ is a Dedekind domain. First, the boring part.

Proposition 37.4.4 ($\mathcal{O}_K$ integrally closed and Noetherian)

For any number field K , the ring $\mathcal{O}_K$ is integrally closed and Noetherian.

Proof. Boring, but here it is anyways for completeness.

Since $\mathcal{O}_K \cong \mathbb{Z}^{\oplus n}$, we get that it's Noetherian.

Now we show that $\mathcal{O}_K$ is integrally closed. Suppose that $\eta \in K$ is the root of some polynomial with coefficients in $\mathcal{O}_K$. Thus

$$\eta^n = \alpha_{n-1} \cdot \eta^{n-1} + \alpha_{n-2} \cdot \eta^{n-2} + \cdots + \alpha_0$$

where $\alpha_i \in \mathcal{O}_K$. We want to show that $\eta \in \mathcal{O}_K$ as well.

Well, from the above, $\mathcal{O}_K[\eta]$ is finitely generated... thus $\mathbb{Z}[\eta] \subseteq \mathcal{O}_K[\eta]$ is finitely generated. So $\eta \in \overline{\mathbb{Z}}$, and hence $\eta \in K \cap \overline{\mathbb{Z}} = \mathcal{O}_K$. $\square$

Now let's do the fun part. We'll prove a stronger result, which will re-appear repeatedly.

Theorem 37.4.5 (Important: prime ideals divide rational primes)

Let $\mathcal{O}_K$ be a ring of integers and $\mathfrak{p}$ a nonzero prime ideal inside it. Then $\mathfrak{p}$ contains a rational prime p . Moreover, $\mathfrak{p}$ is maximal.

Proof. Take any $\alpha \neq 0$ in $\mathfrak{p}$, and consider the product of all its conjugates. This is the norm $N(\alpha)$, which is contained in $\mathfrak{p}$. Hence $\mathfrak{p}$ contains some integer; pick the smallest one, say p . We must have that p is a rational prime, since otherwise $\mathfrak{p} \ni p = xy$ implies one of $x, y \in \mathfrak{p}$. This shows the first part.

We now do something pretty tricky to show $\mathfrak{p}$ is maximal. Look at $\mathcal{O}_K/\mathfrak{p}$; since $\mathfrak{p}$ is prime it's supposed to be an integral domain... but we claim that it's actually finite! To do this, we forget that we can multiply on $\mathcal{O}_K$. Recalling that $\mathcal{O}_K \cong \mathbb{Z}^{\oplus n}$ as an abelian group, we obtain a map

$$\mathbb{F}_p^{\oplus n} \cong \mathcal{O}_K/(p) \twoheadrightarrow \mathcal{O}_K/\mathfrak{p}.$$

Hence $|\mathcal{O}_K/\mathfrak{p}| \leq p^n$ is finite. Since finite integral domains are fields (Problem 12A*) we are done. $\square$

Since every nonzero prime $\mathfrak{p}$ is maximal, we now know that $\mathcal{O}_K$ is a Dedekind domain. Note that this tricky proof is essentially inspired by the solution to Problem 12C*.

§37.5 Unique factorization works

Okay, I'll just say it now!

Unique factorization works perfectly in Dedekind domains!

Theorem 37.5.1 (Prime factorization works)

Let $\mathfrak{a}$ be a nonzero proper ideal of a Dedekind domain $\mathcal{A}$. Then $\mathfrak{a}$ can be written as a finite product of nonzero prime ideals $\mathfrak{p}_i$, say

$$\mathfrak{a} = \mathfrak{p}_1^{e_1} \mathfrak{p}_2^{e_2} \cdots \mathfrak{p}_g^{e_g}$$

and this factorization is unique up to the order of the $\mathfrak{p}_i$.

Moreover, $\mathfrak{a}$ divides $\mathfrak{b}$ if and only if for every prime ideal $\mathfrak{p}$, the exponent of $\mathfrak{p}$ in $\mathfrak{a}$ is less than the corresponding exponent in $\mathfrak{b}$.

I won't write out the proof, but I'll describe the basic method of attack. Section 3 of [U108] does a nice job of explaining it. When we proved the fundamental theorem of arithmetic, the basic plot was:

- (1) Show that if p is a rational prime¹ then $p \mid bc$ means $p \mid b$ or $p \mid c$. (This is called Euclid's Lemma.)
- (2) Use strong induction to show that every $N > 1$ can be written as the product of primes (easy).
- (3) Show that if $p_1 \cdots p_m = q_1 \cdots q_n$ for some primes (not necessarily unique), then $p_i = q_i$ for some i , say q_1 .
- (4) Divide both sides by p_1 and use induction.

What happens if we try to repeat the proof here? We get step 1 for free, because we're using a better definition of "prime". We can also do step 3, since it follows from step 1. But step 2 doesn't work, because for abstract Dedekind domains we don't really have a notion of size. And step 4 doesn't work because we don't yet have a notion of what the inverse of a prime ideal is.

Well, it turns out that we *can* define the inverse $\mathfrak{a}^{-1}$ of an ideal, and I'll do so by the end of this chapter. You then need to check that $\mathfrak{a} \cdot \mathfrak{a}^{-1} = (1) = \mathcal{A}$. In fact, even this isn't easy. You have to check it's true for prime ideals $\mathfrak{p}$, *then* prove prime factorization, and then prove that this is true. Moreover, $\mathfrak{a}^{-1}$ is not actually an ideal, so you need to work in the field of fractions K instead of $\mathcal{A}$.

So the main steps in the new situation are as follows:

- (1) First, show that every ideal $\mathfrak{a}$ divides $\mathfrak{p}_1 \cdots \mathfrak{p}_g$ for some finite collection of primes. (This is an application of Zorn's Lemma.)
- (2) Define $\mathfrak{p}^{-1}$ and show that $\mathfrak{p}\mathfrak{p}^{-1} = (1)$.
- (3) Show that a factorization exists (again using Zorn's Lemma).
- (4) Show that it's unique, using the new inverse we've defined.

Finally, let me comment on how nice this is if $\mathcal{A}$ is a PID (like $\mathbb{Z}$). Thus every element $a \in \mathcal{A}$ is in direct correspondence with an ideal (a) . Now suppose (a) factors as a product of ideals $\mathfrak{p}_i = (p_i)$, say,

$$(a) = (p_1)^{e_1} (p_2)^{e_2} \cdots (p_n)^{e_n}.$$

¹ Note that the kindergarten definition of a prime is that " p isn't the product of two smaller integers". This isn't the correct definition of a prime: the definition of a prime is that $p \mid bc$ means $p \mid b$ or $p \mid c$. The kindergarten definition is something called "irreducible". Fortunately, in $\mathbb{Z}$, primes and irreducibles are the same thing, so no one ever told you that your definition of "prime" was wrong.

This verbatim reads

$$a = up_1^{e_1} p_2^{e_2} \cdots p_n^{e_n}$$

where u is some unit (recall Definition 12.3.1). Hence, Dedekind domains which are PID's satisfy unique factorization for *elements*, just like in $\mathbb{Z}$. (In fact, the converse of this is true.)

§37.6 The factoring algorithm

Let's look at some examples from quadratic fields. Recall that if $K = \mathbb{Q}(\sqrt{d})$, then

$$\mathcal{O}_K = \begin{cases} \mathbb{Z}[\sqrt{d}] & d \equiv 2, 3 \pmod{4} \\ \mathbb{Z}\left[\frac{1+\sqrt{d}}{2}\right] & d \equiv 1 \pmod{4}. \end{cases}$$

Also, recall that the norm of $a + b\sqrt{-d}$ is given by $a^2 + db^2$.

Example 37.6.1 (Factoring 6 in the integers of $\mathbb{Q}(\sqrt{-5})$)

Let $\mathcal{O}_K = \mathbb{Z}[\sqrt{-5}]$ arise from $K = \mathbb{Q}(\sqrt{-5})$. We've already seen that

$$(6) = (2) \cdot (3) = (1 + \sqrt{-5})(1 - \sqrt{-5})$$

and you can't get any further with these principal ideals. But let

$$\mathfrak{p} = (1 + \sqrt{-5}, 2) = (1 - \sqrt{-5}, 2) \quad \text{and} \quad \mathfrak{q}_1 = (1 + \sqrt{-5}, 3), \quad \mathfrak{q}_2 = (1 - \sqrt{-5}, 3).$$

Then it turns out $(6) = \mathfrak{p}^2 \mathfrak{q}_1 \mathfrak{q}_2$. More specifically, $(2) = \mathfrak{p}^2$, $(3) = \mathfrak{q}_1 \mathfrak{q}_2$, and $(1 + \sqrt{-5}) = \mathfrak{p} \mathfrak{q}_1$ and $(1 - \sqrt{-5}) = \mathfrak{p} \mathfrak{q}_2$. (Proof in just a moment.)

I want to stress that all our ideals are computed relative to $\mathcal{O}_K$. So for example,

$$(2) = \{2x \mid x \in \mathcal{O}_K\}.$$

How do we know in this example that $\mathfrak{p}$ is prime/maximal? (Again, these are the same since we're in a Dedekind domain.) Answer: look at $\mathcal{O}_K/\mathfrak{p}$ and see if it's a field. There is a trick to this: we can express

$$\mathcal{O}_K = \mathbb{Z}[\sqrt{-5}] \cong \mathbb{Z}[x]/(x^2 + 5).$$

So when we take *that* mod $\mathfrak{p}$, we get that

$$\mathcal{O}_K/\mathfrak{p} = \mathbb{Z}[x]/(x^2 + 5, 2, 1 + x) \cong \mathbb{F}_2[x]/(x^2 + 5, x + 1)$$

as rings.

Question 37.6.2. Conclude that $\mathcal{O}_K/\mathfrak{p} \cong \mathbb{F}_2$, and satisfy yourself that $\mathfrak{q}_1$ and $\mathfrak{q}_2$ are also maximal.

I should give an explicit example of an ideal multiplication: let's compute

$$\begin{aligned} \mathfrak{q}_1 \mathfrak{q}_2 &= ((1 + \sqrt{-5})(1 - \sqrt{-5}), 3(1 + \sqrt{-5}), 3(1 - \sqrt{-5}), 9) \\ &= (6, 3 + 3\sqrt{-5}, 3 - 3\sqrt{-5}, 9) \\ &= (6, 3 + 3\sqrt{-5}, 3 - 3\sqrt{-5}, 3) \\ &= (3) \end{aligned}$$

where we first did $9 - 6 = 3$ (think Euclidean algorithm!), then noted that all the other generators don't contribute anything we don't already have with the 3 (again these are ideals computed in $\mathcal{O}_K$). You can do the computation for $\mathfrak{p}^2$, $\mathfrak{p}\mathfrak{q}_1$, $\mathfrak{p}\mathfrak{q}_2$ in the same way.

Finally, it's worth pointing out that we should quickly verify that $\mathfrak{p} \neq (x)$ for some x ; in other words, that $\mathfrak{p}$ is not principal. Assume for contradiction that it is. Then x divides both $1 + \sqrt{-5}$ and 2, in the sense that $1 + \sqrt{-5} = \alpha_1 x$ and $2 = \alpha_2 x$ for some $\alpha_1, \alpha_2 \in \mathcal{O}_K$. (Principal ideals are exactly the “multiples” of x , so $(x) = x\mathcal{O}_K$.) Taking the norms, we find that $N_{K/\mathbb{Q}}(x)$ divides both

$$N_{K/\mathbb{Q}}(1 + \sqrt{-5}) = 6 \quad \text{and} \quad N_{K/\mathbb{Q}}(2) = 4.$$

Since $\mathfrak{p} \neq (1)$, x cannot be a unit, so its norm must be 2. But there are no elements of norm $2 = a^2 + 5b^2$ in $\mathcal{O}_K$.

Example 37.6.3 (Factoring 3 in the integers of $\mathbb{Q}(\sqrt{-17})$)

Let $\mathcal{O}_K = \mathbb{Z}[\sqrt{-17}]$ arise from $K = \mathbb{Q}(\sqrt{-17})$. We know $\mathcal{O}_K \cong \mathbb{Z}[x]/(x^2 + 17)$. Now

$$\mathcal{O}_K/3\mathcal{O}_K \cong \mathbb{Z}[x]/(3, x^2 + 17) \cong \mathbb{F}_3[x]/(x^2 - 1).$$

This already shows that (3) cannot be a prime (i.e. maximal) ideal, since otherwise our result should be a field. Anyways, we have a projection

$$\mathcal{O}_K \twoheadrightarrow \mathbb{F}_3[x]/((x-1)(x+1)).$$

Let $\mathfrak{q}_1$ be the pre-image $(x-1)$ in the image, that is,

$$\mathfrak{q}_1 = (3, \sqrt{-17} - 1).$$

Similarly,

$$\mathfrak{q}_2 = (3, \sqrt{-17} + 1).$$

We have $\mathcal{O}_K/\mathfrak{q}_1 \cong \mathbb{F}_3$, so $\mathfrak{q}_1$ is maximal (prime). Similarly $\mathfrak{q}_2$ is prime. Magically, you can check explicitly that

$$\mathfrak{q}_1\mathfrak{q}_2 = (3).$$

Hence this is the factorization of (3) into prime ideals.

The fact that $\mathfrak{q}_1\mathfrak{q}_2 = (3)$ looks magical, but it's really true:

$$\begin{aligned} \mathfrak{q}_1\mathfrak{q}_2 &= (3, \sqrt{-17} - 1)(3, \sqrt{-17} + 1) \\ &= (9, 3\sqrt{-17} + 3, 3\sqrt{-17} - 3, 18) \\ &= (9, 3\sqrt{-17} + 3, 6) \\ &= (3, 3\sqrt{-17} + 3, 6) \\ &= (3). \end{aligned}$$

In fact, it turns out this always works in general: given a rational prime p , there is an algorithm to factor p in any $\mathcal{O}_K$ of the form $\mathbb{Z}[\theta]$.

Theorem 37.6.4 (Factoring algorithm / Dedekind-Kummer theorem)

Let K be a number field. Let $\theta \in \mathcal{O}_K$ with $[\mathcal{O}_K : \mathbb{Z}[\theta]] = j < \infty$, and let p be a prime not dividing j . Then $(p) = p\mathcal{O}_K$ is factored as follows:

Let f be the minimal polynomial of θ and factor $\bar{f} \bmod p$ as

$$\bar{f} \equiv \prod_{i=1}^g (\bar{f}_i)^{e_i} \pmod{p}.$$

Then $\mathfrak{p}_i = (f_i(\theta), p)$ is prime for each i and the factorization of (p) is

$$\mathcal{O}_K \supseteq (p) = \prod_{i=1}^g \mathfrak{p}_i^{e_i}.$$

In particular, if K is monogenic with $\mathcal{O}_K = \mathbb{Z}[\theta]$ then $j = 1$ and the theorem applies for all primes p .

In almost all our applications in this book, K will be monogenic; i.e. $j = 1$. Here $\bar{\psi}$ denotes the image in $\mathbb{F}_p[x]$ of a polynomial $\psi \in \mathbb{Z}[x]$.

Question 37.6.5. There are many possible pre-images f_i we could have chosen (for example if $\bar{f}_i = x^2 + 1 \pmod{3}$, we could pick $f_i = x^2 + 3x + 7$.) Why does this not affect the value of $\mathfrak{p}_i$?

Note that earlier, we could check the factorization worked for any particular case. The proof that this works is much the same, but we need one extra tool, the ideal norm. After that we leave the proposition as Problem 37E.

This algorithm gives us a concrete way to compute prime factorizations of (p) in any monogenic number field with $\mathcal{O}_K = \mathbb{Z}[\theta]$. To summarize the recipe:

1. Find the minimal polynomial of θ , say $f \in \mathbb{Z}[x]$.
2. Factor $f \bmod p$ into irreducible polynomials $\bar{f}_1^{e_1} \bar{f}_2^{e_2} \dots \bar{f}_g^{e_g}$.
3. Compute $\mathfrak{p}_i = (f_i(\theta), p)$ for each i .

Then your $(p) = \mathfrak{p}_1^{e_1} \dots \mathfrak{p}_g^{e_g}$.

Exercise 37.6.6. Factor (29) in $\mathbb{Q}(i)$ using the above algorithm.

§37.7 Fractional ideals

Prototypical example for this section: Analog to $\mathbb{Q}$ for $\mathbb{Z}$, allowing us to take inverses of ideals. Prime factorization works in the nicest way possible.

We now have a neat theory of factoring ideals of $\mathcal{A}$, just like factoring the integers. Now note that our factorization of $\mathbb{Z}$ naturally gives a way to factor elements of $\mathbb{Q}$; just factor the numerator and denominator separately.

Let's make the analogy clearer. The analogue of a rational number is as follows.

Definition 37.7.1. Let $\mathcal{A}$ be a Dedekind domain with field of fractions K . A **fractional ideal** J of K is a set of the form

$$J = \frac{1}{x} \cdot \mathfrak{a} \quad \text{where } x \in \mathcal{A}, \text{ and } \mathfrak{a} \text{ is an integral ideal.}$$

For emphasis, ideals of $\mathcal{A}$ will be sometimes referred to as **integral ideals**.

You might be a little surprised by this definition: one would expect that a fractional ideal should be of the form $\frac{\mathfrak{a}}{\mathfrak{b}}$ for some integral ideals $\mathfrak{a}, \mathfrak{b}$. But in fact, it suffices to just take $x \in \mathcal{A}$ in the denominator. The analogy is that when we looked at $\mathcal{O}_K$, we found that we only needed integer denominators: $\frac{1}{4-\sqrt{3}} = \frac{1}{13}(4 + \sqrt{3})$. Similarly here, it will turn out that we only need to look at $\frac{1}{x} \cdot \mathfrak{a}$ rather than $\frac{\mathfrak{a}}{\mathfrak{b}}$, and so we define it this way from the beginning. See Problem 37D[†] for a different equivalent definition.

Example 37.7.2 ($\frac{5}{2}\mathbb{Z}$ is a fractional ideal)

The set

$$\frac{5}{2}\mathbb{Z} = \left\{ \frac{5}{2}n \mid n \in \mathbb{Z} \right\} = \frac{1}{2}(5\mathbb{Z})$$

is a fractional ideal of $\mathbb{Z}$.

Now, as we prescribed, the fractional ideals form a multiplicative group:

Theorem 37.7.3 (Fractional ideals form a group)

Let $\mathcal{A}$ be a Dedekind domain and K its field of fractions. For any integral ideal $\mathfrak{a}$, the set

$$\mathfrak{a}^{-1} = \{x \in K \mid x\mathfrak{a} \subseteq (1) = \mathcal{A}\}$$

is a fractional ideal with $\mathfrak{a}\mathfrak{a}^{-1} = (1)$.

Definition 37.7.4. Thus nonzero fractional ideals of K form a group under multiplication with identity $(1) = \mathcal{A}$. This **ideal group** is denoted J_K .

Example 37.7.5 ($(3)^{-1}$ in $\mathbb{Z}$)

Please check that in $\mathbb{Z}$ we have

$$(3)^{-1} = \left\{ \frac{1}{3}n \mid n \in \mathbb{Z} \right\} = \frac{1}{3}\mathbb{Z}.$$

It follows that every fractional ideal J can be uniquely written as

$$J = \prod_i \mathfrak{p}_i^{n_i} \cdot \prod_i \mathfrak{q}_i^{-m_i}$$

where n_i and m_i are positive integers. In fact, $\mathfrak{a}$ is an integral ideal if and only if all its exponents are nonnegative, just like the case with integers. So, a perhaps better way to think about fractional ideals is as products of prime ideals, possibly with negative exponents.

§37.8 The ideal norm

One last tool is the ideal norm, which gives us a notion of the “size” of an ideal.

Definition 37.8.1. The **ideal norm** (or absolute norm) of a nonzero ideal $\mathfrak{a} \subseteq \mathcal{O}_K$ is defined as $|\mathcal{O}_K/\mathfrak{a}|$ and denoted $N(\mathfrak{a})$.

Example 37.8.2 (Ideal norm of (5) in the Gaussian integers)

Let $K = \mathbb{Q}(i)$, $\mathcal{O}_K = \mathbb{Z}[i]$. Consider the ideal (5) in $\mathcal{O}_K$. We have that

$$\mathcal{O}_K/(5) \cong \{a + bi \mid a, b \in \mathbb{Z}_5\}$$

so (5) has ideal norm 25, corresponding to the fact that $\mathcal{O}_K/(5)$ has $5^2 = 25$ elements.

Example 37.8.3 (Ideal norm of (2 + i) in the Gaussian integers)

You’ll notice that

$$\mathcal{O}_K/(2 + i) \cong \mathbb{F}_5$$

since mod $2 + i$ we have both $5 \equiv 0$ and $i \equiv -2$. (Indeed, since $(2 + i)$ is prime we had better get a field!) Thus $N((2 + i)) = 5$; similarly $N((2 - i)) = 5$.

Thus the ideal norm measures how “roomy” the ideal is: that is, (5) is a lot more spaced out in $\mathbb{Z}[i]$ than it is in $\mathbb{Z}$. (This intuition will be important when we will actually view $\mathcal{O}_K$ as a lattice.)

Question 37.8.4. What are the ideals with ideal norm one?

Our example with (5) suggests several properties of the ideal norm which turn out to be true:

Lemma 37.8.5 (Properties of the absolute norm)

Let $\mathfrak{a}$ be a nonzero ideal of $\mathcal{O}_K$.

- (a) $N(\mathfrak{a})$ is finite.
- (b) For any other nonzero ideal $\mathfrak{b}$, $N(\mathfrak{a}\mathfrak{b}) = N(\mathfrak{a})N(\mathfrak{b})$.
- (c) If $\mathfrak{a} = (a)$ is principal, then $N(\mathfrak{a}) = N_{K/\mathbb{Q}}(a)$.

I unfortunately won’t prove these properties, though we already did (a) in our proof that $\mathcal{O}_K$ was a Dedekind domain.

The fact that N is completely multiplicative lets us also consider the norm of a fractional ideal J by the natural extension

$$J = \prod_i \mathfrak{p}_i^{n_i} \cdot \prod_i \mathfrak{q}_i^{-m_i} \implies N(J) \stackrel{\text{def}}{=} \frac{\prod_i N(\mathfrak{p}_i)^{n_i}}{\prod_i N(\mathfrak{q}_i)^{m_i}}.$$

Thus N is a natural group homomorphism $J_K \rightarrow \mathbb{Q} \setminus \{0\}$.

§37.9 Problems to think about

Problem 37A. Show that there are three different factorizations of 77 in $\mathcal{O}_K$, where $K = \mathbb{Q}(\sqrt{-13})$.

Problem 37B. Let $K = \mathbb{Q}(\sqrt[3]{2})$; one can show that $\mathcal{O}_K = \mathbb{Z}[\sqrt[3]{2}]$. Find the factorization of (5) in $\mathcal{O}_K$.

Problem 37C (Fermat's little theorem). Let $\mathfrak{p}$ be a prime ideal in some ring of integers $\mathcal{O}_K$. Show that for $\alpha \in \mathcal{O}_K$,

$$\alpha^{N(\mathfrak{p})} \equiv \alpha \pmod{\mathfrak{p}}.$$

Problem 37D[†]. Let $\mathcal{A}$ be a Dedekind domain with field of fractions K , and pick $J \subseteq K$. Show that J is a fractional ideal if and only if

- (i) J is closed under addition and multiplication by elements of $\mathcal{A}$, and
- (ii) J is finitely generated as an abelian group.

More succinctly: J is a fractional ideal $\iff J$ is a finitely generated $\mathcal{A}$ -module.

Problem 37E. In the notation of Theorem 37.6.4, let $I = \prod_{i=1}^g \mathfrak{p}_i^{e_i}$. Assume for simplicity that K is monogenic, hence $\mathcal{O}_K = \mathbb{Z}[\theta]$.

- (a) Prove that each $\mathfrak{p}_i$ is prime.
- (b) Show that (p) divides I .
- (c) Use the norm to show that $(p) = I$.

38 Minkowski bound and class groups

We now have a neat theory of unique factorization of ideals. In the case of a PID, this in fact gives us unique factorization of elements. (We call this a **unique factorization domain**.) Sweet.

We'll define, in a moment, something called the *class group* which measures how far $\mathcal{O}_K$ is from being a PID; the bigger the class group, the farther $\mathcal{O}_K$ is from being a PID. In particular, the $\mathcal{O}_K$ is a PID if it has trivial class group.

Then we will provide some inequalities which let us put restrictions on the class group; for instance, this will let us show in some cases that the class group must be trivial. Astonishingly, the proof will use Minkowski's theorem, a result from geometry.

§38.1 The class group

Prototypical example for this section: PID's have trivial class group

Let K be a number field, and let J_K denote the multiplicative group of fractional ideals of $\mathcal{O}_K$. Let P_K denote the multiplicative group of **principal fractional ideals**: those of the form $(x) = x\mathcal{O}_K$ for some $x \in K$.

Question 38.1.1. Check that P_K is also a multiplicative group. (This is really easy: name $x\mathcal{O}_K \cdot y\mathcal{O}_K$ and $(x\mathcal{O}_K)^{-1}$.)

As J_K is abelian, we can now define the **class group** to be the quotient

$$\text{Cl}_K \stackrel{\text{def}}{=} J_K / P_K.$$

The elements of Cl_K are called **classes**.

Equivalently,

The class group Cl_K is the set of nonzero fractional ideals modulo scaling by a constant in K .

In particular, Cl_K is trivial if all ideals are principal, since the nonzero principal ideals are the same up to scaling.

The size of the class group is called the **class number**. It's a beautiful theorem that the class number is always finite, and the bulk of this chapter will build up to this result. It requires several ingredients.

§38.2 The discriminant of a number field

Prototypical example for this section: Quadratic fields.

Let's say I have $K = \mathbb{Q}(\sqrt{2})$. As we've seen before, this means $\mathcal{O}_K = \mathbb{Z}[\sqrt{2}]$, meaning

$$\mathcal{O}_K = \{a + b\sqrt{2} \mid a, b \in \mathbb{Z}\}.$$

The key insight now is that you might think of this as a *lattice*: geometrically, we want to think about this the same way we think about $\mathbb{Z}^2$.


Perversely, we might try to embed this into $\mathbb{Q}^2$ by sending $a + b\sqrt{2}$ to (a, b) . But this is a little stupid, since we're rudely making K , which somehow lives inside $\mathbb{Q}$ and is “one-dimensional” in that sense, into a two-dimensional space. It also depends on a choice of basis, which we don't like. A better way is to think about the fact that there are two embeddings $\sigma_1 : K \rightarrow \mathbb{C}$ and $\sigma_2 : K \rightarrow \mathbb{C}$, namely the identity, and conjugation:

$$\begin{aligned}\sigma_1(a + b\sqrt{2}) &= a + b\sqrt{2} \\ \sigma_2(a + b\sqrt{2}) &= a - b\sqrt{2}.\end{aligned}$$

Fortunately for us, these embeddings both have real image. This leads us to consider the set of points

$$(\sigma_1(\alpha), \sigma_2(\alpha)) \in \mathbb{R}^2 \quad \text{as } \alpha \in K.$$

This lets us visualize what $\mathcal{O}_K$ looks like in $\mathbb{R}^2$. The points of K are dense in $\mathbb{R}^2$, but the points of $\mathcal{O}_K$ cut out a lattice.



To see how big the lattice is, we look at how $\{1, \sqrt{2}\}$, the generators of $\mathcal{O}_K$, behave. The point corresponding to $a + b\sqrt{2}$ in the lattice is

$$a \cdot (1, 1) + b \cdot (\sqrt{2}, -\sqrt{2}).$$

The **mesh** of the lattice¹ is defined as the hypervolume of the “fundamental parallelepiped” I’ve colored blue above. For this particular case, it ought to be equal to the area of that parallelogram, which is

$$\det \begin{bmatrix} 1 & -\sqrt{2} \\ 1 & \sqrt{2} \end{bmatrix} = 2\sqrt{2}.$$

The definition of the discriminant is precisely this, except with an extra square factor (since permutation of rows could lead to changes in sign in the matrix above). Problem 39B* shows that the squaring makes Δ_K an integer.

To make the next definition, we invoke:

¹Most authors call this the volume, but I think this is not the right word to use – lattices have “volume” zero since they are just a bunch of points! In contrast, the English word “mesh” really does refer to the width of a “gap”.

Theorem 38.2.1 (The n embeddings of a number field)

Let K be a number field of degree n . Then there are exactly n field homomorphisms $K \hookrightarrow \mathbb{C}$, say $\sigma_1, \dots, \sigma_n$, which fix $\mathbb{Q}$.

Proof. Deferred to Theorem 41.3.1, once we have the tools of Galois theory. $\square$

In fact, in Theorem 41.3.4 we see that for $\alpha \in K$, we have that $\sigma_i(\alpha)$ runs over the conjugates of α as $i = 1, \dots, n$. It follows that

$$\mathrm{Tr}_{K/\mathbb{Q}}(\alpha) = \sum_{i=1}^n \sigma_i(\alpha) \quad \text{and} \quad \mathrm{N}_{K/\mathbb{Q}}(\alpha) = \prod_{i=1}^n \sigma_i(\alpha).$$

This allows us to define:

Definition 38.2.2. Suppose $\alpha_1, \dots, \alpha_n$ is a $\mathbb{Z}$ -basis of $\mathcal{O}_K$. The **discriminant** of the number field K is defined by

$$\Delta_K \stackrel{\text{def}}{=} \det \begin{bmatrix} \sigma_1(\alpha_1) & \dots & \sigma_n(\alpha_1) \\ \vdots & \ddots & \vdots \\ \sigma_1(\alpha_n) & \dots & \sigma_n(\alpha_n) \end{bmatrix}^2.$$

This does not depend on the choice of the $\{\alpha_i\}$; we will not prove this here.

Example 38.2.3 (Discriminant of $K = \mathbb{Q}(\sqrt{2})$)

We have $\mathcal{O}_K = \mathbb{Z}[\sqrt{2}]$ and as discussed above the discriminant is

$$\Delta_K = (-2\sqrt{2})^2 = 8.$$

Example 38.2.4 (Discriminant of $\mathbb{Q}(i)$)

Let $K = \mathbb{Q}(i)$. We have $\mathcal{O}_K = \mathbb{Z}[i] = \mathbb{Z} \oplus i\mathbb{Z}$. The embeddings are the identity and complex conjugation which take 1 to $(1, 1)$ and i to $(i, -i)$. So

$$\Delta_K = \det \begin{bmatrix} 1 & 1 \\ i & -i \end{bmatrix}^2 = (-2i)^2 = -4.$$

This example illustrates that the discriminant need not be positive for number fields which wander into the complex plane (the lattice picture is a less perfect analogy). But again, as we'll prove in the problems the discriminant is always an integer.

Example 38.2.5 (Discriminant of $\mathbb{Q}(\sqrt{5})$)

Let $K = \mathbb{Q}(\sqrt{5})$. This time, $\mathcal{O}_K = \mathbb{Z} \oplus \frac{1+\sqrt{5}}{2}\mathbb{Z}$, and so the discriminant is going to look a little bit different. The embeddings are still $a + b\sqrt{5} \mapsto a + b\sqrt{5}, a - b\sqrt{5}$.

Applying this to the $\mathbb{Z}$ -basis $\left\{1, \frac{1+\sqrt{5}}{2}\right\}$, we get

$$\Delta_K^2 = \det \begin{bmatrix} 1 & 1 \\ \frac{1+\sqrt{5}}{2} & \frac{1-\sqrt{5}}{2} \end{bmatrix}^2 = (-\sqrt{5})^2 = 5.$$

Exercise 38.2.6. Extend all this to show that if $K = \mathbb{Q}(\sqrt{d})$ for $d \neq 1$ squarefree, we have

$$\Delta_K = \begin{cases} d & \text{if } d \equiv 1 \pmod{4} \\ 4d & \text{if } d \equiv 2, 3 \pmod{4}. \end{cases}$$

Actually, let me point out something curious: recall that the polynomial discriminant of $Ax^2 + Bx + C$ is $B^2 - 4AC$. Then:

- In the $d \equiv 1 \pmod{4}$ case, Δ_K is the discriminant of $x^2 - x - \frac{d-1}{4}$, which is the minimal polynomial of $\frac{1}{2}(1 + \sqrt{d})$. Of course, $\mathcal{O}_K = \mathbb{Z}[\frac{1}{2}(1 + \sqrt{d})]$.
- In the $d \equiv 2, 3 \pmod{4}$ case, Δ_K is the discriminant of $x^2 - d$ which is the minimal polynomial of $\sqrt{d}$. Once again, $\mathcal{O}_K = \mathbb{Z}[\sqrt{d}]$.

This is not a coincidence! Problem 39C* asserts that this is true in general; hence the name “discriminant”.

§38.3 The signature of a number field

Prototypical example for this section: $\mathbb{Q}(\sqrt[100]{2})$ has signature $(2, 49)$.

In the example of $K = \mathbb{Q}(i)$, we more or less embedded K into the space $\mathbb{C}$. However, K is a degree two extension, so what we’d really like to do is embed it into $\mathbb{R}^2$. To do so, we’re going to take advantage of complex conjugation.

Let K be a number field and $\sigma_1, \dots, \sigma_n$ be its embeddings. We distinguish between the **real embeddings** (which map all of K into $\mathbb{R}$) and the **complex embeddings** (which map some part of K outside $\mathbb{R}$). Notice that if σ is a complex embedding, then so is the conjugate $\bar{\sigma} \neq \sigma$; hence complex embeddings come in pairs.

Definition 38.3.1. Let K be a number field of degree n , and set

r_1 = number of real embeddings

r_2 = number of pairs of complex embeddings.

The **signature** of K is the pair (r_1, r_2) . Observe that $r_1 + 2r_2 = n$.

Example 38.3.2 (Basic examples of signatures)

(a) $\mathbb{Q}$ has signature $(1, 0)$.

(b) $\mathbb{Q}(\sqrt{2})$ has signature $(2, 0)$.

(c) $\mathbb{Q}(i)$ has signature $(0, 1)$.

(d) Let $K = \mathbb{Q}(\sqrt[3]{2})$, and let ω be a cube root of unity. The elements of K are

$$K = \{a + b\sqrt[3]{2} + c\sqrt[3]{4} \mid a, b, c \in \mathbb{Q}\}.$$

Then the signature is $(1, 1)$, because the three embeddings are

$$\sigma_1 : \sqrt[3]{2} \mapsto \sqrt[3]{2}, \quad \sigma_2 : \sqrt[3]{2} \mapsto \sqrt[3]{2}\omega, \quad \sigma_3 : \sqrt[3]{2} \mapsto \sqrt[3]{2}\omega^2.$$

The first of these is real and the latter two are conjugate pairs.

Example 38.3.3 (Even more signatures)

In the same vein $\mathbb{Q}(\sqrt[99]{2})$ and $\mathbb{Q}(\sqrt[100]{2})$ have signatures $(1, 49)$ and $(2, 49)$.

Question 38.3.4. Verify the signatures of the above two number fields.

From now on, we will number the embeddings of K in such a way that

$$\sigma_1, \sigma_2, \dots, \sigma_{r_1}$$

are the real embeddings, while

$$\sigma_{r_1+1} = \overline{\sigma_{r_1+2}}, \quad \sigma_{r_1+2} = \overline{\sigma_{r_1+3}}, \quad \dots, \quad \sigma_{r_1+r_2} = \overline{\sigma_{r_1+r_2+1}}.$$

are the r_2 pairs of complex embeddings. We define the **canonical embedding** of K as

$$K \xhookrightarrow{\ell} \mathbb{R}^{r_1} \times \mathbb{C}^{r_2} \quad \text{by} \quad \alpha \mapsto (\sigma_1(\alpha), \dots, \sigma_{r_1}(\alpha), \sigma_{r_1+1}(\alpha), \dots, \sigma_{r_1+r_2}(\alpha)).$$

All we've done is omit, for the complex case, the second of the embeddings in each conjugate pair. This is no big deal, since they are just conjugates; the above tuple is all the information we need.

For reasons that will become obvious in a moment, I'll let τ denote the isomorphism

$$\tau : \mathbb{R}^{r_1} \times \mathbb{C}^{r_2} \xrightarrow{\sim} \mathbb{R}^{r_1+2r_2} = \mathbb{R}^n$$

by breaking each complex number into its real and imaginary part, as

$$\begin{aligned} \alpha \mapsto & (\sigma_1(\alpha), \dots, \sigma_{r_1}(\alpha), \\ & \operatorname{Re} \sigma_{r_1+1}(\alpha), \operatorname{Im} \sigma_{r_1+1}(\alpha), \\ & \operatorname{Re} \sigma_{r_1+2}(\alpha), \operatorname{Im} \sigma_{r_1+2}(\alpha), \\ & \dots, \\ & \operatorname{Re} \sigma_{r_1+r_2}(\alpha), \operatorname{Im} \sigma_{r_1+r_2}(\alpha)). \end{aligned}$$

Example 38.3.5 (Example of canonical embedding)

As before let $K = \mathbb{Q}(\sqrt[3]{2})$ and set

$$\sigma_1 : \sqrt[3]{2} \mapsto \sqrt[3]{2}, \quad \sigma_2 : \sqrt[3]{2} \mapsto \sqrt[3]{2}\omega, \quad \sigma_3 : \sqrt[3]{2} \mapsto \sqrt[3]{2}\omega^2$$

where $\omega = -\frac{1}{2} + \frac{\sqrt{3}}{2}i$, noting that we've already arranged indices so $\sigma_1 = \operatorname{id}$ is real while σ_2 and σ_3 are a conjugate pair. So the embeddings $K \xhookrightarrow{\ell} \mathbb{R} \times \mathbb{C} \xrightarrow{\sim} \mathbb{R}^3$ are given by

$$\alpha \xhookrightarrow{\ell} (\sigma_1(\alpha), \sigma_2(\alpha)) \xrightarrow{\tau} (\sigma_1(\alpha), \operatorname{Re} \sigma_2(\alpha), \operatorname{Im} \sigma_2(\alpha)).$$

For concreteness, taking $\alpha = 9 + \sqrt[3]{2}$ gives

$$\begin{aligned} 9 + \sqrt[3]{2} & \xhookrightarrow{\ell} \left(9 + \sqrt[3]{2}, 9 + \sqrt[3]{2}\omega \right) \\ & = \left(9 + \sqrt[3]{2}, 9 - \frac{1}{2}\sqrt[3]{2} + \frac{\sqrt{108}}{2}i \right) \in \mathbb{R} \times \mathbb{C} \\ & \xrightarrow{\tau} \left(9 + \sqrt[3]{2}, 9 - \frac{1}{2}\sqrt[3]{2}, \frac{\sqrt{108}}{2} \right) \in \mathbb{R}^3. \end{aligned}$$

Now, the whole point of this is that we want to consider the resulting lattice when we take $\mathcal{O}_K$. In fact, we have:

Lemma 38.3.6

Consider the composition of the embeddings $K \hookrightarrow \mathbb{R}^{r_1} \times \mathbb{C}^{r_2} \xrightarrow{\sim} \mathbb{R}^n$. Then as before, $\mathcal{O}_K$ becomes a lattice L in $\mathbb{R}^n$, with mesh equal to

$$\frac{1}{2^{r_2}} \sqrt{|\Delta_K|}.$$

Proof. Fun linear algebra problem (you just need to manipulate determinants). Left as Problem 38D. $\square$

From this we can deduce:

Lemma 38.3.7

Consider the composition of the embeddings $K \hookrightarrow \mathbb{R}^{r_1} \times \mathbb{C}^{r_2} \xrightarrow{\sim} \mathbb{R}^n$. Let $\mathfrak{a}$ be an ideal in $\mathcal{O}_K$. Then the image of $\mathfrak{a}$ is a lattice $L_{\mathfrak{a}}$ in $\mathbb{R}^n$ with mesh equal to

$$\frac{N(\mathfrak{a})}{2^{r_2}} \sqrt{|\Delta_K|}.$$

Sketch of Proof. Let

$$d = N(\mathfrak{a}) \stackrel{\text{def}}{=} [\mathcal{O}_K : \mathfrak{a}].$$

Then in the lattice $L_{\mathfrak{a}}$, we somehow only take $\frac{1}{d}$ th of the points which appear in the lattice L , which is why the area increases by a factor of $N(\mathfrak{a})$. To make this all precise I would need to do a lot more with lattices and geometry than I have space for in this chapter, so I will omit the details. But I hope you can see why this is intuitively true. $\square$

§38.4 Minkowski's theorem

Now I can tell you why I insisted we move from $\mathbb{R}^{r_1} \times \mathbb{C}^{r_2}$ to $\mathbb{R}^n$. In geometry, there's a really cool theorem of Minkowski's that goes as follows.

Theorem 38.4.1 (Minkowski)

Let $S \subseteq \mathbb{R}^n$ be a convex set containing 0 which is centrally symmetric (meaning that $x \in S \iff -x \in S$). Let L be a lattice with mesh d . If either

- (a) The volume of S exceeds $2^n d$, or
- (b) The volume of S equals $2^n d$ and S is compact,

then S contains a nonzero lattice point of L .

Question 38.4.2. Show that the condition $0 \in S$ is actually extraneous in the sense that any nonempty, convex, centrally symmetric set contains the origin.

Sketch of Proof. Part (a) is surprisingly simple and has a very olympiad-esque solution: it's basically Pigeonhole on areas. We'll prove part (a) in the special case $n = 2$, $L = \mathbb{Z}^2$ for simplicity as the proof can easily be generalized to any lattice and any n . Thus we want to show that any such convex set S with area more than 4 contains a lattice point.

Dissect the plane into 2×2 squares

$$[2a - 1, 2a + 1] \times [2b - 1, 2b + 1]$$

and overlay all these squares on top of each other. By the Pigeonhole Principle, we find there exist two points $p \neq q \in S$ which map to the same point. Since S is symmetric, $-q \in S$. Then $\frac{1}{2}(p - q) \in S$ (convexity) and is a nonzero lattice point.

I'll briefly sketch part (b): the idea is to consider $(1 + \varepsilon)S$ for $\varepsilon > 0$ (this is " S magnified by a small factor $1 + \varepsilon$ "). This satisfies condition (a). So for each $\varepsilon > 0$ the set of nonzero lattice points in $(1 + \varepsilon)S$, say S_ε , is a *finite nonempty set* of (discrete) points (the "finite" part follows from the fact that $(1 + \varepsilon)S$ is bounded). So there has to be some point that's in S_ε for every $\varepsilon > 0$ (why?), which implies it's in S . $\square$

§38.5 The trap box


The last ingredient we need is a set to apply Minkowski's theorem to. I propose:

Definition 38.5.1. Let M be a positive real. In $\mathbb{R}^{r_1} \times \mathbb{C}^{r_2}$, define the box S to be the set of points $(x_1, \dots, x_{r_1}, z_1, \dots, z_{r_2})$ such that

$$\sum_{i=1}^{r_1} |x_i| + 2 \sum_{j=1}^{r_2} |z_j| \leq M.$$

Note that this depends on the value of M .

Think of this box as a *mousetrap*: anything that falls in it is going to have a small norm, and our goal is to use Minkowski to lure some nonzero element into it.



That is, suppose $\alpha \in \mathfrak{a}$ falls into the box I've defined above, which means

$$M \geq \sum_{i=1}^{r_1} |\sigma_i(\alpha)| + 2 \sum_{i=r_1+1}^{r_1+r_2} |\sigma_i(\alpha)| = \sum_{i=1}^n |\sigma_i(\alpha)|,$$

where we are remembering that the last few σ 's come in conjugate pairs. This looks like the trace, but the absolute values are in the way. So instead, we apply AM-GM to obtain:

Lemma 38.5.2 (Effect of the mousetrap)

Let $\alpha \in \mathcal{O}_K$, and suppose $\iota(\alpha)$ is in S (where $\iota : K \hookrightarrow \mathbb{R}^{r_1} \times \mathbb{C}^{r_2}$ as usual). Then

$$N_{K/\mathbb{Q}}(\alpha) = \prod_{i=1}^n |\sigma_i(\alpha)| \leq \left(\frac{M}{n}\right)^n.$$

The last step we need to do is compute the volume of the box. This is again some geometry I won't do, but take my word for it:

Lemma 38.5.3 (Size of the mousetrap)

Let $\tau : \mathbb{R}^{r_1} \times \mathbb{C}^{r_2} \xrightarrow{\sim} \mathbb{R}^n$ as before. Then the image of S under τ is a convex, compact, centrally symmetric set with volume

$$2^{r_1} \cdot \left(\frac{\pi}{2}\right)^{r_2} \cdot \frac{M^n}{n!}.$$

Question 38.5.4. (Sanity check) Verify that the above is correct for the signatures $(r_1, r_2) = (2, 0)$ and $(r_1, r_2) = (0, 1)$, which are the possible signatures when $n = 2$.

§38.6 The Minkowski bound

We can now put everything we have together to obtain the great Minkowski bound.

Theorem 38.6.1 (Minkowski bound)

Let $\mathfrak{a} \subseteq \mathcal{O}_K$ be any nonzero ideal. Then there exists $0 \neq \alpha \in \mathfrak{a}$ such that

$$N_{K/\mathbb{Q}}(\alpha) \leq \left(\frac{4}{\pi}\right)^{r_2} \frac{n!}{n^n} \sqrt{|\Delta_K|} \cdot N(\mathfrak{a}).$$

Proof. This is a matter of putting all our ingredients together. Let's see what things we've defined already:

$$K \xrightarrow{\iota} \mathbb{R}^{r_1} \times \mathbb{C}^{r_2} \xrightarrow{\tau} \mathbb{R}^n$$

$$\text{box } S \longmapsto \tau(S) \quad \text{with volume } 2^{r_1} \left(\frac{\pi}{2}\right)^{r_2} \frac{M^n}{n!}$$

$$\mathcal{O}_K \longmapsto \text{Lattice } L \quad \text{with mesh } 2^{-r_2} \sqrt{|\Delta_K|}$$

$$\mathfrak{a} \longmapsto \text{Lattice } L_{\mathfrak{a}} \quad \text{with mesh } 2^{-r_2} \sqrt{|\Delta_K|} N(\mathfrak{a})$$

Pick a value of M such that the mesh of $L_{\mathfrak{a}}$ equals 2^{-n} of the volume of the box. Then Minkowski's theorem gives that some $0 \neq \alpha \in \mathfrak{a}$ lands inside the box — the mousetrap is configured to force $N_{K/\mathbb{Q}}(\alpha) \leq \frac{1}{n^n} M^n$. The correct choice of M is

$$M^n = M^n \cdot 2^n \cdot \frac{\text{mesh}}{\text{vol box}} = 2^n \cdot \frac{n!}{2^{r_1} \cdot \left(\frac{\pi}{2}\right)^{r_2}} \cdot 2^{-r_2} \sqrt{|\Delta_K|} N(\mathfrak{a})$$

which gives the bound after some arithmetic. $\square$

§38.7 The class group is finite

Definition 38.7.1. Let $M_K = \left(\frac{4}{\pi}\right)^{r_2} \frac{n!}{n^n} \sqrt{|\Delta_K|}$ for brevity. Note that it is a constant depending on K .

So that's cool and all, but what we really wanted was to show that the class group is finite. How can the Minkowski bound help? Well, you might notice that we can rewrite it to say

$$N((\alpha) \cdot \mathfrak{a}^{-1}) \leq M_K$$

where M_K is some constant depending on K , and $\alpha \in \mathfrak{a}$.

Question 38.7.2. Show that $(\alpha) \cdot \mathfrak{a}^{-1}$ is an integral ideal. (Unwind definitions.)

But in the class group we *mod out* by principal ideals like (α) . If we shut our eyes for a moment and mod out, the above statement becomes “ $N(\mathfrak{a}^{-1}) \leq M_K$ ”. The precise statement of this is

Corollary 38.7.3

Let K be a number field, and pick a fractional ideal J . Then we can find α such that $\mathfrak{b} = (\alpha) \cdot J$ is integral and $N(\mathfrak{b}) \leq M_K$.

Proof. For fractional ideals I and J write $I \sim J$ to mean that $I = (\alpha)J$ for some α ; then Cl_K is just modding out by $\sim$. Let J be a fractional ideal. Then J^{-1} is some other fractional ideal. By definition, for some $\alpha \in \mathcal{O}_K$ we have that αJ^{-1} is an integral ideal $\mathfrak{a}$. The Minkowski bound tells us that for some $x \in \mathfrak{a}$, we have $N(x\mathfrak{a}^{-1}) \leq M_K$. But $x\mathfrak{a}^{-1} \sim \mathfrak{a}^{-1} = (\alpha J^{-1})^{-1} \sim J$. $\square$

Corollary 38.7.4 (Finiteness of class group)

Class groups are always finite.

Proof. For every class in Cl_K , we can identify an integral ideal $\mathfrak{a}$ with norm less than M_K . We just have to show there are finitely many such integral ideals; this will mean there are finitely many classes.

Suppose we want to build such an ideal $\mathfrak{a} = \mathfrak{p}_1^{e_1} \dots \mathfrak{p}_m^{e_m}$. Recall that a prime ideal $\mathfrak{p}_i$ must have some rational prime p inside it, meaning $\mathfrak{p}_i$ divides (p) and p divides $N(\mathfrak{p}_i)$. So let's group all the $\mathfrak{p}_i$ we want to build $\mathfrak{a}$ with based on which (p) they came from.

$$\begin{array}{cccc} (2) & (3) & (5) & \dots \\ \vdots & \vdots & \vdots & \\ \vdots & \vdots & \vdots & \\ \vdots & \vdots & \vdots & \end{array}$$

To be more dramatic: imagine you have a *cherry tree*; each branch corresponds to a prime (p) and contains as cherries (prime ideals) the factors of (p) (finitely many). Your bucket (the ideal $\mathfrak{a}$ you're building) can only hold a total weight (norm) of M_K . So you

can't even touch the branches higher than M_K . You can repeat cherries (oops), but the weight of a cherry on branch (p) is definitely $\geq p$; all this means that the number of ways to build $\mathfrak{a}$ is finite. $\square$

§38.8 Computation of class numbers

Definition 38.8.1. The order of Cl_K is called the **class number** of K .

Remark 38.8.2. If $\text{Cl}_K = 1$, then $\mathcal{O}_K$ is a PID, hence a UFD.

By computing the actual value of M_K , we can quite literally build the entire “cherry tree” mentioned in the previous proof. Let's give an example how!

Proposition 38.8.3

The field $\mathbb{Q}(\sqrt{-67})$ has class number 1.

Proof. Since $K = \mathbb{Q}(\sqrt{-67})$ has signature $(0, 1)$ and discriminant $\Delta_K = -67$ (since $-67 \equiv 1 \pmod{4}$) we can compute

$$M_K = \left(\frac{4}{\pi}\right)^1 \cdot \frac{2!}{2^2} \sqrt{67} \approx 5.2.$$

That means we can cut off the cherry tree after $(2), (3), (5)$, since any cherries on these branches will necessarily have norm $\geq M_K$. We now want to factor each of these in $\mathcal{O}_K = \mathbb{Z}[\theta]$, where $\theta = \frac{1+\sqrt{-67}}{2}$ has minimal polynomial $x^2 - x + 17$. But something miraculous happens:

- When we try to reduce $x^2 - x + 17 \pmod{2}$, we get an irreducible polynomial $x^2 - x + 1$. By the factoring algorithm (Theorem 37.6.4) this means (2) is prime.
- Similarly, reducing mod 3 gives $x^2 - x + 2$, which is irreducible. This means (3) is prime.
- Finally, for the same reason, (5) is prime.

It's our lucky day; all of the ideals $(2), (3), (5)$ are prime (already principal). To put it another way, each of the three branches has only one (large) cherry on it. That means any time we put together an integral ideal with norm $\leq M_K$, it is actually principal. In fact, these guys have norm 4, 9, 25 respectively...so we can't even touch (3) and (5) , and the only ideals we can get are (1) and (2) (with norms 1 and 4).

Now we claim that's all. Pick a fractional ideal J . By Corollary 38.7.3, we can find an integral ideal $\mathfrak{b} \sim J$ with $N(\mathfrak{b}) \leq M_K$. But by the above, either $\mathfrak{b} = (1)$ or $\mathfrak{b} = (4)$, both of which are principal, and hence trivial in Cl_K . So J is trivial in Cl_K too, as needed. $\square$

Let's do a couple more.

Theorem 38.8.4 (Gaussian integers $\mathbb{Z}[i]$ form a UFD)

The field $\mathbb{Q}(i)$ has class number 1.

Proof. This is $\mathcal{O}_K$ where $K = \mathbb{Q}(i)$, so we just want Cl_K to be trivial. We have $M_K = \frac{2}{\pi}\sqrt{4} < 2$. So every class has an integral ideal of norm $\mathfrak{b}$ satisfying

$$N(\mathfrak{b}) \left(\frac{4}{\pi}\right)^1 \cdot \frac{2!}{2^2} \cdot \sqrt{4} = \frac{4}{\pi} < 2.$$

Well, that's silly: we don't have any branches to pick from at all. In other words, we can only have $\mathfrak{b} = (1)$. $\square$

Here's another example of something that still turns out to be unique factorization, but this time our cherry tree will actually have cherries that can be picked.

Proposition 38.8.5 ($\mathbb{Z}[\sqrt{7}]$ is a UFD)

The field $\mathbb{Q}(\sqrt{7})$ has class number 1.

Proof. First we compute the Minkowski bound.

Question 38.8.6. Check that $M_K \approx 2.646$.

So this time, the only branch is (2). Let's factor (2) as usual: the polynomial $x^2 + 7$ reduces as $(x-1)(x+1) \pmod{2}$, and hence

$$(2) = (2, \sqrt{7}-1)(2, \sqrt{7}+1).$$

Oops! We now have two cherries, and they both seem reasonable. But actually, I claim that

$$(2, \sqrt{7}-1) = (3-\sqrt{7}).$$

Question 38.8.7. Prove this.

So both the cherries are principal ideals, and as before we conclude that Cl_K is trivial. But note that this time, the prime ideal (2) actually splits; we got lucky that the two cherries were principal but this won't always work. $\square$

How about some nontrivial class groups? First, we use a lemma that will help us with narrowing down the work in our cherry tree.

Lemma 38.8.8 (Ideals divide their norms)

Let $\mathfrak{b}$ be an integral ideal with $N(\mathfrak{b}) = n$. Then $\mathfrak{b}$ divides the ideal (n) .

Proof. By definition, $n = |\mathcal{O}_K/\mathfrak{b}|$. Treating $\mathcal{O}_K/\mathfrak{b}$ as an (additive) abelian group and using Lagrange's theorem, we find

$$0 \equiv \underbrace{\alpha + \cdots + \alpha}_{n \text{ times}} = n\alpha \pmod{\mathfrak{b}} \quad \text{for all } \alpha \in \mathcal{O}_K.$$

Thus $(n) \subseteq \mathfrak{b}$, done. $\square$

Now we can give such an example.

Proposition 38.8.9 (Class group of $\mathbb{Q}(\sqrt{-17})$)

The number field $K = \mathbb{Q}(\sqrt{-17})$ has class group $\mathbb{Z}_4$.

You are not obliged to read the entire proof in detail, as it is somewhat gory. The idea is just that there are some cherries which are not trivial in the class group.

Proof. Since $\Delta_K = -68$, we compute the Minkowski bound

$$M_K = \frac{4}{\pi} \sqrt{17} < 6.$$

Now, it suffices to factor with (2), (3), (5). The minimal polynomial of $\sqrt{-17}$ is $x^2 + 17$, so as usual

$$\begin{aligned}(2) &= (2, \sqrt{-17} + 1)^2 \\ (3) &= (3, \sqrt{-17} - 1)(3, \sqrt{-17} + 1) \\ (5) &= (5)\end{aligned}$$

corresponding to the factorizations of $x^2 + 17$ modulo each of 2, 3, 5. Set $\mathfrak{p} = (2, \sqrt{-17} + 1)$ and $\mathfrak{q}_1 = (3, \sqrt{-17} - 1)$, $\mathfrak{q}_2 = (3, \sqrt{-17} + 1)$. We can compute

$$N(\mathfrak{p}) = 2 \quad \text{and} \quad N(\mathfrak{q}_1) = N(\mathfrak{q}_2) = 3.$$

In particular, they are not principal. The ideal (5) is out the window; it has norm 25. Hence, the three cherries are $\mathfrak{p}$, $\mathfrak{q}_1$, $\mathfrak{q}_2$.

The possible ways to arrange these cherries into ideals with norm ≤ 5 are

$$\{(1), \mathfrak{p}, \mathfrak{q}_1, \mathfrak{q}_2, \mathfrak{p}^2\}.$$

However, you can compute

$$\mathfrak{p}^2 = (2)$$

so $\mathfrak{p}^2$ and (1) are in the same class group; that is, they are trivial. In particular, the class group has order at most 4.

From now on, let $[\mathfrak{a}]$ denote the class (member of the class group) that $\mathfrak{a}$ is in. Since $\mathfrak{p}$ isn't principal (so $[\mathfrak{p}] \neq [(1)]$), it follows that $\mathfrak{p}$ has order two. So Lagrange's theorem says that Cl_K has order either 2 or 4.

Now we claim $[\mathfrak{q}_1]^2 \neq [(1)]$, which implies that $\mathfrak{q}_1$ has order greater than 2. If not, $\mathfrak{q}_1^2$ is principal. We know $N(\mathfrak{q}_1) = 3$, so this can only occur if $\mathfrak{q}_1^2 = (3)$; this would force $\mathfrak{q}_1 = \mathfrak{q}_2$. This is impossible since $\mathfrak{q}_1 + \mathfrak{q}_2 = (1)$.

Thus, $\mathfrak{q}_1$ has even order greater than 2. So it has to have order 4. From this we deduce

$$\text{Cl}_K \cong \mathbb{Z}_4. \quad \square$$

Remark 38.8.10. When we did this at Harvard during Math 129, there was a five-minute interruption in which students (jokingly) complained about the difficulty of evaluating $\frac{4}{\pi} \sqrt{17}$. Excerpt:

- “Will we be allowed to bring a small calculator on the exam?” – Student 1
 “What does the size have to do with anything? You could have an Apple Watch” – Professor
 “Just use the fact that $\pi \geq 3$ ” – me
 “Even [other professor] doesn't know that, how are we supposed to?” – Student 2
 “You have to do this yourself!” – Professor
 “This is an outrage.” – Student 1

§38.9 Problems to think about

Problem 38A. Show that $K = \mathbb{Q}(\sqrt{-163})$ has trivial class group, and hence $\mathcal{O}_K = \mathbb{Z}[\sqrt{-163}]$ is a unique factorization domain.²

Problem 38B. Determine the class group of $\mathbb{Q}(\sqrt{-31})$.

Problem 38C (China TST 1998). Let n be a positive integer. A polygon in the plane (not necessarily convex) has area greater than n . Prove that one can translate it so that it contains at least $n + 1$ lattice points.

Problem 38D (Lemma 38.3.6). Consider the composition of the embeddings $K \hookrightarrow \mathbb{R}^{r_1} \times \mathbb{C}^{r_2} \xrightarrow{\sim} \mathbb{R}^n$. Show that the image of $\mathcal{O}_K \subseteq K$ has mesh equal to

$$\frac{1}{2^{r_2}} \sqrt{|\Delta_K|}.$$

Problem 38E. Let $p \equiv 1 \pmod{4}$ be a prime. Show that there are unique integers $a > b > 0$ such that $a^2 + b^2 = p$.

Problem 38F (Korea 2014). Let p be an odd prime and k a positive integer such that $p \mid k^2 + 5$. Prove that there exist positive integers m, n such that $p^2 = m^2 + 5n^2$.

² In fact, $n = 163$ is the largest number for which $\mathbb{Q}(\sqrt{-n})$ has trivial class group. The complete list is 1, 2, 3, 7, 11, 19, 43, 67, 163, the [Heegner numbers](#). You might notice Euler's prime-generating polynomial $t^2 + t + 41$ when doing the above problem. Not a coincidence!

39 More properties of the discriminant

I'll remind you that the discriminant of a number field K is given by

$$\Delta_K \stackrel{\text{def}}{=} \det \begin{bmatrix} \sigma_1(\alpha_1) & \cdots & \sigma_n(\alpha_1) \\ \vdots & \ddots & \vdots \\ \sigma_1(\alpha_n) & \cdots & \sigma_n(\alpha_n) \end{bmatrix}^2$$

where $\alpha_1, \dots, \alpha_n$ is a $\mathbb{Z}$ -basis for K , and the σ_i are the n embeddings of K into $\mathbb{C}$.

Several examples, properties, and equivalent definitions follow.

§39.1 Problems to think about

Problem 39A* (Discriminant of cyclotomic field). Let p be an odd rational prime and ζ_p a primitive p th root of unity. Let $K = \mathbb{Q}(\zeta_p)$. Show that

$$\Delta_K = (-1)^{\frac{p-1}{2}} p^{p-2}.$$



Problem 39B* (Trace representation of Δ_K). Let $\alpha_1, \dots, \alpha_n$ be a basis for $\mathcal{O}_K$. Prove that

$$\Delta_K = \det \begin{bmatrix} \text{Tr}_{K/\mathbb{Q}}(\alpha_1^2) & \text{Tr}_{K/\mathbb{Q}}(\alpha_1\alpha_2) & \cdots & \text{Tr}_{K/\mathbb{Q}}(\alpha_1\alpha_n) \\ \text{Tr}_{K/\mathbb{Q}}(\alpha_2\alpha_1) & \text{Tr}_{K/\mathbb{Q}}(\alpha_2^2) & \cdots & \text{Tr}_{K/\mathbb{Q}}(\alpha_2\alpha_n) \\ \vdots & \vdots & \ddots & \vdots \\ \text{Tr}_{K/\mathbb{Q}}(\alpha_n\alpha_1) & \text{Tr}_{K/\mathbb{Q}}(\alpha_n\alpha_2) & \cdots & \text{Tr}_{K/\mathbb{Q}}(\alpha_n\alpha_n) \end{bmatrix}.$$

In particular, Δ_K is an integer.

Problem 39C* (Root representation of Δ_K). The **discriminant** of a quadratic polynomial $Ax^2 + Bx + C$ is defined as $B^2 - 4AC$. More generally, the polynomial discriminant of a polynomial $f \in \mathbb{Z}[x]$ of degree n is

$$\Delta(f) \stackrel{\text{def}}{=} c^{2n-2} \prod_{1 \leq i < j \leq n} (z_i - z_j)^2$$

where $z_1, \dots, z_n$ are the roots of f , and c is the leading coefficient of f .

Suppose K is monogenic with $\mathcal{O}_K = \mathbb{Z}[\theta]$. Let f denote the minimal polynomial of θ (hence monic). Show that

$$\Delta_K = \Delta(f).$$

Problem 39D. Show that if $K \neq \mathbb{Q}$ is a number field then $|\Delta_K| > 1$.

Problem 39E (Brill's theorem). For a number field K with signature (r_1, r_2) , show that $\Delta_K > 0$ if and only if r_2 is even.



Problem 39F (Stickelberger theorem). Let K be a number field. Prove that

$$\Delta_K \equiv 0 \text{ or } 1 \pmod{4}.$$

40 Bonus: Let's solve Pell's equation!

This is an optional aside, and can be safely ignored. (On the other hand, it's pretty short.)

§40.1 Units

Prototypical example for this section: ± 1 , roots of unity, $3 - 2\sqrt{2}$ and its powers.

Recall according to Problem 36A* that $\alpha \in \mathcal{O}_K$ is invertible if and only if

$$N_{K/\mathbb{Q}}(\alpha) = \pm 1.$$

We let $\mathcal{O}_K^\times$ denote the set of units of $\mathcal{O}_K$.

Question 40.1.1. Show that $\mathcal{O}_K^\times$ is a group under multiplication. Hence we name it the **unit group** of $\mathcal{O}_K$.

What are some examples of units?

Example 40.1.2 (Examples of units in a number field)

1. ± 1 are certainly units, present in any number field.
2. If $\mathcal{O}_K$ contains a root of unity ω (i.e. $\omega^n = 1$), then ω is a unit. (In fact, ± 1 are special cases of this.)
3. Of course, not all units of $\mathcal{O}_K$ are roots of unity. For example, if $\mathcal{O}_K = \mathbb{Z}[\sqrt{3}]$ (from $K = \mathbb{Q}(\sqrt{3})$) then the number $2 + \sqrt{3}$ is a unit, as its norm is

$$N_{K/\mathbb{Q}}(2 + \sqrt{3}) = 2^2 - 3 \cdot 1^2 = 1.$$

Alternatively, just note that the inverse $2 - \sqrt{3} \in \mathcal{O}_K$ as well:

$$(2 - \sqrt{3})(2 + \sqrt{3}) = 1.$$

Either way, $2 - \sqrt{3}$ is a unit.

4. Given any unit $u \in \mathcal{O}_K^\times$, all its powers are also units. So for example, $(3 - 2\sqrt{2})^n$ is always a unit of $\mathbb{Z}[\sqrt{2}]$, for any n . If u is not a root of unity, then this generates infinitely many new units in $\mathcal{O}_K^\times$.

Question 40.1.3. Verify the claims above that

- (a) Roots of unity are units, and
- (b) Powers of units are units.

One can either proceed from the definition or use the characterization $N_{K/\mathbb{Q}}(\alpha) = \pm 1$. If one definition seems more natural to you, use the other.

§40.2 Dirichlet's unit theorem

Prototypical example for this section: The units of $\mathbb{Z}[\sqrt{3}]$ are $\pm(2 + \sqrt{3})^n$.

Definition 40.2.1. Let $\mu(\mathcal{O}_K)$ denote the set of roots of unity contained in a number field K (equivalently, in $\mathcal{O}_K$).

Example 40.2.2 (Examples of $\mu(\mathcal{O}_K)$)

(a) If $K = \mathbb{Q}(i)$, then $\mathcal{O}_K = \mathbb{Z}[i]$. So

$$\mu(\mathcal{O}_K) = \{\pm 1, \pm i\} \quad \text{where } K = \mathbb{Q}(i).$$

(b) If $K = \mathbb{Q}(\sqrt{3})$, then $\mathcal{O}_K = \mathbb{Z}[\sqrt{3}]$. So

$$\mu(\mathcal{O}_K) = \{\pm 1\} \quad \text{where } K = \mathbb{Q}(\sqrt{3}).$$

(c) If $K = \mathbb{Q}(\sqrt{-3})$, then $\mathcal{O}_K = \mathbb{Z}[\frac{1}{2}(1 + \sqrt{-3})]$. So

$$\mu(\mathcal{O}_K) = \left\{ \pm 1, \frac{\pm 1 \pm \sqrt{-3}}{2} \right\} \quad \text{where } K = \mathbb{Q}(\sqrt{-3})$$

where the $\pm$'s in the second term need not depend on each other; in other words $\mu(\mathcal{O}_K) = \{z \mid z^6 = 1\}$.

Exercise 40.2.3. Show that we always have that $\mu(\mathcal{O}_K)$ comprises the roots to $x^n - 1$ for some integer n . (First, show it is a finite group under multiplication.)

We now quote, without proof, the so-called Dirichlet's unit theorem, which gives us a much more complete picture of what the units in $\mathcal{O}_K$ are. Legend says that Dirichlet found the proof of this theorem during an Easter concert in the Sistine Chapel.

Theorem 40.2.4 (Dirichlet's unit theorem)

Let K be a number field with signature (r_1, r_2) and set

$$s = r_1 + r_2 - 1.$$

Then there exist units $u_1, \dots, u_s$ such that every unit $\alpha \in \mathcal{O}_K^\times$ can be written *uniquely* in the form

$$\alpha = \omega \cdot u_1^{n_1} \dots u_s^{n_s}$$

for $\omega \in \mu(\mathcal{O}_K)$ is a root of unity, and $n_1, \dots, n_s \in \mathbb{Z}$.

More succinctly:

We have $\mathcal{O}_K^\times \cong \mathbb{Z}^{r_1+r_2-1} \times \mu(\mathcal{O}_K)$.

A choice of $u_1, \dots, u_s$ is called a choice of **fundamental units**.

Here are some example applications.

Example 40.2.5 (Some unit groups)

- (a) Let $K = \mathbb{Q}(i)$ with signature $(0, 1)$. Then we obtain $s = 0$, so Dirichlet's Unit theorem says that there are no units other than the roots of unity. Thus

$$\mathcal{O}_K^\times = \{\pm 1, \pm i\} \quad \text{where } K = \mathbb{Q}(i).$$

This is not surprising, since $a + bi \in \mathbb{Z}[i]$ is a unit if and only if $a^2 + b^2 = 1$.

- (b) Let $K = \mathbb{Q}(\sqrt{3})$, which has signature $(2, 0)$. Then $s = 1$, so we expect exactly one fundamental unit. A fundamental unit is $2 + \sqrt{3}$ (or $2 - \sqrt{3}$, its inverse) with norm 1, and so we find

$$\mathcal{O}_K^\times = \left\{ \pm (2 + \sqrt{3})^n \mid n \in \mathbb{Z} \right\}.$$

- (c) Let $K = \mathbb{Q}(\sqrt[3]{2})$ with signature $(1, 1)$. Then $s = 1$, so we expect exactly one fundamental unit. The choice $1 + \sqrt[3]{2} + \sqrt[3]{4}$. So

$$\mathcal{O}_K^\times = \left\{ \pm \left(1 + \sqrt[3]{2} + \sqrt[3]{4} \right)^n \mid n \in \mathbb{Z} \right\}.$$

I haven't actually shown you that these are fundamental units, and indeed computing fundamental units is in general hard.

§40.3 Finding fundamental units

Here is a table with some fundamental units.

d	Unit
$d = 2$	$1 + \sqrt{2}$
$d = 3$	$2 + \sqrt{3}$
$d = 5$	$\frac{1}{2}(1 + \sqrt{5})$
$d = 6$	$5 + 2\sqrt{6}$
$d = 7$	$8 + 3\sqrt{7}$
$d = 10$	$3 + \sqrt{10}$
$d = 11$	$10 + 3\sqrt{11}$

In general, determining fundamental units is computationally hard.

However, once I tell you what the fundamental unit is, it's not too bad (at least in the case $s = 1$) to verify it. For example, suppose we want to show that $10 + 3\sqrt{11}$ is a fundamental unit of $K = \mathbb{Q}(\sqrt{11})$, which has ring of integers $\mathbb{Z}[\sqrt{11}]$. If not, then for some $n > 1$, we would have to have

$$10 + 3\sqrt{11} = \pm (x + y\sqrt{11})^n.$$

For this to happen, at the very least we would need $|y| < 3$. We would also have $x^2 - 11y^2 = \pm 1$. So one can just verify (using $y = 1, 2$) that this fails.

The point is that: Since $(10, 3)$ is the *smallest* (in the sense of $|y|$) integer solution to $x^2 - 11y^2 = \pm 1$, it must be the fundamental unit. This holds more generally, although in the case that $d \equiv 1 \pmod{4}$ a modification must be made as x, y might be half-integers (like $\frac{1}{2}(1 + \sqrt{5})$).

Theorem 40.3.1 (Fundamental units of pell equations)

Assume d is a squarefree integer.

- (a) If $d \equiv 2, 3 \pmod{4}$, and (x, y) is a minimal integer solution to $x^2 - dy^2 = \pm 1$, then $x + y\sqrt{d}$ is a fundamental unit.
- (b) If $d \equiv 1 \pmod{4}$, and (x, y) is a minimal *half-integer* solution to $x^2 - dy^2 = \pm 1$, then $x + y\sqrt{d}$ is a fundamental unit. (Equivalently, the minimal integer solution to $a^2 - db^2 = \pm 4$ gives $\frac{1}{2}(a + b\sqrt{d})$.)

(Any reasonable definition of “minimal” will work, such as sorting by $|y|$.)

§40.4 Pell's equation

This class of results completely eradicates Pell's Equation. After all, solving

$$a^2 - d \cdot b^2 = \pm 1$$

amounts to finding elements of $\mathbb{Z}[\sqrt{d}]$ with norm ± 1 . It's a bit weirder in the $d \equiv 1 \pmod{4}$ case, since in that case $K = \mathbb{Q}(\sqrt{d})$ gives $\mathcal{O}_K = \mathbb{Z}[\frac{1}{2}(1 + \sqrt{d})]$, and so the fundamental unit may not actually be a solution. (For example, when $d = 5$, we get the solution $(\frac{1}{2}, \frac{1}{2})$.) Nonetheless, all *integer* solutions are eventually generated.

To make this all concrete, here's a simple example.

Example 40.4.1 ($x^2 - 5y^2 = \pm 1$)

Set $K = \mathbb{Q}(\sqrt{5})$, so $\mathcal{O}_K = \mathbb{Z}[\frac{1}{2}(1 + \sqrt{5})]$. By Dirichlet's unit theorem, $\mathcal{O}_K^\times$ is generated by a single element u . The choice

$$u = \frac{1}{2} + \frac{1}{2}\sqrt{5}$$

serves as a fundamental unit, as there are no smaller integer solutions to $a^2 - 5b^2 = \pm 4$.

The first several powers of u are

n	u^n	Norm
-2	$\frac{1}{2}(3 - \sqrt{5})$	1
-1	$\frac{1}{2}(1 - \sqrt{5})$	-1
0	1	1
1	$\frac{1}{2}(1 + \sqrt{5})$	-1
2	$\frac{1}{2}(3 + \sqrt{5})$	1
3	$2 + \sqrt{5}$	-1
4	$\frac{1}{2}(7 + 3\sqrt{5})$	1
5	$\frac{1}{2}(11 + 5\sqrt{5})$	-1
6	$9 + 4\sqrt{5}$	1

One can see that the first integer solution is $(2, 1)$, which gives -1 . The first solution with $+1$ is $(9, 4)$. Continuing the pattern, we find that every third power of u gives an integer solution (see also Problem 40B), with the odd ones giving a solution to $x^2 - 5y^2 = -1$ and the even ones a solution to $x^2 - 5y^2 = +1$. All solutions are generated this way, up to $\pm$ signs (by considering $\pm u^{\pm n}$).

§40.5 Problems to think about

Problem 40A (Fictitious account of the battle of Hastings). Determine the number of soldiers in the following battle:

The men of Harold stood well together, as their wont was, and formed thirteen squares, with a like number of men in every square thereof, and woe to the hardy Norman who ventured to enter their redoubts; for a single blow of Saxon war-hatched would break his lance and cut through his coat of mail . . . when Harold threw himself into the fray the Saxons were one might square of men, shouting the battle-cries, “Ut!”, “Olicrosse!”, “Godemite!”

Problem 40B. Let $d > 0$ be a squarefree integer, and let u denote the fundamental unit of $\mathbb{Q}(\sqrt{d})$. Show that either $u \in \mathbb{Z}[\sqrt{d}]$, or $u^n \in \mathbb{Z}[\sqrt{d}] \iff 3 \mid n$.

Problem 40C. Show that there are no integer solutions to

$$x^2 - 34y^2 = -1$$

despite the fact that -1 is a quadratic residue mod 34.

XI

Algebraic NT II: Galois and Ramification Theory

41 Things Galois	385
41.1 Motivation	385
41.2 Field extensions, algebraic closures, and splitting fields	386
41.3 Embeddings into algebraic closures for number fields	387
41.4 Everyone hates characteristic 2: separable vs irreducible	388
41.5 Automorphism groups and Galois extensions	390
41.6 Fundamental theorem of Galois theory	392
41.7 Problems to think about	394
41.8 (Optional) Proof that Galois extensions are splitting	395
42 Finite fields	397
42.1 Example of a finite field	397
42.2 Finite fields have prime power order	398
42.3 All finite fields are isomorphic	399
42.4 The Galois theory of finite fields	400
42.5 Problems to think about	401
43 Ramification theory	402
43.1 Ramified / inert / split primes	402
43.2 Primes ramify if and only if they divide Δ_K	403
43.3 Inertial degrees	403
43.4 The magic of Galois extensions	404
43.5 (Optional) Decomposition and inertia groups	406
43.6 Tangential remark: more general Galois extensions	408
43.7 Problems to think about	408
44 The Frobenius element	409
44.1 Frobenius elements	409
44.2 Conjugacy classes	410
44.3 Chebotarev density theorem	412
44.4 Example: Frobenius elements of cyclotomic fields	412
44.5 Frobenius elements behave well with restriction	413
44.6 Application: Quadratic reciprocity	414
44.7 Frobenius elements control factorization	416
44.8 Example application: IMO 2003 problem 6	419
44.9 Problems to think about	420
45 Bonus: A Bit on Artin Reciprocity	421
45.1 Infinite primes	421
45.2 Modular arithmetic with infinite primes	421
45.3 Infinite primes in extensions	423
45.4 Frobenius element and Artin symbol	424
45.5 Artin reciprocity	426
45.6 Problems to think about	429

41 Things Galois

§41.1 Motivation

Prototypical example for this section: $\mathbb{Q}(\sqrt{2})$ and $\mathbb{Q}(\sqrt[3]{2})$.

The key idea in Galois theory is that of *embeddings*, which give us another way to get at the idea of the “conjugate” we described earlier.

Let K be a number field. An **embedding** $\sigma : K \hookrightarrow \mathbb{C}$, is an *injective field homomorphism*: it needs to preserve addition and multiplication, and in particular it should fix 1.

Question 41.1.1. Show that in this context, $\sigma(q) = q$ for any rational number q .

Example 41.1.2 (Examples of embeddings)

- (a) If $K = \mathbb{Q}(i)$, the two embeddings of K into $\mathbb{C}$ are $z \mapsto z$ (the identity) and $z \mapsto \bar{z}$ (complex conjugation).
- (b) If $K = \mathbb{Q}(\sqrt{2})$, the two embeddings of K into $\mathbb{C}$ are $a + b\sqrt{2} \mapsto a + b\sqrt{2}$ (the identity) and $a + b\sqrt{2} \mapsto a - b\sqrt{2}$ (conjugation).
- (c) If $K = \mathbb{Q}(\sqrt[3]{2})$, there are three embeddings:
 - The identity embedding, which sends $1 \mapsto 1$ and $\sqrt[3]{2} \mapsto \sqrt[3]{2}$.
 - An embedding which sends $1 \mapsto 1$ and $\sqrt[3]{2} \mapsto \omega \sqrt[3]{2}$, where ω is a cube root of unity. Note that this is enough to determine the rest of the embedding.
 - An embedding which sends $1 \mapsto 1$ and $\sqrt[3]{2} \mapsto \omega^2 \sqrt[3]{2}$.

I want to make several observations about these embeddings, which will form the core ideas of Galois theory. Pay attention here!

- First, you’ll notice some duality between roots: in the first example, i gets sent to $\pm i$, $\sqrt{2}$ gets sent to $\pm\sqrt{2}$, and $\sqrt[3]{2}$ gets sent to the other roots of $x^3 - 2$. This is no coincidence, and one can show this occurs in general. Specifically, suppose α has minimal polynomial

$$0 = c_n \alpha^n + c_{n-1} \alpha^{n-1} + \cdots + c_1 \alpha + c_0$$

where the c_i are rational. Then applying any embedding σ to both sides gives

$$\begin{aligned} 0 &= \sigma(c_n \alpha^n + c_{n-1} \alpha^{n-1} + \cdots + c_1 \alpha + c_0) \\ &= \sigma(c_n) \sigma(\alpha)^n + \sigma(c_{n-1}) \sigma(\alpha)^{n-1} + \cdots + \sigma(c_1) \sigma(\alpha) + \sigma(c_0) \\ &= c_n \sigma(\alpha)^n + c_{n-1} \sigma(\alpha)^{n-1} + \cdots + c_1 \sigma(\alpha) + c_0 \end{aligned}$$

where in the last step we have used the fact that $c_i \in \mathbb{Q}$, so they are fixed by σ . So, *roots of minimal polynomials go to other roots of that polynomial.*

- Next, I want to draw out a contrast between the second and third examples. Specifically, in example (b) where we consider embeddings $K = \mathbb{Q}(\sqrt{2})$ to $\mathbb{C}$. The image of these embeddings lands entirely in K : that is, we could just as well have looked at $K \rightarrow K$ rather than looking at $K \rightarrow \mathbb{C}$. However, this is not true in (c): indeed $\mathbb{Q}(\sqrt[3]{2}) \subset \mathbb{R}$, but the non-identity embeddings have complex outputs!

The key difference is to again think about conjugates. Key observation:

The field $K = \mathbb{Q}(\sqrt[3]{2})$ is “deficient” because the minimal polynomial $x^3 - 2$ has two other roots $\omega\sqrt[3]{2}$ and $\omega^2\sqrt[3]{2}$ not contained in K .

On the other hand $K = \mathbb{Q}(\sqrt{2})$ is just fine because both roots of $x^2 - 2$ are contained inside K . Finally, one can actually fix the deficiency in $K = \mathbb{Q}(\sqrt[3]{2})$ by completing it to a field $\mathbb{Q}(\sqrt[3]{2}, \omega)$. Fields like $\mathbb{Q}(i)$ or $\mathbb{Q}(\sqrt{2})$ which are “self-contained” are called *Galois extensions*, as we’ll explain shortly.

- Finally, you’ll notice that in the examples above, *the number of embeddings from K to $\mathbb{C}$ happens to be the degree of K* . This is an important theorem, Theorem 41.3.1.

In this chapter we’ll develop these ideas in full generality, for any field other than $\mathbb{Q}$.

§41.2 Field extensions, algebraic closures, and splitting fields

Prototypical example for this section: $\mathbb{Q}(\sqrt[3]{2})/\mathbb{Q}$ is an extension, $\mathbb{C}$ is an algebraic closure of any number field.

First, we define a notion of one field sitting inside another, in order to generalize the notion of a number field.

Definition 41.2.1. Let K and F be fields. If $F \subseteq K$, we write K/F and say K is a **field extension** of F .

Thus K is automatically an F -vector space (just like $\mathbb{Q}(\sqrt{2})$ is automatically a $\mathbb{Q}$ -vector space). The **degree** is the dimension of this space, denoted $[K : F]$. If $[K : F]$ is finite, we say K/F is a **finite (field) extension**.

That’s really all. There’s nothing tricky at all.

Question 41.2.2. What do you call a finite extension of $\mathbb{Q}$?

Degrees of finite extensions are multiplicative.

Theorem 41.2.3 (Field extensions have multiplicative degree)

Let $F \subseteq K \subseteq L$ be fields with L/K , K/F finite. Then

$$[L : K][K : F] = [L : F].$$

Proof. Basis bash: you can find a basis of L over K , and then expand that into a basis L over F . (Diligent readers can fill in details.) $\square$

Next, given a field (like $\mathbb{Q}(\sqrt[3]{2})$) we want something to embed it into (in our case $\mathbb{C}$). So we just want a field that contains all the roots of all the polynomials:

Theorem 41.2.4 (Algebraic closures)

Let F be a field. Then there exists a field extension $\overline{F}$ containing F , called an **algebraic closure**, such that all polynomials in $\overline{F}[x]$ factor completely.

Example 41.2.5 ($\mathbb{C}$)

$\mathbb{C}$ is an algebraic closure of $\mathbb{Q}$, $\mathbb{R}$ and even itself.

Abuse of Notation 41.2.6. Some authors also require the algebraic closure to be *minimal by inclusion*: for example, given $\mathbb{Q}$ they would want only $\overline{\mathbb{Q}}$ (the algebraic numbers). It's a theorem that such a minimal algebraic closure is unique, and so these authors will refer to *the* algebraic closure of K .

I like $\mathbb{C}$, so I'll use the looser definition.

§41.3 Embeddings into algebraic closures for number fields

Now that I've defined all these ingredients, I can prove:

Theorem 41.3.1 (The n embeddings of a number field)

Let K be a number field of degree n . Then there are exactly n field homomorphisms $K \hookrightarrow \mathbb{C}$, say $\sigma_1, \dots, \sigma_n$ which fix $\mathbb{Q}$.

Remark 41.3.2. Note that a nontrivial homomorphism of fields is necessarily injective (the kernel is an ideal). This justifies the use of " $\hookrightarrow$ ", and we call each σ_i an **embedding** of K into $\mathbb{C}$.

Proof. This is actually kind of fun! Recall that any irreducible polynomial over $\mathbb{Q}$ has distinct roots (Lemma 36.4.2). We'll adjoin elements $\alpha_1, \alpha_2, \dots, \alpha_n$ one at a time to $\mathbb{Q}$, until we eventually get all of K , that is,

$$K = \mathbb{Q}(\alpha_1, \dots, \alpha_n).$$

Diagrammatically, this is

$$\begin{array}{ccccccc}
 \mathbb{Q} & \hookrightarrow & \mathbb{Q}(\alpha_1) & \hookrightarrow & \mathbb{Q}(\alpha_1, \alpha_2) & \hookrightarrow & \dots \hookrightarrow K \\
 \downarrow \text{id} & & \downarrow \tau_1 & & \downarrow \tau_2 & & \downarrow \tau_n = \sigma \\
 \mathbb{C} & \longrightarrow & \mathbb{C} & \longrightarrow & \mathbb{C} & \longrightarrow & \dots \longrightarrow \mathbb{C}
 \end{array}$$

First, we claim there are exactly

$$[\mathbb{Q}(\alpha_1) : \mathbb{Q}]$$

ways to pick τ_1 . Observe that τ_1 is determined by where it sends α_1 (since it has to fix $\mathbb{Q}$). Letting p_1 be the minimal polynomial of α_1 , we see that there are $\deg p_1$ choices for τ_1 , one for each (distinct) root of p_1 . That proves the claim.

Similarly, given a choice of τ_1 , there are

$$[\mathbb{Q}(\alpha_1, \alpha_2) : \mathbb{Q}(\alpha_1)]$$

ways to pick τ_2 . (It's a little different: τ_1 need not be the identity. But it's still true that τ_2 is determined by where it sends α_2 , and as before there are $[\mathbb{Q}(\alpha_1, \alpha_2) : \mathbb{Q}(\alpha_1)]$ possible ways.)

Multiplying these all together gives the desired $[K : \mathbb{Q}]$. $\square$

Remark 41.3.3. A shorter proof can be given using the primitive element theorem. Do you see how?

It's common to see expressions like “let K be a number field of degree n , and $\sigma_1, \dots, \sigma_n$ its n embeddings” without further explanation. The relation between these embeddings and the Galois conjugates is given as follows.

Theorem 41.3.4 (Embeddings are evenly distributed over conjugates)

Let K be a number field of degree n with n embeddings $\sigma_1, \dots, \sigma_n$, and let $\alpha \in K$ have m Galois conjugates over $\mathbb{Q}$.

Then $\sigma_j(\alpha)$ is “evenly distributed” over each of these m conjugates: for any Galois conjugate β , exactly $\frac{n}{m}$ of the embeddings send α to β .

Proof. In the previous proof, adjoin $\alpha_1 = \alpha$ first. $\square$

So, now we can define the trace and norm over $\mathbb{Q}$ in a nice way: given a number field K , we set

$$\mathrm{Tr}_{K/\mathbb{Q}}(\alpha) = \sum_{i=1}^n \sigma_i(\alpha) \quad \text{and} \quad \mathrm{N}_{K/\mathbb{Q}}(\alpha) = \prod_{i=1}^n \sigma_i(\alpha)$$

where σ_i are the n embeddings of K into $\mathbb{C}$.

§41.4 Everyone hates characteristic 2: separable vs irreducible

Prototypical example for this section: $\mathbb{Q}$ has characteristic zero, hence irreducible polynomials are separable.

Now, we want a version of the above theorem for any field F . If you read the proof, you'll see that the only thing that ever uses anything about the field $\mathbb{Q}$ is Lemma 36.4.2, where we use the fact that

Irreducible polynomials over F have no double roots.

Let's call a polynomial with no double roots **separable**; thus we want irreducible polynomials to be separable. We did this for $\mathbb{Q}$ in the last chapter by taking derivatives. Should work for any field, right?

Nope. Suppose we took the derivative of some polynomial like $2x^3 + 24x + 9$, namely $6x^2 + 24$. In $\mathbb{C}$ it's obvious that the derivative of a nonconstant polynomial f' isn't zero. But suppose we considered the above as a polynomial in $\mathbb{F}_3$, i.e. modulo 3. Then the derivative is zero. Oh, no!

We have to impose a condition that prevents something like this from happening.

Definition 41.4.1. For a field F , the **characteristic** of F is the smallest positive integer p such that,

$$\underbrace{1_F + \dots + 1_F}_{p \text{ times}} = 0$$

or zero if no such integer p exists.

Example 41.4.2 (Field characteristics)

Old friends $\mathbb{R}$, $\mathbb{Q}$, $\mathbb{C}$ all have characteristic zero. But $\mathbb{F}_p$, the integers modulo p , is a field of characteristic p .

Exercise 41.4.3. Let F be a field of characteristic p . Show that if $p > 0$ then p is a prime number. (A proof is given next chapter.)

With the assumption of characteristic zero, our earlier proof works.

Lemma 41.4.4 (Separability in characteristic zero)

Any irreducible polynomial in a characteristic zero field is separable.

Unfortunately, this lemma is false if the “characteristic zero” condition is dropped.

Remark 41.4.5. The reason it’s called *separable* is (I think) this picture: I have a polynomial and I want to break it into irreducible parts. Normally, if I have a double root in a polynomial, that means it’s not irreducible. But in characteristic $p > 0$ this fails. So inseparable polynomials are strange when you think about them: somehow you have double roots that can’t be separated from each other.

We can get this to work for any field extension in which separability is not an issue.

Definition 41.4.6. A **separable extension** K/F is one in which every irreducible polynomial in F is separable (for example, if F has characteristic zero). A field F is **perfect** if any finite field extension K/F is separable.

In fact, as we see in the next chapter:

Theorem 41.4.7 (Finite fields are perfect)

Suppose F is a field with finitely many elements. Then it is perfect.

Thus, we will almost never have to worry about separability since every field we see in the Napkin is either finite or characteristic 0. So the inclusion of the word “separable” is mostly a formality.

Proceeding onwards, we obtain

Theorem 41.4.8 (The n embeddings of any separable extension)

Let K/F be a separable extension of degree n and let $\overline{F}$ be an algebraic closure of F . Then there are exactly n field homomorphisms $K \hookrightarrow \overline{F}$, say $\sigma_1, \dots, \sigma_n$, which fix F .

In any case, this lets us define the trace for *any* separable normal extension.

Definition 41.4.9. Let K/F be a separable extension of degree n , and let $\sigma_1, \dots, \sigma_n$ be the n embeddings into an algebraic closure of F . Then we define

$$\mathrm{Tr}_{K/F}(\alpha) = \sum_{i=1}^n \sigma_i(\alpha) \quad \text{and} \quad \mathrm{N}_{K/F}(\alpha) = \prod_{i=1}^n \sigma_i(\alpha).$$

When $F = \mathbb{Q}$ and the algebraic closure is $\mathbb{C}$, this coincides with our earlier definition!

§41.5 Automorphism groups and Galois extensions

Prototypical example for this section: $\mathbb{Q}(\sqrt{2})$ is Galois but $\mathbb{Q}(\sqrt[3]{2})$ is not.

We now want to get back at the idea we stated at the beginning of this section that $\mathbb{Q}(\sqrt[3]{2})$ is deficient in a way that $\mathbb{Q}(\sqrt{2})$ is not.

First, we define the “internal” automorphisms.

Definition 41.5.1. Suppose K/F is a finite extension. Then $\text{Aut}(K/F)$ is the set of field isomorphisms $\sigma : K \rightarrow K$ which fix F . In symbols

$$\text{Aut}(K/F) = \{\sigma : K \rightarrow K \mid \sigma \text{ is identity on } F\}.$$

This is a group under function composition!

Note that this time, we have a condition that F is fixed by σ . (This was not there before when we considered $F = \mathbb{Q}$, because we got it for free.)

Example 41.5.2 (Old examples of automorphism groups)

Reprising the example at the beginning of the chapter in the new notation, we have:

- (a) $\text{Aut}(\mathbb{Q}(i)/\mathbb{Q}) \cong \mathbb{Z}_2$, with elements $z \mapsto z$ and $z \mapsto \bar{z}$.
- (b) $\text{Aut}(\mathbb{Q}(\sqrt{2})/\mathbb{Q}) \cong \mathbb{Z}_2$ in the same way.
- (c) $\text{Aut}(\mathbb{Q}(\sqrt[3]{2})/\mathbb{Q})$ is the trivial group, with only the identity embedding!

Example 41.5.3 (Automorphism group of $\mathbb{Q}(\sqrt{2}, \sqrt{3})$)

Here’s a new example: let $K = \mathbb{Q}(\sqrt{2}, \sqrt{3})$. It turns out that $\text{Aut}(K/\mathbb{Q}) = \{1, \sigma, \tau, \sigma\tau\}$, where

$$\sigma : \begin{cases} \sqrt{2} & \mapsto -\sqrt{2} \\ \sqrt{3} & \mapsto \sqrt{3} \end{cases} \quad \text{and} \quad \tau : \begin{cases} \sqrt{2} & \mapsto \sqrt{2} \\ \sqrt{3} & \mapsto -\sqrt{3} \end{cases}.$$

In other words, $\text{Aut}(K/\mathbb{Q})$ is the Klein Four Group.

First, let’s repeat the proof of the observation that these embeddings shuffle around roots (akin to the first observation in the introduction):

Lemma 41.5.4 (Root shuffling in $\text{Aut}(K/F)$)

Let $f \in F[x]$, suppose K/F is a finite extension, and assume $\alpha \in K$ is a root of f . Then for any $\sigma \in \text{Aut}(K/F)$, $\sigma(\alpha)$ is also a root of f .

Proof. Let $f(x) = c_n x^n + c_{n-1} x^{n-1} + \cdots + c_0$, where $c_i \in F$. Thus,

$$0 = \sigma(f(\alpha)) = \sigma(c_n \alpha^n + \cdots + c_0) = c_n \sigma(\alpha)^n + \cdots + c_0 = f(\sigma(\alpha)). \quad \square$$

In particular, taking f to be the minimal polynomial of α we deduce

An embedding $\sigma \in \text{Aut}(K/F)$ sends an $\alpha \in K$ to one of its various Galois conjugates (over F).

Next, let's look again at the “deficiency” of certain fields. Look at $K = \mathbb{Q}(\sqrt[3]{2})$. So, again $K/\mathbb{Q}$ is deficient for two reasons. First, while there are three maps $\mathbb{Q}(\sqrt[3]{2}) \hookrightarrow \mathbb{C}$, only one of them lives in $\text{Aut}(K/\mathbb{Q})$, namely the identity. In other words, $|\text{Aut}(K/\mathbb{Q})|$ is *too small*. Secondly, K is missing some Galois conjugates ($\omega\sqrt[3]{2}$ and $\omega^2\sqrt[3]{2}$).

The way to capture the fact that there are missing Galois conjugates is the notion of a splitting field.

Definition 41.5.5. Let F be a field and $p(x) \in F[x]$ a polynomial of degree n . Then $p(x)$ has roots $\alpha_1, \dots, \alpha_n$ in an algebraic closure of F . The **splitting field** of F is defined as $F(\alpha_1, \dots, \alpha_n)$.

In other words, the splitting field is the smallest field in which $p(x)$ splits.

Example 41.5.6 (Examples of splitting fields)

- (a) The splitting field of $x^2 - 5$ over $\mathbb{Q}$ is $\mathbb{Q}(\sqrt{5})$. This is a degree 2 extension.
- (b) The splitting field of $x^2 + x + 1$ over $\mathbb{Q}$ is $\mathbb{Q}(\omega)$, where ω is a cube root of unity. This is a degree 3 extension.
- (c) The splitting field of $x^2 + 3x + 2 = (x + 1)(x + 2)$ is just $\mathbb{Q}$! There's nothing to do.

Example 41.5.7 (Splitting fields: a cautionary tale)

The splitting field of $x^3 - 2$ over $\mathbb{Q}$ is in fact

$$\mathbb{Q}(\sqrt[3]{2}, \omega)$$

and not just $\mathbb{Q}(\sqrt[3]{2})$! One must really adjoin *all* the roots, and it's not necessarily the case that these roots will generate each other.

To be clear:

- For $x^2 - 5$, we adjoin $\sqrt{5}$ and this will automatically include $-\sqrt{5}$.
- For $x^2 + x + 1$, we adjoin ω and get the other root ω^2 for free.
- But for $x^3 - 2$, if we adjoin $\sqrt[3]{2}$, we do NOT get $\omega\sqrt[3]{2}$ and $\omega^2\sqrt[3]{2}$ for free. Indeed, $\mathbb{Q}(\sqrt[3]{2}) \subset \mathbb{R}$!

Note that in particular, the splitting field of $x^3 - 2$ over $\mathbb{Q}$ is *degree six*, not just degree three.

In general, **the splitting field of a polynomial can be an extension of degree up to $n!$** . The reason is that if $p(x)$ has n roots and none of them are “related” to each other, then any permutation of the roots will work.

Now, we obtain:

Theorem 41.5.8 (Galois extensions are splitting)

For finite extensions K/F , $|\text{Aut}(K/F)|$ divides $[K : F]$, with equality if and only if K is the *splitting field* of some separable polynomial with coefficients in F .

The proof of this is deferred to an optional section at the end of the chapter. If K/F is a finite extension and $|\text{Aut}(K/F)| = [K : F]$, we say the extension K/F is **Galois**. In that case, we denote $\text{Aut}(K/F)$ by $\text{Gal}(K/F)$ instead and call this the **Galois group** of K/F .

Example 41.5.9 (Examples and non-examples of Galois extensions)

- (a) The extension $\mathbb{Q}(\sqrt{2})/\mathbb{Q}$ is Galois, since it's the splitting field of $x^2 - 2$ over $\mathbb{Q}$. The Galois group has order two, $\sqrt{2} \mapsto \pm\sqrt{2}$.
- (b) The extension $\mathbb{Q}(\sqrt{2}, \sqrt{3})/\mathbb{Q}$ is Galois, since it's the splitting field of $(x^2 - 5)^2 - 6$ over $\mathbb{Q}$. As discussed before, the Galois group is $\mathbb{Z}_2 \times \mathbb{Z}_2$.
- (c) The extension $\mathbb{Q}(\sqrt[3]{2})/\mathbb{Q}$ is *not* Galois.

To explore $\mathbb{Q}(\sqrt[3]{2})$ one last time:

Example 41.5.10 (Galois closures, and the automorphism group of $\mathbb{Q}(\sqrt[3]{2}, \omega)$)

Let's return to the field $K = \mathbb{Q}(\sqrt[3]{2}, \omega)$, which is a field with $[K : \mathbb{Q}] = 6$. Consider the two automorphisms:

$$\sigma : \begin{cases} \sqrt[3]{2} & \mapsto \omega \sqrt[3]{2} \\ \omega & \mapsto \omega \end{cases} \quad \text{and} \quad \tau : \begin{cases} \sqrt[3]{2} & \mapsto \sqrt[3]{2} \\ \omega & \mapsto \omega^2. \end{cases}$$

Notice that $\sigma^3 = \tau^2 = \text{id}$. From this one can see that the automorphism group of K must have order 6 (it certainly has order ≤ 6 ; now use Lagrange's theorem). So, $K/\mathbb{Q}$ is Galois! Actually one can check explicitly that

$$\text{Gal}(K/\mathbb{Q}) \cong S_3$$

is the symmetric group on 3 elements, with order $3! = 6$.

This example illustrates the fact that given a non-Galois field extension, one can “add in” missing conjugates to make it Galois. This is called taking a **Galois closure**.

§41.6 Fundamental theorem of Galois theory

After all this stuff about Galois Theory, I might as well tell you the fundamental theorem, though I won't prove it. Basically, it says that if K/F is Galois with Galois group G , then:

Subgroups of G correspond exactly to fields E with $F \subseteq E \subseteq K$.

To tell you how the bijection goes, I have to define a fixed field.

Definition 41.6.1. Let K be a field and H a subgroup of $\text{Aut}(K/F)$. We define the **fixed field** of H , denoted K^H , as

$$K^H \stackrel{\text{def}}{=} \{x \in K : \sigma(x) = x \ \forall \sigma \in H\}.$$

Question 41.6.2. Verify quickly that K^H is actually a field.


Now let's look at examples again. Consider $K = \mathbb{Q}(\sqrt{2}, \sqrt{3})$, where

$$G = \text{Gal}(K/\mathbb{Q}) = \{\text{id}, \sigma, \tau, \sigma\tau\}$$

is the Klein four group (where $\sigma(\sqrt{2}) = -\sqrt{2}$ but $\sigma(\sqrt{3}) = \sqrt{3}$; τ goes the other way).

Question 41.6.3. Let $H = \{\text{id}, \sigma\}$. What is K^H ?

In that case, the diagram of fields between $\mathbb{Q}$ and K matches exactly with the subgroups of G , as follows:



We see that subgroups correspond to fixed fields. That, and much more, holds in general.

Theorem 41.6.4 (Fundamental theorem of Galois theory)

Let K/F be a Galois extension with Galois group $G = \text{Gal}(K/F)$.

(a) There is a bijection between field towers $F \subseteq E \subseteq K$ and subgroups $H \subseteq G$:

$$\left\{ \begin{array}{c} K \\ | \\ E \\ | \\ F \end{array} \right\} \Longleftrightarrow \left\{ \begin{array}{c} 1 \\ | \\ H \\ | \\ G \end{array} \right\}$$

The bijection sends H to its fixed field K^H , and hence is inclusion reversing.

(b) Under this bijection, we have $[K : E] = |H|$ and $[E : F] = [G : H]$.

(c) K/E is always Galois, and its Galois group is $\text{Gal}(K/E) = H$.

(d) E/F is Galois if and only if H is normal in G . If so, $\text{Gal}(E/F) = G/H$.

Exercise 41.6.5. Suppose we apply this theorem for

$$K = \mathbb{Q}(\sqrt[3]{2}, \omega).$$

Verify that the fact $E = \mathbb{Q}(\sqrt[3]{2})$ is not Galois corresponds to the fact that S_3 does not have normal subgroups of order 2.

§41.7 Problems to think about

Problem 41A* (Galois group of the cyclotomic field). Let p be an odd rational prime and ζ_p a primitive p th root of unity. Let $K = \mathbb{Q}(\zeta_p)$. Show that

$$\text{Gal}(K/\mathbb{Q}) \cong (\mathbb{Z}/p\mathbb{Z})^*.$$

Problem 41B (Greek constructions). Prove that the three Greek constructions

- (a) doubling the cube,
- (b) squaring the circle, and
- (c) trisecting an angle

are all impossible. (Assume π is transcendental.)

Problem 41C (Hong Kong Olympiad). Prove that there are no rational numbers p, q, r satisfying

$$\cos\left(\frac{2\pi}{7}\right) = p + \sqrt{q} + \sqrt[3]{r}.$$

Problem 41D. Show that the only automorphism of $\mathbb{R}$ is the identity. Hence $\text{Aut}(\mathbb{R})$ is the trivial group.



Problem 41E (Artin's primitive element theorem). Let K be a number field. Show that $K \cong \mathbb{Q}(\gamma)$ for some γ .

§41.8 (Optional) Proof that Galois extensions are splitting

We prove Theorem 41.5.8. First, we extract a useful fragment from the fundamental theorem.

Theorem 41.8.1 (Fixed field theorem)

Let K be a field and G a subgroup of $\text{Aut}(K)$. Then $[K : K^G] = |G|$.

The inequality itself is not difficult:

Exercise 41.8.2. Show that $[K : F] \geq |\text{Aut}(K/F)|$, and that equality holds if and only if the set of elements fixed by all $\sigma \in \text{Aut}(K/F)$ is exactly F . (Use Theorem 41.8.1.)

The equality case is trickier.

The easier direction is when K is a splitting field. Assume $K = F(\alpha_1, \dots, \alpha_n)$ is the splitting field of some separable polynomial $p \in F[x]$ with n distinct roots $\alpha_1, \dots, \alpha_n$. Adjoin them one by one:

$$\begin{array}{ccccccc}
 F & \hookrightarrow & F(\alpha_1) & \hookrightarrow & F(\alpha_1, \alpha_2) & \hookrightarrow & \dots \hookrightarrow K \\
 \text{id} \downarrow & & \tau_1 \downarrow & & \tau_2 \downarrow & & \dots \downarrow \tau_n = \sigma \\
 F & \hookrightarrow & F(\alpha_1) & \hookrightarrow & F(\alpha_1, \alpha_2) & \hookrightarrow & \dots \hookrightarrow K
 \end{array}$$

(Does this diagram look familiar?) Every map $K \rightarrow K$ which fixes F corresponds to an above commutative diagram. As before, there are exactly $[F(\alpha_1) : F]$ ways to pick τ_1 . (You need the fact that the minimal polynomial p_1 of α_1 is separable for this: there need to be exactly $\deg p_1 = [F(\alpha_1) : F]$ distinct roots to nail p_1 into.) Similarly, given a choice of τ_1 , there are $[F(\alpha_1, \alpha_2) : F(\alpha_1)]$ ways to pick τ_2 . Multiplying these all together gives the desired $[K : F]$.

Now assume K/F is Galois. First, we state:

Lemma 41.8.3

Let K/F be Galois, and $p \in F[x]$ irreducible. If any root of p (in $\overline{F}$) lies in K , then all of them do, and in fact p is separable.

Proof. Let $\alpha \in K$ be the prescribed root. Consider the set

$$S = \{\sigma(\alpha) \mid \sigma \in \text{Gal}(K/F)\}.$$

(Note that $\alpha \in S$ since $\text{Gal}(K/F) \ni \text{id}$.) By construction, any $\tau \in \text{Gal}(K/F)$ fixes S . So if we construct

$$\tilde{p}(x) = \prod_{\beta \in S} (x - \beta),$$

then by Vieta's Formulas, we find that all the coefficients of $\tilde{p}$ are fixed by elements of σ . By the *equality case* we specified in the exercise, it follows that $\tilde{p}$ has coefficients in F ! (This is where we use the condition.) Also, by Lemma 41.5.4, $\tilde{p}$ divides p .

Yet p was irreducible, so it is the minimal polynomial of α in $F[x]$, and therefore we must have that p divides $\tilde{p}$. Hence $p = \tilde{p}$. Since $\tilde{p}$ was built to be separable, so is p . $\square$

Now we're basically done – pick a basis $\omega_1, \dots, \omega_n$ of K/F , and let p_i be their minimal polynomials; by the above, we don't get any roots outside K . Consider $P = p_1 \dots p_n$, removing any repeated factors. The roots of P are $\omega_1, \dots, \omega_n$ and some other guys in K . So K is the splitting field of P .

DRAFT (Evan Chen)
Updated February 15, 2017

42 Finite fields

In this short chapter, we classify all fields with finitely many elements and compute the Galois groups. Nothing in here is very hard, and so most of the proofs are just sketches; if you like, you should check the details yourself.

The whole point of this chapter is to prove:

- A finite field F must have order p^n , with p prime and n an integer.
- In this case, F has characteristic p .
- The extension $F/\mathbb{F}_p$ is Galois, and $\text{Gal}(F/\mathbb{F}_p)$ is a cyclic group of order n . The generator is the automorphism

$$\sigma : F \rightarrow F \quad \text{by} \quad x \mapsto x^p.$$

If you're in a hurry you can just remember these results and skip to the next chapter.

§42.1 Example of a finite field

Before diving in, we give some examples.

Recall that the *characteristic* of a field F is the smallest positive integer p such that

$$\underbrace{1_F + \cdots + 1_F}_{p \text{ times}} = 0$$

or 0 if no such integer p exists.

Example 42.1.1 (Base field)

Let $\mathbb{F}_p$ denote the field of integers modulo p . This is a field with p elements, with characteristic p .

Example 42.1.2 (The finite field of nine elements)

Let

$$F \cong \mathbb{F}_3[X]/(X^2 + 1) \cong \mathbb{Z}[i]/(3).$$

We can think of its elements as

$$\{a + bi \mid 0 \leq a, b \leq 2\}.$$

Since (3) is prime in $\mathbb{Z}[i]$, the ring of integers of $\mathbb{Q}(i)$, we see F is a field with $3^2 = 9$ elements inside it. Note that, although this field has 9 elements, every element x has the property that

$$3x = \underbrace{x + \cdots + x}_{3 \text{ times}} = 0.$$

In particular, F has characteristic 3.

§42.2 Finite fields have prime power order

Lemma 42.2.1

If the characteristic of F isn't zero, it must be a prime number.

Proof. Assume not, so $n = ab$ for $a, b < n$. Then let

$$A = \underbrace{1_F + \cdots + 1_F}_{a \text{ times}} \neq 0$$

and

$$B = \underbrace{1_F + \cdots + 1_F}_{b \text{ times}} \neq 0.$$

Then $AB = 0$, contradicting the fact that F is a field. $\square$

We like fields of characteristic zero, but unfortunately for finite fields we are doomed to have nonzero characteristic.

Lemma 42.2.2 (Finite fields have prime power orders)

Let F be a finite field of characteristic p . Then

- (a) p is not zero, and hence prime.
- (b) F is a finite extension of $\mathbb{F}_p$, and in particular it is an $\mathbb{F}_p$ -vector space.
- (c) $|F| = p^n$ for some prime p , integer n .

Proof. Very briefly, since this is easy:

- (a) Apply Lagrange's theorem (or pigeonhole principle!) to $(F, +)$ to get the characteristic isn't zero.
- (b) The additive subgroup of $(F, +)$ generated by 1_F is an isomorphic copy of $\mathbb{F}_p$.
- (c) Since it's a field extension, F is a finite-dimensional vector space over $\mathbb{F}_p$, with some basis $e_1, \dots, e_n$. It follows that there are p^n elements of F . $\square$

Remark 42.2.3. An amusing alternate proof of (c) by contradiction: if a prime $q \neq p$ divides $|F|$, then by Cauchy's theorem (Problem 11A*) on $(F, +)$ there's a (nonzero) element x of order q . Evidently

$$x \cdot \underbrace{(1_F + \cdots + 1_F)}_{q \text{ times}} = 0$$

then, but $x \neq 0$, and hence the characteristic of F also divides q , which is impossible.

An important point in the above proof is that

Lemma 42.2.4 (Finite fields are field extensions of $\mathbb{F}_p$)

If $|F| = p^n$ is a finite field, then there is an isomorphic copy of $\mathbb{F}_p$ sitting inside F . Thus F is a field extension of $\mathbb{F}_p$.

We want to refer a lot to this copy of $\mathbb{F}_p$, so in what follows:

Abuse of Notation 42.2.5. Every integer n can be identified as an element of F , namely

$$n \stackrel{\text{def}}{=} \underbrace{1_F + \cdots + 1_F}_{n \text{ times}}.$$

Note that (as expected) this depends only on $n \pmod{p}$.

Convince yourself that this gives:

Theorem 42.2.6 (Freshman's dream)

For any $a, b \in F$ we have

$$(a + b)^p = a^p + b^p.$$

Proof. Use the Binomial theorem, and the fact that $\binom{p}{i}$ is divisible by p for $0 < i < p$. $\square$

Exercise 42.2.7. Convince yourself that this proof works.

§42.3 All finite fields are isomorphic

We next proceed to prove “Fermat’s little theorem”:

Theorem 42.3.1 (Fermat's little theorem in finite fields)

Let F be a finite field of order p^n . Then every element $x \in F$ satisfies

$$x^{p^n} - x = 0.$$

Proof. If $x = 0$ it’s true; otherwise, use Lagrange’s theorem on the abelian group $(F, \times)$ to get $x^{p^n-1} = 1_F$. $\square$

We can now prove:

Theorem 42.3.2 (Complete classification of finite fields)

A field F is a finite field with p^n elements if and only if it is a splitting field of $x^{p^n} - x$ over $\mathbb{F}_p$.

Proof. By “Fermat’s little theorem”, all the elements of F satisfy this polynomial. So we just have to show that the roots of this polynomial are distinct (i.e. that it is separable).

To do this, we use the derivative trick again: the derivative of this polynomial is

$$p^n \cdot x^{p^n-1} - 1 = -1$$

which has no roots at all, so the polynomial cannot have any double roots. $\square$

Note that the polynomial $x^{p^n} - x \pmod{p}$ is far from irreducible, but the computation above shows that it’s separable.

Example 42.3.3 (The finite field of order nine again)

The polynomial $x^9 - x$ is separable modulo 3 and has factorization

$$x(x+1)(x+2)(x^2+1)(x^2+x+2)(x^2+2x+2) \pmod{3}.$$

So if F has order 9, then we intuitively expect it to be the field generated by adjoining all the roots: 0, 1, 2, as well as $\pm i$, $1 \pm i$, $2 \pm i$. Indeed, that's the example we had at the beginning of this chapter.

This is a little imprecise since the complex number i doesn't really make sense (just what does " $i \pmod{3}$ " mean?), but it carries the right idea.

§42.4 The Galois theory of finite fields

Let F be a finite field of p^n elements. By the above theorem, it's the splitting field of a separable polynomial, hence we know that $F/\mathbb{F}_p$ is a Galois extension. We would like to find the Galois group.

In fact, we are very lucky: it is cyclic. First, we exhibit one such element $\sigma_p \in \text{Gal}(F/\mathbb{F}_p)$:

Theorem 42.4.1 (The p th power automorphism)

The map $\sigma_p : F \rightarrow F$ defined by

$$\sigma_p(x) : x \mapsto x^p$$

is an automorphism, and moreover fixes $\mathbb{F}_p$.

Proof. It's a homomorphism since it fixes 1_F , respects multiplication, and respects addition.

Question 42.4.2. Why does it respect addition?

Next, we claim that it is injective. To see this, note that

$$x^p = y^p \iff x^p - y^p = 0 \iff (x - y)^p = 0 \iff x = y.$$

Here we have again used the Freshman's Dream. Since F is finite, this injective map is automatically bijective. The fact that it fixes $\mathbb{F}_p$ is Fermat's little theorem. $\square$

Now we're done:

Theorem 42.4.3 (Galois group of the extension $F/\mathbb{F}_p$)

Let F have size p^n . Then $\text{Gal}(F/\mathbb{F}_p) \cong \mathbb{Z}_n$ with generator σ_p .

Proof. Since $[F : \mathbb{F}_p] = n$, the Galois group G has order n . So we just need to show $\sigma_p \in G$ has order n .

Note that σ_p applied k times gives $x \mapsto x^{p^k}$. Hence, σ_p applied n times is the identity, as all elements of F satisfy $x^{p^n} = x$. But if $k < n$, then σ_p applied k times cannot be the identity or $x^{p^k} - x$ would have too many roots. $\square$

We can see an example of this again with the finite field of order 9.

Example 42.4.4 (Galois group of finite field of order 9)

Let F be the finite field of order 9, represented as $K = \mathbb{Z}[i]/(3)$. Let $\sigma_3 : F \rightarrow F$ be $x \mapsto x^3$. We can witness the fate of all nine elements:

$$\begin{array}{ccccccc}
 0 & 1 & 2 & i & 1+i & 2+i & \\
 & & & \uparrow & \uparrow & \uparrow & \\
 & & & \sigma & \sigma & \sigma & \\
 & & & \downarrow & \downarrow & \downarrow & \\
 & & & -i & 1-i & 2-i &
 \end{array}$$

(As claimed, 0, 1, 2 are the fixed points, so I haven't drawn arrows for them.) As predicted, the Galois group has order two:

$$\text{Gal}(K/\mathbb{F}_3) = \{\text{id}, \sigma_3\} \cong \mathbb{Z}_2.$$

This concludes the proof of all results stated at the beginning of this chapter.

§42.5 Problems to think about

43 Ramification theory

We're very interested in how rational primes p factor in a bigger number field K . Some examples of this behavior: in $\mathbb{Z}[i]$ (which is a UFD!), we have factorizations

$$\begin{aligned}(2) &= (1+i)^2 \\ (3) &= (3) \\ (5) &= (2+i)(2-i).\end{aligned}$$

In this chapter we'll learn more about how primes break down when they're thrown into bigger number fields. Using weapons from Galois Theory, this will culminate in a proof of Quadratic Reciprocity.

§43.1 Ramified / inert / split primes

Prototypical example for this section: In $\mathbb{Z}[i]$, 2 is ramified, 3 is inert, and 5 splits.

Let p be a rational prime, and toss it into $\mathcal{O}_K$. Thus we get a factorization into prime ideals

$$p \cdot \mathcal{O}_K = \mathfrak{p}_1^{e_1} \cdots \mathfrak{p}_g^{e_g}.$$

We say that each $\mathfrak{p}_i$ is **above** (p) .¹ Pictorially, you might draw this as follows:

$$\begin{array}{ccccc} K & & \supset & \mathcal{O}_K & \mathfrak{p}_i \\ | & & & | & | \\ \mathbb{Q} & & \supset & \mathbb{Z} & (p) \end{array}$$

Some names for various behavior that can happen:

- We say p is **ramified** if $e_i > 1$ for some i . For example 2 is ramified in $\mathbb{Z}[i]$.
- We say p is **inert** if $g = 1$ and $e_1 = 1$; i.e. (p) remains prime. For example 3 is inert in $\mathbb{Z}[i]$.
- We say p is **split** if $g > 1$. For example 5 is split in $\mathbb{Z}[i]$.

Question 43.1.1. More generally, for a prime p in $\mathbb{Z}[i]$:

- p is ramified exactly when $p = 2$.
- p is inert exactly when $p \equiv 3 \pmod{4}$.
- p is split exactly when $p \equiv 1 \pmod{4}$.

Prove this.

¹Reminder that $p \cdot \mathcal{O}_K$ and (p) mean the same thing, and I'll use both interchangeably.

§43.2 Primes ramify if and only if they divide Δ_K

The most unusual case is ramification: Just like we don't expect a randomly selected polynomial to have a double root, we don't expect a randomly selected prime to be ramified. In fact, the key to understanding ramification is the discriminant.

For the sake of discussion, let's suppose that K is monogenic, $\mathcal{O}_K = \mathbb{Z}[\theta]$, where θ has minimal polynomial f . Let p be a rational prime we'd like to factor. If f factors as $f_1^{e_1} \dots f_g^{e_g}$, then we know that the prime factorization of (p) is given by

$$p \cdot \mathcal{O}_K = \prod_i (p, f_i(\theta))^{e_i}.$$

In particular, p ramifies exactly when f has a double root mod p ! To detect whether this happens, we look at the polynomial discriminant of f , namely

$$\Delta(f) = \prod_{i < j} (z_i - z_j)^2$$

and see whether it is zero mod p – thus p ramifies if and only if this is true.

It turns out that the naïve generalization to any number field works if we replace $\Delta(f)$ by just the discriminant Δ_K of K ; these are the same by Problem 39C*. That is,

Theorem 43.2.1 (Discriminant detects ramification)

Let p be a rational prime and K a number field. Then p is ramified if and only if p divides Δ_K .

Example 43.2.2 (Ramification in the Gaussian integers)

Let $K = \mathbb{Q}(i)$ so $\mathcal{O}_K = \mathbb{Z}[i]$ and $\Delta_K = 4$. As predicted, the only prime ramifying in $\mathbb{Z}[i]$ is 2, the only prime factor of Δ_K .

In particular, only finitely many primes ramify.

§43.3 Inertial degrees

Prototypical example for this section: (7) has inertial degree 2 in $\mathbb{Z}[i]$ and $(2+i)$ has inertial degree 1 in $\mathbb{Z}[i]$.

Recall that we were able to define an ideal norm $N(\mathfrak{a}) = |\mathcal{O}_K/\mathfrak{a}|$ measuring how “roomy” the ideal $\mathfrak{a}$ is. For example, (5) has ideal norm $5^2 = 25$ in $\mathbb{Z}[i]$, since

$$\mathbb{Z}[i]/(5) \cong \{a + bi \mid a, b \in \mathbb{Z}_5\}$$

has $5^2 = 25$ elements.

Now, let's look at

$$p \cdot \mathcal{O}_K = \mathfrak{p}_1^{e_1} \dots \mathfrak{p}_g^{e_g}$$

in $\mathcal{O}_K$, where K has degree n . Taking the ideal norms of both sides, we have that

$$p^n = N(\mathfrak{p}_1)^{e_1} \dots N(\mathfrak{p}_g)^{e_g}.$$

We conclude that $\mathfrak{p}_i = p^{f_i}$ for some integer $f_i \geq 1$, and moreover that

$$n = \sum_{i=1}^g e_i f_i.$$

Definition 43.3.1. We say f_i is the **inertial degree** of $\mathfrak{p}_i$, and e_i is the **ramification index**.

Example 43.3.2 (Examples of inertial degrees)

Work in $\mathbb{Z}[i]$, which is degree 2. The inertial degree detects how “spacy” the given $\mathfrak{p}$ is when interpreted in $\mathcal{O}_K$.

- (a) The prime $7 \cdot \mathbb{Z}[i]$ has inertial degree 2. Indeed, $\mathbb{Z}[i]/(7)$ has $7^2 = 49$ elements, those of the form $a + bi$ for a, b modulo 7. It gives “two degrees” of space.
- (b) Let $(5) = (2 + i)(2 - i)$. The inertial degrees of $(2 + i)$ and $(2 - i)$ are both 1. Indeed, $\mathbb{Z}[i]/(2 + i)$ only gives “one degree” of space, since each of its elements can be viewed as integers modulo 5, and there are only $5^1 = 5$ elements.


If you understand this, it should be intuitively clear why the sum of $e_i f_i$ should equal n .

§43.4 The magic of Galois extensions

OK, that’s all fine and well. But something *really magical* happens when we add the additional hypothesis that $K/\mathbb{Q}$ is *Galois*: all the inertial degrees and ramification degrees are equal. We set about proving this.

Let $K/\mathbb{Q}$ be Galois with $G = \text{Gal}(K/\mathbb{Q})$. Note that if $\mathfrak{p} \subseteq \mathcal{O}_K$ is a prime above p , then the image $\sigma(\mathfrak{p})$ is also prime for any $\sigma \in G$ (since σ is an automorphism!). Moreover, since $p \in \mathfrak{p}$ and σ fixes $\mathbb{Q}$, we know that $p \in \sigma(\mathfrak{p})$ as well.

Thus, by the pointwise mapping, **the Galois group acts on the prime ideals above a rational prime p** . Picture:



The notation $\sigma(\mathfrak{p})$ is hideous in this context, since we’re really thinking of σ as just doing a group action, and so we give the shorthand:

Abuse of Notation 43.4.1. Let $\sigma\mathfrak{p}$ be shorthand for $\sigma^{-1}(\mathfrak{p})$.

Since the σ 's are all bijections (they are automorphisms!), it should come as no surprise that the prime ideals which are in the same orbit are closely related. But miraculously, it turns out there is only one orbit!

Theorem 43.4.2 (Galois group acts transitively)

Let $K/\mathbb{Q}$ be Galois with $G = \text{Gal}(K/\mathbb{Q})$. Let $\{\mathfrak{p}_i\}$ be the set of distinct prime ideals in the factorization of $p \cdot \mathcal{O}_K$ (in $\mathcal{O}_K$).

Then G acts transitively on the $\mathfrak{p}_i$: for every i and j , we can find σ such that $\sigma\mathfrak{p}_i = \mathfrak{p}_j$.

Proof. Fairly slick. Suppose for contradiction that no $\sigma \in G$ sends $\mathfrak{p}_1$ to $\mathfrak{p}_2$, say. By the Chinese remainder theorem, we can find an $x \in \mathcal{O}_K$ such that

$$\begin{aligned} x &\equiv 0 \pmod{\mathfrak{p}_1} \\ x &\equiv 1 \pmod{\mathfrak{p}_i} \text{ for } i \geq 2 \end{aligned}$$

Then, compute the norm

$$N_{K/\mathbb{Q}}(x) = \prod_{\sigma \in \text{Gal}(K/\mathbb{Q})} \sigma(x).$$

Each $\sigma(x)$ is in K because $K/\mathbb{Q}$ is Galois!

Since $N_{K/\mathbb{Q}}(x)$ is an integer and divisible by $\mathfrak{p}_1$, we should have that $N_{K/\mathbb{Q}}(x)$ is divisible by p . Thus it should be divisible by $\mathfrak{p}_2$ as well. But by the way we selected x , we have $x \notin \sigma^{-1}\mathfrak{p}_2$ for every $\sigma \in G$! So $\sigma(x) \notin \mathfrak{p}_2$ for any σ , which is a contradiction. $\square$

Theorem 43.4.3 (Inertial degree and ramification indices are all equal)

Assume $K/\mathbb{Q}$ is Galois. Then for any rational prime p we have

$$p \cdot \mathcal{O}_K = (\mathfrak{p}_1 \mathfrak{p}_2 \cdots \mathfrak{p}_g)^e$$

for some e , where the $\mathfrak{p}_i$ are distinct prime ideals with the same inertial degree f . Hence

$$[K : \mathbb{Q}] = efg.$$

Proof. To see that the inertial degrees are equal, note that each σ induces an isomorphism

$$\mathcal{O}_K/\mathfrak{p} \cong \mathcal{O}_K/\sigma(\mathfrak{p}).$$

Because the action is transitive, all f_i are equal.

Exercise 43.4.4. Using the fact that $\sigma \in \text{Gal}(K/\mathbb{Q})$, show that

$$\sigma^{-1}(p \cdot \mathcal{O}_K) = p \cdot \sigma^{-1}(\mathcal{O}_K) = p \cdot \mathcal{O}_K.$$

So for every σ , we have that $p \cdot \mathcal{O}_K = \prod \mathfrak{p}_i^{e_i} = \prod (\sigma\mathfrak{p}_i)^{e_i}$. Since the action is transitive, all e_i are equal. $\square$

Let's see an illustration of this.

Example 43.4.5 (Factoring 5 in a Galois/non-Galois extension)

Let $p = 5$ be a prime.

- (a) Let $E = \mathbb{Q}(\sqrt[3]{2})$. One can show that $\mathcal{O}_E = \mathbb{Z}[\sqrt[3]{2}]$, so we use the Factoring Algorithm on the minimal polynomial $x^3 - 2$. Since $x^3 - 2 \equiv (x - 3)(x^2 + 3x + 9) \pmod{5}$ is the irreducible factorization, we have that

$$(5) = (5, \sqrt[3]{2} - 3)(5, \sqrt[3]{4} + 3\sqrt[3]{2} + 9)$$

which have inertial degrees 1 and 2, respectively. The fact that this is not uniform reflects that E is not Galois.

- (b) Now let $K = \mathbb{Q}(\sqrt[3]{2}, \omega)$, which is the splitting field of $x^3 - 2$ over $\mathbb{Q}$; now K is Galois. It turns out that

$$\mathcal{O}_K = \mathbb{Z}[\varepsilon] \quad \text{where } \varepsilon \text{ is a root of } t^6 + 3t^5 - 5t^3 + 3t + 1.$$

(this takes a lot of work to obtain, so we won't do it here). Modulo 5 this has an irreducible factorization $(x^2 + x + 2)(x^2 + 3x + 3)(x^2 + 4x + 1) \pmod{5}$, so by the Factorization Algorithm,

$$(5) = (5, \varepsilon^2 + \varepsilon + 2)(5, \varepsilon^2 + 3\varepsilon + 3)(5, \varepsilon^2 + 4\varepsilon + 1).$$

This time all inertial degrees are 2, as the theorem predicts for K Galois.

§43.5 (Optional) Decomposition and inertia groups

Let p be a rational prime. Thus

$$p \cdot \mathcal{O}_K = (\mathfrak{p}_1 \dots \mathfrak{p}_g)^e$$

and all the $\mathfrak{p}_i$ have inertial degree f . Let $\mathfrak{p}$ denote a choice of the $\mathfrak{p}_1$.

We can look at both the fields $\mathcal{O}_K/\mathfrak{p}$ and $\mathbb{Z}/p = \mathbb{F}_p$. Naturally, since $\mathcal{O}_K/\mathfrak{p}$ is a finite field we can view it as a field extension of $\mathbb{F}_p$. So we can get the diagram

$$\begin{array}{ccccccc} K & \supset & \mathcal{O}_K & \mathfrak{p} & \mathcal{O}_K/\mathfrak{p} \cong \mathbb{F}_{p^f} \\ | & & | & | & | \\ \mathbb{Q} & \supset & \mathbb{Z} & (p) & \mathbb{F}_p \end{array}$$

At the far right we have finite field extensions, which we know are *really* well behaved. So we ask:

How are $\text{Gal}((\mathcal{O}_K/\mathfrak{p})/\mathbb{F}_p)$ and $\text{Gal}(K/\mathbb{Q})$ related?

Absurdly enough, there is an explicit answer: **it's just the stabilizer of $\mathfrak{p}$, at least when p is unramified.**

Definition 43.5.1. Let $D_{\mathfrak{p}} \subseteq \text{Gal}(K/\mathbb{Q})$ be the stabilizer of $\mathfrak{p}$, that is

$$D_{\mathfrak{p}} \stackrel{\text{def}}{=} \{\sigma \in \text{Gal}(K/\mathbb{Q}) \mid \sigma\mathfrak{p} = \mathfrak{p}\}.$$

We say $D_{\mathfrak{p}}$ is the **decomposition group** of $\mathfrak{p}$.

Then, every $\sigma \in D_{\mathfrak{p}}$ induces an automorphism of $\mathcal{O}_K/\mathfrak{p}$ by

$$\alpha \mapsto \sigma(\alpha) \pmod{\mathfrak{p}}.$$

So there's a natural map

$$D_{\mathfrak{p}} \xrightarrow{\theta} \text{Gal}((\mathcal{O}_K/\mathfrak{p})/\mathbb{F}_p)$$

by declaring $\theta(\sigma)$ to just be “ $\sigma \pmod{\mathfrak{p}}$ ”. The fact that $\sigma \in D_{\mathfrak{p}}$ (i.e. σ fixes $\mathfrak{p}$) ensures this map is well-defined.

Theorem 43.5.2 (Decomposition group and Galois group)

Define θ as above. Then

- θ is surjective, and
- its kernel is a group of order e , the ramification index.

In particular, if p is unramified then $D_{\mathfrak{p}} \cong \text{Gal}((\mathcal{O}_K/\mathfrak{p})/\mathbb{F}_p)$.

(The proof is not hard, but a bit lengthy and in my opinion not very enlightening.)

If p is unramified, then taking modulo $\mathfrak{p}$ gives $D_{\mathfrak{p}} \cong \text{Gal}((\mathcal{O}_K/\mathfrak{p})/\mathbb{F}_p)$.

But we know exactly what $\text{Gal}((\mathcal{O}_K/\mathfrak{p})/\mathbb{F}_p)$ is! We already have $\mathcal{O}_K/\mathfrak{p} \cong \mathbb{F}_{p^f}$, and the Galois group is

$$\text{Gal}((\mathcal{O}_K/\mathfrak{p})/\mathbb{F}_p) \cong \text{Gal}(\mathbb{F}_{p^f}/\mathbb{F}_p) \cong \langle x \mapsto x^p \rangle \cong \mathbb{Z}_f.$$

So $D_{\mathfrak{p}} \cong \mathbb{Z}_f$ as well.

Let's now go back to

$$D_{\mathfrak{p}} \xrightarrow{\theta} \text{Gal}((\mathcal{O}_K/\mathfrak{p})/\mathbb{F}_p).$$

The kernel of θ is called the **inertia group** and denoted $I_{\mathfrak{p}} \subseteq D_{\mathfrak{p}}$; it has order e .

This gives us a pretty cool sequence of subgroups $\{1\} \subseteq I \subseteq D \subseteq G$ where G is the Galois group (I'm dropping the $\mathfrak{p}$ -subscripts now). Let's look at the corresponding *fixed fields* via the Fundamental theorem of Galois theory. Picture:

$$\begin{array}{ccccc}
 \mathfrak{p} \subseteq \mathcal{O}_K \subseteq & K & \longleftrightarrow & \{1\} \\
 & \downarrow \text{Ramify} & & \downarrow e \\
 & K^I & & I \\
 & \downarrow \text{Inert} & & \downarrow f \\
 & K^D & & D \\
 & \downarrow \text{Split} & & \downarrow g \\
 (p) \subseteq \mathbb{Z} \subseteq & \mathbb{Q} & \longleftrightarrow & G
 \end{array}$$

Something curious happens:

- When (p) is lifted into K^D it splits completely into g unramified primes. Each of these has inertial degree 1.
- When the primes in K^D are lifted to K^I , they remain inert, and now have inertial degree f .
- When then lifted to K , they ramify with exponent e (but don't split at all).

Picture: In other words, the process of going from 1 to efg can be very nicely broken into the three steps above. To draw this in the picture, we get

$$\begin{array}{ccccccc}
 (p) & \longrightarrow & \mathfrak{p}_1 \dots \mathfrak{p}_g & \longrightarrow & \mathfrak{p}_1 \dots \mathfrak{p}_g & \longrightarrow & (\mathfrak{p}_1 \dots \mathfrak{p}_g)^e \\
 \\
 \{f_i\} : & & 1, \dots, 1 & & f, \dots, f & & f, \dots, f \\
 \\
 \mathbb{Q} & \xrightarrow{\text{Split}} & K^D & \xrightarrow{\text{Inert}} & K^I & \xrightarrow{\text{Ramify}} & K
 \end{array}$$

In any case, in the “typical” case that there is no ramification, we just have $K^I = K$.

§43.6 Tangential remark: more general Galois extensions

All the discussion about Galois extensions carries over if we replace $K/\mathbb{Q}$ by some different Galois extension K/F . Instead of a rational prime p breaking down in $\mathcal{O}_K$, we would have a prime ideal $\mathfrak{p}$ of F breaking down as

$$\mathfrak{p} \cdot \mathcal{O}_L = (\mathfrak{P}_1 \dots \mathfrak{P}_g)^e$$

in $\mathcal{O}_L$ and then all results hold verbatim. (The $\mathfrak{P}_i$ are primes in L above $\mathfrak{p}$.) Instead of $\mathbb{F}_p$ we would have $\mathcal{O}_F/\mathfrak{p}$.

The reason I choose to work with $F = \mathbb{Q}$ is that capital Gothic P 's ($\mathfrak{P}$) look *really* terrifying.

§43.7 Problems to think about

Problem 43A[†]. Prove that no rational prime p can remain inert in $K = \mathbb{Q}(\sqrt[3]{2}, \omega)$, the splitting field of $x^3 - 2$. How does this generalize?

44 The Frobenius element

Throughout this chapter $K/\mathbb{Q}$ is a Galois extension with Galois group G , p is an *unramified* rational prime in K , and $\mathfrak{p}$ is a prime above it. Picture:

$$\begin{array}{ccccccc} K & \supset & \mathcal{O}_K & \mathfrak{p} & \mathcal{O}_K/\mathfrak{p} \cong \mathbb{F}_{p^f} \\ | & & | & | & | \\ \mathbb{Q} & \supset & \mathbb{Z} & (p) & \mathbb{F}_p \end{array}$$

If p is unramified, then one can show there is a unique $\sigma \in \text{Gal}(L/K)$ such that $\sigma(\alpha) \equiv \alpha^p \pmod{\mathfrak{p}}$ for every prime p .

§44.1 Frobenius elements

Prototypical example for this section: $\text{Frob}_{\mathfrak{p}}$ in $\mathbb{Z}[i]$ depends on $p \pmod{4}$.

Here is the theorem statement again:

Theorem 44.1.1 (The Frobenius element)

Assume $K/\mathbb{Q}$ is Galois with Galois group G . Let p be a rational prime unramified in K , and $\mathfrak{p}$ a prime above it. There is a *unique* element $\text{Frob}_{\mathfrak{p}} \in G$ with the property that

$$\text{Frob}_{\mathfrak{p}}(\alpha) \equiv \alpha^p \pmod{\mathfrak{p}}.$$

It is called the **Frobenius element** at $\mathfrak{p}$, and has order f .

The *uniqueness* part is pretty important: it allows us to show that a given $\sigma \in \text{Gal}(L/K)$ is the Frobenius element by just observing that it satisfies the above functional equation.

Let's see an example of this:

Example 44.1.2 (Frobenius elements of the Gaussian integers)

Let's actually compute some Frobenius elements for $K = \mathbb{Q}(i)$, which has $\mathcal{O}_K = \mathbb{Z}[i]$. This is a Galois extension, with $G = \mathbb{Z}_2^\times$, corresponding to the identity and complex conjugation.

If p is an odd prime with $\mathfrak{p}$ above it, then $\text{Frob}_{\mathfrak{p}}$ is the unique element such that

$$(a + bi)^p \equiv \text{Frob}_{\mathfrak{p}}(a + bi) \pmod{\mathfrak{p}}$$

in $\mathbb{Z}[i]$. In particular,

$$\text{Frob}_{\mathfrak{p}}(i) = i^p = \begin{cases} i & p \equiv 1 \pmod{4} \\ -i & p \equiv 3 \pmod{4}. \end{cases}$$

From this we see that $\text{Frob}_{\mathfrak{p}}$ is the identity when $p \equiv 1 \pmod{4}$ and $\text{Frob}_{\mathfrak{p}}$ is complex conjugation when $p \equiv 3 \pmod{4}$.

Note that we really only needed to compute $\text{Frob}_{\mathfrak{p}}$ on i . If this seems too good to be true, a philosophical reason is “freshman’s dream” where $(x + y)^p \equiv x^p + y^p \pmod{p}$ (and hence $\pmod{\mathfrak{p}}$). So if σ satisfies the functional equation on generators, it satisfies the functional equation everywhere.

We also have an important lemma:


Lemma 44.1.3 (Order of the Frobenius element)

Let $\text{Frob}_{\mathfrak{p}}$ be a Frobenius element from an extension $K/\mathbb{Q}$. Then the order of $\mathfrak{p}$ is equal to the inertial degree $f_{\mathfrak{p}}$. In particular, (p) splits completely in $\mathcal{O}_K$ if and only if $\text{Frob}_{\mathfrak{p}} = \text{id}$.

Exercise 44.1.4. Prove this lemma as by using the fact that $\mathcal{O}_K/\mathfrak{p}$ is the finite field of order $f_{\mathfrak{p}}$, and the Frobenius element is just $x \mapsto x^p$ on this field.

Let us now prove the main theorem. This will only make sense in the context of decomposition groups, so readers which skipped that part should omit this proof.

Proof of existence of Frobenius element. The entire theorem is just a rephrasing of the fact that the map θ defined in the last section is an isomorphism when p is unramified. Picture:



In here we can restrict our attention to $D_{\mathfrak{p}}$ since we need to have $\sigma(\alpha) \equiv \alpha \pmod{\mathfrak{p}}$ when $\alpha \equiv 0 \pmod{\mathfrak{p}}$. Thus we have the isomorphism

$$D_{\mathfrak{p}} \xrightarrow{\theta} \text{Gal}((\mathcal{O}_K/\mathfrak{p})/\mathbb{F}_p).$$

But we already know $\text{Gal}((\mathcal{O}_K/\mathfrak{p})/\mathbb{F}_p)$, according to the string of isomorphisms

$$\text{Gal}((\mathcal{O}_K/\mathfrak{p})/\mathbb{F}_p) \cong \text{Gal}(\mathbb{F}_{p^f}/\mathbb{F}_p) \cong \langle T = x \mapsto x^p \rangle \cong \mathbb{Z}_f.$$

So the unique such element is the pre-image of T under θ . □

§44.2 Conjugacy classes

Now suppose $\mathfrak{p}_1$ and $\mathfrak{p}_2$ are *two* primes above an unramified rational prime p . Then we can define $\text{Frob}_{\mathfrak{p}_1}$ and $\text{Frob}_{\mathfrak{p}_2}$. Since the Galois group acts transitively, we can select $\sigma \in \text{Gal}(K/\mathbb{Q})$ be such that

$$\sigma(\mathfrak{p}_1) = \mathfrak{p}_2.$$

We claim that

$$\text{Frob}_{\mathfrak{p}_2} = \sigma \circ \text{Frob}_{\mathfrak{p}_1} \circ \sigma^{-1}.$$

Note that this is an equation in G .

Question 44.2.1. Prove this.

More generally, for a given unramified rational prime p , we obtain:

Theorem 44.2.2 (Conjugacy classes in Galois groups)

The set

$$\{\text{Frob}_{\mathfrak{p}} \mid \mathfrak{p} \text{ above } p\}$$

is one of the conjugacy classes of G .

Proof. We've used the fact that $G = \text{Gal}(K/\mathbb{Q})$ is transitive to show that $\text{Frob}_{\mathfrak{p}_1}$ and $\text{Frob}_{\mathfrak{p}_2}$ are conjugate if they both lie above p ; hence it's *contained* in some conjugacy class. So it remains to check that for any $\mathfrak{p}$, σ , we have $\sigma \circ \text{Frob}_{\mathfrak{p}} \circ \sigma^{-1} = \text{Frob}_{\mathfrak{p}'}$ for some $\mathfrak{p}'$. For this, just take $\mathfrak{p}' = \sigma\mathfrak{p}$. Hence the set is indeed a conjugacy class. $\square$

In summary,

$\text{Frob}_{\mathfrak{p}}$ is determined up to conjugation by the prime p from which $\mathfrak{p}$ arises.

So even though the Gothic letters look scary, the content of $\text{Frob}_{\mathfrak{p}}$ really just comes from the more friendly-looking rational prime p .

Example 44.2.3 (Frobenius elements in $\mathbb{Q}(\sqrt[3]{2}, \omega)$)

With those remarks, here is a more involved example of a Frobenius map. Let $K = \mathbb{Q}(\sqrt[3]{2}, \omega)$ be the splitting field of

$$t^3 - 2 = (t - \sqrt[3]{2})(t - \omega\sqrt[3]{2})(t - \omega^2\sqrt[3]{2}).$$

Thus $K/\mathbb{Q}$ is Galois. We've seen in an earlier example that

$$\mathcal{O}_K = \mathbb{Z}[\varepsilon] \quad \text{where } \varepsilon \text{ is a root of } t^6 + 3t^5 - 5t^3 + 3t + 1.$$

Let's consider the prime 5 which factors (trust me here) as

$$(5) = (5, \varepsilon^2 + \varepsilon + 2)(5, \varepsilon^2 + 3\varepsilon + 3)(5, \varepsilon^2 + 4\varepsilon + 1) = \mathfrak{p}_1\mathfrak{p}_2\mathfrak{p}_3.$$

Note that all the prime ideals have inertial degree 2. Thus $\text{Frob}_{\mathfrak{p}_i}$ will have order 2 for each i .

Note that

$$\text{Gal}(K/\mathbb{Q}) = \text{permutations of } \{\sqrt[3]{2}, \omega\sqrt[3]{2}, \omega^2\sqrt[3]{2}\} \cong S_3.$$

In this S_3 there are 3 elements of order three: fixing one root and swapping the other two. These correspond to each of $\text{Frob}_{\mathfrak{p}_1}$, $\text{Frob}_{\mathfrak{p}_2}$, $\text{Frob}_{\mathfrak{p}_3}$.

In conclusion, the conjugacy class $\{\text{Frob}_{\mathfrak{p}_1}, \text{Frob}_{\mathfrak{p}_2}, \text{Frob}_{\mathfrak{p}_3}\}$ associated to (5) is the cycle type $(\bullet)(\bullet\bullet)$ in S_3 .

§44.3 Chebotarev density theorem

Natural question: can we represent every conjugacy class in this way? In other words, is every element of G equal to $\text{Frob}_{\mathfrak{p}}$ for some $\mathfrak{p}$?


Miraculously, not only is the answer “yes”, but in fact it does so in the nicest way possible: the $\text{Frob}_{\mathfrak{p}}$ ’s are “equally distributed” when we pick a random $\mathfrak{p}$.

Theorem 44.3.1 (Chebotarev density theorem over $\mathbb{Q}$)

Let C be a conjugacy class of $G = \text{Gal}(K/\mathbb{Q})$. The density of (unramified) primes p such that $\{\text{Frob}_{\mathfrak{p}} \mid \mathfrak{p} \text{ above } p\} = C$ is exactly $|C| / |G|$. In particular, for any $\sigma \in G$ there are infinitely many rational primes p with $\mathfrak{p}$ above p so that $\text{Frob}_{\mathfrak{p}} = \sigma$.

By density, I mean that the proportion of primes $p \leq x$ that work approaches $\frac{|C|}{|G|}$ as $x \rightarrow \infty$. Note that I’m throwing out the primes that ramify in K . This is no issue, since the only primes that ramify are those dividing Δ_K , of which there are only finitely many.

In other words, if I pick a random prime p and look at the resulting conjugacy class, it’s a lot like throwing a dart at G : the probability of hitting any conjugacy class depends just on the size of the class.



Remark 44.3.2. Happily, this theorem (and preceding discussion) also works if we replace $K/\mathbb{Q}$ with any Galois extension K/F ; in that case we replace “ $\mathfrak{p}$ over p ” with “ $\mathfrak{P}$ over $\mathfrak{p}$ ”. In that case, we use $N(\mathfrak{p}) \leq x$ rather than $p \leq x$ as the way to define density.

§44.4 Example: Frobenius elements of cyclotomic fields

Let q be a prime, and consider $L = \mathbb{Q}(\zeta_q)$, with q a primitive q th root of unity. You should recall from various starred problems that

- $\Delta_L = \pm q^{q-2}$,
- $\mathcal{O}_L = \mathbb{Z}[\zeta_q]$, and
- The map

$$\sigma_n : L \rightarrow L \quad \text{by} \quad \zeta_q \mapsto \zeta_q^n$$

is an automorphism of L whenever $\gcd(n, q) = 1$, and depends only on $n \pmod{q}$. In other words, the automorphisms of $L/\mathbb{Q}$ just shuffle around the q th roots of unity. In fact the Galois group consists exactly of the elements $\{\sigma_n\}$, namely

$$\text{Gal}(L/\mathbb{Q}) = \{\sigma_n \mid n \not\equiv 0 \pmod{q}\}.$$

As a group,

$$\mathrm{Gal}(L/\mathbb{Q}) = \mathbb{Z}_q^\times \cong \mathbb{Z}_{q-1}.$$

This is surprisingly nice, because **elements of $\mathrm{Gal}(L/\mathbb{Q})$ look a lot like Frobenius elements already**. Specifically:

Lemma 44.4.1 (Cyclotomic Frobenius elements)

In the cyclotomic setting $L = \mathbb{Q}(\zeta_q)$, let p be a rational unramified prime and $\mathfrak{p}$ above it. Then

$$\mathrm{Frob}_{\mathfrak{p}} = \sigma_p.$$

Proof. Observe that σ_p satisfies the functional equation (check on generators). Done by uniqueness. $\square$

Question 44.4.2. Conclude that a rational prime p splits completely in $\mathcal{O}_L$ if and only if $p \equiv 1 \pmod{m}$.

§44.5 Frobenius elements behave well with restriction

Let $L/\mathbb{Q}$ and $K/\mathbb{Q}$ be Galois extensions, and consider the setup

$$\begin{array}{ccccc} L & \supseteq & \mathfrak{P} & \cdots \longrightarrow & \mathrm{Frob}_{\mathfrak{P}} \in \mathrm{Gal}(L/\mathbb{Q}) \\ | & & | & & \\ K & \supseteq & \mathfrak{p} & \cdots \longrightarrow & \mathrm{Frob}_{\mathfrak{p}} \in \mathrm{Gal}(K/\mathbb{Q}) \\ | & & | & & \\ \mathbb{Q} & \supseteq & (p) & & \end{array}$$

Here $\mathfrak{p}$ is above (p) and $\mathfrak{P}$ is above $\mathfrak{p}$. We may define

$$\mathrm{Frob}_{\mathfrak{p}} : K \rightarrow K \quad \text{and} \quad \mathrm{Frob}_{\mathfrak{P}} : L \rightarrow L$$

and want to know how these are related.

Theorem 44.5.1 (Restrictions of Frobenius elements)

Assume $L/\mathbb{Q}$ and $K/\mathbb{Q}$ are both Galois. Let $\mathfrak{P}$ and $\mathfrak{p}$ be unramified as above. Then $\mathrm{Frob}_{\mathfrak{P}}|_K = \mathrm{Frob}_{\mathfrak{p}}$, i.e. for every $\alpha \in K$,

$$\mathrm{Frob}_{\mathfrak{p}}(\alpha) = \mathrm{Frob}_{\mathfrak{P}}(\alpha).$$

Proof. We know

$$\mathrm{Frob}_{\mathfrak{P}}(\alpha) \equiv \alpha^p \pmod{\mathfrak{P}} \quad \forall \alpha \in \mathcal{O}_L$$

from the definition.

Question 44.5.2. Deduce that

$$\text{Frob}_{\mathfrak{p}}(\alpha) \equiv \alpha^p \pmod{\mathfrak{p}} \quad \forall \alpha \in \mathcal{O}_K.$$

(This is weaker than the previous statement in two ways!)

Thus $\text{Frob}_{\mathfrak{p}}$ restricted to $\mathcal{O}_K$ satisfies the characterizing property of $\text{Frob}_{\mathfrak{p}}$. $\square$

In short, the point of this section is that

Frobenius elements upstairs restrict to Frobenius elements downstairs.

§44.6 Application: Quadratic reciprocity

We now aim to prove:

Theorem 44.6.1 (Quadratic reciprocity)

Let p and q be distinct odd primes. Then

$$\left(\frac{p}{q}\right) \left(\frac{q}{p}\right) = (-1)^{\frac{p-1}{2} \cdot \frac{q-1}{2}}.$$

(See, e.g. [Le] for an exposition on quadratic reciprocity, if you're not familiar with it.)

Step 1: Setup

For this proof, we first define

$$L = \mathbb{Q}(\zeta_q)$$

where ζ_q is a primitive q th root of unity. Then $L/\mathbb{Q}$ is Galois, with Galois group G .

Question 44.6.2. Show that G has a unique subgroup H of order two.

In fact, we can describe it exactly: viewing $G \cong \mathbb{Z}_q^\times$, we have

$$H = \{\sigma_n \mid n \text{ quadratic residue mod } q\}.$$

By the fundamental theorem of Galois Theory, there ought to be a degree 2 extension of $\mathbb{Q}$ inside $\mathbb{Q}(\zeta_q)$ (that is, a quadratic field). Call it $\mathbb{Q}(\sqrt{q^*})$, for q^* squarefree:

$$\begin{array}{ccc} L = \mathbb{Q}(\zeta_q) & \longleftrightarrow & \{1\} \\ \downarrow \scriptstyle \frac{q-1}{2} & & \downarrow \\ K = \mathbb{Q}(\sqrt{q^*}) & \longleftrightarrow & H \\ \downarrow \scriptstyle 2 & & \downarrow \\ \mathbb{Q} & \longleftrightarrow & G \end{array}$$

Exercise 44.6.3. Note that if a rational prime ℓ ramifies in K , then it ramifies in L . Use this to show that

$$q^* = \pm q \text{ and } q^* \equiv 1 \pmod{4}.$$

Together these determine the value of q^* .

(Actually, it is true in general Δ_K divides Δ_L in a tower $L/K/\mathbb{Q}$.)

Step 2: Reformulation

Now we are going to prove:

Theorem 44.6.4 (Quadratic reciprocity, equivalent formulation)

For distinct odd primes p, q we have

$$\left(\frac{p}{q}\right) = \left(\frac{q^*}{p}\right).$$

Exercise 44.6.5. Using the fact that $\left(\frac{-1}{p}\right) = (-1)^{\frac{p-1}{2}}$, show that this is equivalent to quadratic reciprocity as we know it.

We look at the rational prime p in $\mathbb{Z}$. Either it splits into two in K or is inert; either way let $\mathfrak{p}$ be a prime factor in the resulting decomposition (so $\mathfrak{p}$ is either $p \cdot \mathcal{O}_K$ in the inert case, or one of the primes in the split case). Then let $\mathfrak{P}$ be above $\mathfrak{p}$. It could possibly also split in K : the picture looks like

$$\begin{array}{ccc} \mathcal{O}_L = \mathbb{Z}[\zeta_q] & \supset \mathfrak{P} & \longrightarrow \mathbb{Z}[\zeta_p]/\mathfrak{P} \cong \mathbb{F}_{p^f} \\ \mathcal{O}_K = \mathbb{Z}\left[\frac{1+\sqrt{q^*}}{2}\right] & \supseteq \mathfrak{p} & \longrightarrow \mathbb{F}_p \text{ or } \mathbb{F}_{p^2} \\ \mathbb{Z} & \supseteq (p) & \longrightarrow \mathbb{F}_p \end{array}$$

Question 44.6.6. Why is p not ramified in either K or L ?

Step 3: Introducing the Frobenius

Now, we take the Frobenius

$$\sigma_p = \text{Frob}_{\mathfrak{P}} \in \text{Gal}(L/\mathbb{Q}).$$

We claim that

$$\text{Frob}_{\mathfrak{P}} \in H \iff p \text{ splits in } K.$$

To see this, note that $\text{Frob}_{\mathfrak{P}}$ is in H if and only if it acts as the identity on K . But $\text{Frob}_{\mathfrak{P}}|_K$ is $\text{Frob}_{\mathfrak{p}}$! So

$$\text{Frob}_{\mathfrak{P}} \in H \iff \text{Frob}_{\mathfrak{p}} = \text{id}_K.$$

Finally note that $\text{Frob}_{\mathfrak{p}}$ has order 1 if p splits ($\mathfrak{p}$ has inertial degree 1) and order 2 if p is inert. This completes the proof of the claim.

Finishing up

We already know by Lemma 44.4.1 that $\text{Frob}_{\mathfrak{p}} = \sigma_p \in H$ if and only if p is a quadratic residue. On the other hand,

Exercise 44.6.7. Show that p splits in $\mathcal{O}_K = \mathbb{Z}[\frac{1}{2}(1 + \sqrt{q^*})]$ if and only if $\left(\frac{q^*}{p}\right) = 1$. (Use the factoring algorithm. You need the fact that $p \neq 2$ here.)

In other words

$$\left(\frac{p}{q}\right) = 1 \iff \sigma_p \in H \iff p \text{ splits in } \mathbb{Z}\left[\frac{1}{2}(1 + \sqrt{q^*})\right] \iff \left(\frac{q^*}{p}\right) = 1.$$

This completes the proof.

§44.7 Frobenius elements control factorization

Prototypical example for this section: $\text{Frob}_{\mathfrak{p}}$ controlled the splitting of p in the proof of quadratic reciprocity; the same holds in general.

In the proof of quadratic reciprocity, we used the fact that Frobenius elements behaved well with restriction in order to relate the splitting of p with properties of $\text{Frob}_{\mathfrak{p}}$.

In fact, there is a much stronger statement for any intermediate field $\mathbb{Q} \subseteq E \subseteq K$ which works even if $E/\mathbb{Q}$ is not Galois. It relies on the notion of a *factorization pattern*. Here is how it goes.

Set $n = [E : \mathbb{Q}]$, and let p be a rational prime unramified in K . Then p can be broken in E as


$$p \cdot \mathcal{O}_E = \mathfrak{p}_1 \mathfrak{p}_2 \cdots \mathfrak{p}_g$$

with inertial degrees $f_1, \dots, f_g$: (these inertial degrees might be different since $E/\mathbb{Q}$ isn't Galois). The numbers $f_1 + \cdots + f_g = n$ form a partition of the number n . For example, in the quadratic reciprocity proof we had $n = 2$, with possible partitions $1 + 1$ (if p split) and 2 (if p was inert). We call this the **factorization pattern** of p in E .

Next, we introduce a Frobenius $\text{Frob}_{\mathfrak{p}}$ above (p) , all the way in K ; this is an element of $G = \text{Gal}(K/\mathbb{Q})$. Then let H be the group corresponding to the field E . Diagram:

$$\begin{array}{ccc}
 K & \longleftrightarrow & \{1\} \\
 \downarrow & & \downarrow \\
 E & \longleftrightarrow & H \\
 \downarrow & & \downarrow \\
 \mathbb{Q} & \longleftrightarrow & G
 \end{array}
 \quad
 \begin{array}{l}
 \text{Frob}_{\mathfrak{p}} \\
 \\
 \mathfrak{p}_1 \cdots \mathfrak{p}_g \quad f_1 + \cdots + f_g = n \\
 \downarrow \\
 (p)
 \end{array}$$

Then $\text{Frob}_{\mathfrak{p}}$ induces a *permutation* of the n left cosets gH by left multiplication (after all, $\text{Frob}_{\mathfrak{p}}$ is an element of G too!). Just as with any permutation, we may look at the resulting cycle decomposition, which has a natural “cycle structure”: a partition of n .



The theorem is that these coincide:

Theorem 44.7.1 (Frobenius elements control decomposition)

Let $\mathbb{Q} \subseteq E \subseteq K$ an extension of number fields and assume $K/\mathbb{Q}$ is Galois (though $E/\mathbb{Q}$ need not be). Pick an unramified rational prime p ; let $G = \text{Gal}(K/\mathbb{Q})$ and H the corresponding intermediate subgroup. Finally, let $\mathfrak{P}$ be a prime above p in K .

Then the *factorization pattern* of p in E is given by the *cycle structure* of $\text{Frob}_{\mathfrak{P}}$ acting on the left cosets of H .

Often, we take $E = K$, in which case this is just asserting that the decomposition of the prime p is controlled by a Frobenius element over it.

An important special case is when $E = \mathbb{Q}(\alpha)$, because as we will see it is let us determine how the minimal polynomial of α factors modulo p . To motivate this, let's go back a few chapters and think about the Factoring Algorithm.

Let α be an algebraic integer and f its minimal polynomial (of degree n). Set $E = \mathbb{Q}(\alpha)$ (which has degree n over $\mathbb{Q}$). Suppose we're lucky enough that $\mathcal{O}_E = \mathbb{Z}[\alpha]$, i.e. that E is monogenic. Then we know by the Factoring Algorithm, to factor any p in E , all we have to do is factor f modulo p , since if $f = f_1^{e_1} \dots f_g^{e_g} \pmod{p}$ then we have

$$(p) = \prod_i \mathfrak{p}_i = \prod_i (f_i(\alpha), p)^{e_i}.$$

This gives us complete information about the ramification indices and inertial degrees; the e_i are the ramification indices, and $\deg f_i$ are the inertial degrees (since $\mathcal{O}_E/\mathfrak{p}_i \cong \mathbb{F}_p[X]/(f_i(X))$).

In particular, if p is unramified then all the e_i are equal to 1, and we get

$$n = \deg f = \deg f_1 + \deg f_2 + \dots + \deg f_g.$$

Once again we have a partition of n ; we call this the **factorization pattern** of f modulo p . So, to see the factorization pattern of an unramified p in $\mathcal{O}_E$, we just have to know the factorization pattern of the $f \pmod{p}$.

Turning this on its head, if we want to know the factorization pattern of $f \pmod{p}$, we just need to know how p decomposes. And it turns out these coincide even without the assumption that E is monogenic.

Theorem 44.7.2 (Frobenius controls polynomial factorization)

Let α be an algebraic integer with minimal polynomial f , and let $E = \mathbb{Q}(\alpha)$. Then for any prime p unramified in the splitting field K of f , the following coincide:

- (i) The factorization pattern of p in E .
- (ii) The factorization pattern of $f \pmod{p}$.
- (iii) The cycle structure associated to the action of $\text{Frob}_{\mathfrak{P}} \in \text{Gal}(K/\mathbb{Q})$ on the roots of f , where $\mathfrak{P}$ is above p in K .

Example 44.7.3 (Factoring $x^3 - 2 \pmod{5}$)

Let $\alpha = \sqrt[3]{2}$ and $f = x^3 - 2$, so $E = \mathbb{Q}(\sqrt[3]{2})$. Set $p = 5$ and let finally, let $K = \mathbb{Q}(\sqrt[3]{2}, \omega)$ be the splitting field. Setup:

$$\begin{array}{ccc}
 K = \mathbb{Q}(\sqrt[3]{2}, \omega) & \mathfrak{P} & x^3 - 2 = (x - \sqrt[3]{2})(x - \sqrt[3]{2}\omega)(x - \sqrt[3]{2}\omega^2) \\
 \downarrow 2 & \downarrow & \\
 E = \mathbb{Q}(\sqrt[3]{2}) & \mathfrak{p} & x^3 - 2 = (x - \sqrt[3]{2})(x^2 + \sqrt[3]{2}x + \sqrt[3]{4}) \\
 \downarrow 3 & \downarrow & \\
 \mathbb{Q} & (5) & x^3 - 2 \text{ irreducible over } \mathbb{Q}
 \end{array}$$

The three claimed objects now all have shape $2 + 1$:

- (i) By the Factoring Algorithm, we have $(5) = (5, \sqrt[3]{2} - 3)(5, 9 + 3\sqrt[3]{2} + \sqrt[3]{4})$.
- (ii) We have $x^3 - 2 \equiv (x - 3)(x^2 + 3x + 9) \pmod{5}$.
- (iii) We saw before that $\text{Frob}_{\mathfrak{P}} = (\bullet)(\bullet\bullet)$.

Sketch of Proof. Letting $n = \deg f$. Let H be the subgroup of $G = \text{Gal}(K/\mathbb{Q})$ corresponding to E , so $[G : E] = n$. Pictorially, we have

$$\begin{array}{ccc}
 K & \{1\} & \mathfrak{P} \\
 \downarrow & \downarrow & \downarrow \\
 E = \mathbb{Q}(\alpha) & H & \mathfrak{p} \\
 \downarrow & \downarrow & \downarrow \\
 \mathbb{Q} & G & (p)
 \end{array}$$

We claim that (i), (ii), (iii) are all equivalent to

(iv) The pattern of the action of $\text{Frob}_{\mathfrak{p}}$ on the G/H .

In other words we claim the cosets correspond to the n roots of f in K . Indeed H is just the set of $\tau \in G$ such that $\tau(\alpha) = \alpha$, so there's a bijection between the roots and the cosets G/H by $\tau H \mapsto \tau(\alpha)$. Think of it this way: if $G = S_n$, and $H = \{\tau : \tau(1) = 1\}$, then G/H has order $n!/(n-1)! = n$ and corresponds to the elements $\{1, \dots, n\}$. So there is a natural bijection from (iii) to (iv).

The fact that (i) is in bijection to (iv) was the previous theorem, Theorem 44.7.1. The correspondence (i) $\iff$ (ii) is a fact of Galois theory, so we omit the proof here. $\square$

All this can be done in general with $\mathbb{Q}$ replaced by F ; for example, in [Le02].

§44.8 Example application: IMO 2003 problem 6

As an example of the power we now have at our disposal, let's prove:



Problem 6. Let p be a prime number. Prove that there exists a prime number q such that for every integer n , the number $n^p - p$ is not divisible by q .

We will show, much more strongly, that there exist infinitely many primes q such that $X^p - p$ is irreducible modulo q .

Solution. Okay! First, we draw the tower of fields

$$\mathbb{Q} \subseteq \mathbb{Q}(\sqrt[p]{p}) \subseteq K$$

where K is the splitting field of $f(x) = x^p - p$. Let $E = \mathbb{Q}(\sqrt[p]{p})$ for brevity and note it has degree $[E : \mathbb{Q}] = p$. Let $G = \text{Gal}(K/\mathbb{Q})$.

Question 44.8.1. Show that p divides the order of G . (Look at E .)

Hence by Cauchy's theorem (Problem 11A*, which is a purely group-theoretic fact) we can find a $\sigma \in G$ of order p . By Chebotarev, there exist infinitely many rational (unramified) primes $q \neq p$ and primes $\mathfrak{Q} \subseteq \mathcal{O}_K$ above q such that $\text{Frob}_{\mathfrak{Q}} = \sigma$. (Yes, that's an uppercase Gothic Q . Sorry.)

We claim that all these q work.

By Theorem 44.7.2, the factorization of $f \pmod{q}$ is controlled by the action of $\sigma = \text{Frob}_{\mathfrak{Q}}$ on the roots of f . But σ has prime order p in G ! So all the lengths in the cycle structure have to divide p . Thus the possible factorization patterns of f are

$$p = \underbrace{1 + 1 + \cdots + 1}_{p \text{ times}} \quad \text{or} \quad p = p.$$

So we just need to rule out the $p = 1 + \cdots + 1$ case now: this only happens if f breaks into linear factors mod q . Intuitively this edge case seems highly unlikely (are we really so unlucky that f factors into *linear* factors when we want it to be irreducible?). And indeed this is easy to see: this means that σ fixes all of the roots of f in K , but that means σ fixes K altogether, and hence is the identity of G , contradiction. $\square$

Remark 44.8.2. In fact $K = \mathbb{Q}(\sqrt[p]{p}, \zeta_p)$, and $|G| = p(p-1)$. With a little more group theory, we can show that in fact the density of primes q that work is $\frac{1}{p}$.

§44.9 Problems to think about

Problem 44A. Show that for an odd prime p ,

$$\left(\frac{2}{p}\right) = (-1)^{\frac{1}{8}(p^2-1)}.$$

Problem 44B. Let f be a nonconstant polynomial with integer coefficients. Suppose $f \pmod{p}$ splits completely into linear factors for all sufficiently large primes p . Show that f splits completely into linear factors.

Problem 44C[†] (Dirichlet's theorem on arithmetic progressions). Let a and m be relatively prime positive integers. Show that the density of primes $p \equiv a \pmod{m}$ is exactly $\frac{1}{\phi(m)}$.

Problem 44D. Let n be an odd integer which is not a prime power. Show that the n th cyclotomic polynomial is not irreducible modulo *any* rational prime.



Problem 44E (Putnam 2012 B6). Let p be an odd prime such that $p \equiv 2 \pmod{3}$. Let π be a permutation of $\mathbb{F}_p$ by $\pi(x) = x^3 \pmod{p}$. Show that π is even if and only if $p \equiv 3 \pmod{4}$.

45 Bonus: A Bit on Artin Reciprocity

In this chapter, I'm going to state some big theorems of global class field theory and use them to deduce the Kronecker-Weber plus Hilbert class fields. No proofs, but hopefully still appreciable. For experts: this is global class field theory, without ideles.

Here's the executive summary: let K be a number field. Then all abelian extensions L/K can be understood using solely information intrinsic to K : namely, the ray class groups (generalizing ideal class groups).

§45.1 Infinite primes

Prototypical example for this section: $\mathbb{Q}(\sqrt{-5})$ has a complex infinite prime, $\mathbb{Q}(\sqrt{5})$ has two real infinite ones.

Let K be a number field of degree n and signature (r, s) . We know what a prime ideal of $\mathcal{O}_K$ is, but we now allow for the so-called infinite primes, which I'll describe using the embeddings. Recall there are n embeddings $\sigma : K \rightarrow \mathbb{C}$, which consist of

- r real embeddings where $\text{im } \sigma \subseteq \mathbb{R}$, and
- s pairs of conjugate complex embeddings.

Hence $r + 2s = n$. The first class of embeddings form the **real infinite primes**, while the **complex infinite primes** are the second type. We say K is **totally real** (resp **totally complex**) if all its infinite primes are real (resp complex).

Example 45.1.1 (Examples of infinite primes)

- $\mathbb{Q}$ has a single real infinite prime. We often write it as ∞ .
- $\mathbb{Q}(\sqrt{-5})$ has a single complex infinite prime, and no real infinite primes. Hence totally complex.
- $\mathbb{Q}(\sqrt{5})$ has two real infinite primes, and no complex infinite primes. Hence totally real.

§45.2 Modular arithmetic with infinite primes

A **modulus** is a formal product

$$\mathfrak{m} = \prod_{\mathfrak{p}} \mathfrak{p}^{\nu(\mathfrak{p})}$$

where the product runs over all primes, finite and infinite. (Here $\nu(\mathfrak{p})$ is a nonnegative integer, of which only finitely many are nonzero.) We also require that

- $\nu(\mathfrak{p}) = 0$ for any infinite prime $\mathfrak{p}$, and
- $\nu(\mathfrak{p}) \leq 1$ for any real prime $\mathfrak{p}$.

Obviously, every $\mathfrak{m}$ can be written as $\mathfrak{m} = \mathfrak{m}_0 \mathfrak{m}_\infty$ by separating the finite from the (real) infinite primes.

We say $a \equiv b \pmod{\mathfrak{m}}$ if

- If $\mathfrak{p}$ is a finite prime, then $a \equiv b \pmod{\mathfrak{p}^{\nu(\mathfrak{p})}}$ means exactly what you think it should mean: $a - b \in \mathfrak{p}^{\nu(\mathfrak{p})}$.
- If $\mathfrak{p}$ is a *real* infinite prime $\sigma : K \rightarrow \mathbb{R}$, then $a \equiv b \pmod{\mathfrak{p}}$ means that $\sigma(a/b) > 0$.

Of course, $a \equiv b \pmod{\mathfrak{m}}$ means $a \equiv b$ modulo each prime power in $\mathfrak{m}$. With this, we can define a generalization of the class group:

Definition 45.2.1. Let $\mathfrak{m}$ be a modulus of a number field K .

- Let $I_K(\mathfrak{m})$ to be the set of all fractional ideals of K which are relatively prime to $\mathfrak{m}$, which is an abelian group.
- Let $P_K(\mathfrak{m})$ be the subgroup of $I_K(\mathfrak{m})$ generated by

$$\{\alpha \mathcal{O}_K \mid \alpha \in K^\times \text{ and } \alpha \equiv 1 \pmod{\mathfrak{m}}\}.$$

This is sometimes called a “ray” of principal ideals.

Finally define the **ray class group** of $\mathfrak{m}$ to be $C_K(\mathfrak{m}) = I_K(\mathfrak{m})/P_K(\mathfrak{m})$.

This group is known to always be finite. Note the usual class group is $C_K(1)$.

One last definition that we’ll use right after Artin reciprocity:

Definition 45.2.2. A **congruence subgroup** of $\mathfrak{m}$ is a subgroup H with

$$P_K(\mathfrak{m}) \subseteq H \subseteq I_K(\mathfrak{m}).$$

Thus $C_K(\mathfrak{m})$ is a group which contains a lattice of various quotients $I_K(\mathfrak{m})/H$, where H is a congruence subgroup.

This definition takes a while to get used to, so here are examples.

Example 45.2.3 (Ray class groups in $\mathbb{Q}$, finite modulus)

Consider $K = \mathbb{Q}$ with infinite prime ∞ . Then

- If we take $\mathfrak{m} = 1$ then $I_{\mathbb{Q}}(1)$ is all fractional ideals, and $P_{\mathbb{Q}}(1)$ is all principal fractional ideals. Their quotient is the usual class group of $\mathbb{Q}$, which is trivial.
- Now take $\mathfrak{m} = 8$. Thus $I_{\mathbb{Q}}(8) \cong \{\frac{a}{b}\mathbb{Z} \mid a/b \equiv 1, 3, 5, 7 \pmod{8}\}$. Moreover

$$P_{\mathbb{Q}}(8) \cong \left\{ \frac{a}{b}\mathbb{Z} \mid a/b \equiv 1 \pmod{8} \right\}.$$

You might at first glance think that the quotient is thus $(\mathbb{Z}/8\mathbb{Z})^\times$. But the issue is that we are dealing with *ideals*: specifically, we have

$$7\mathbb{Z} = -7\mathbb{Z} \in P_{\mathbb{Q}}(8)$$

because $-7 \equiv 1 \pmod{8}$. So *actually*, we get

$$C_{\mathbb{Q}}(8) \cong \{1, 3, 5, 7 \bmod 8\} / \{1, 7 \bmod 8\} \cong (\mathbb{Z}/4\mathbb{Z})^\times.$$

More generally,

$$C_{\mathbb{Q}}(m) = (\mathbb{Z}/m\mathbb{Z})^\times / \{\pm 1\}.$$

Example 45.2.4 (Ray class groups in $\mathbb{Q}$, infinite moduli)

Consider $K = \mathbb{Q}$ with infinite prime ∞ again.

- Now take $\mathfrak{m} = \infty$. As before $I_{\mathbb{Q}}(\infty) = \mathbb{Q}^\times$. Now, by definition we have

$$P_{\mathbb{Q}}(\infty) = \left\{ \frac{a}{b} \mathbb{Z} \mid a/b > 0 \right\}.$$

At first glance you might think this was $\mathbb{Q}_{>0}$, but the same behavior with ideals shows in fact $P_{\mathbb{Q}}(\infty) = \mathbb{Q}^\times$. So in this case, $P_{\mathbb{Q}}(\infty)$ still has all principal fractional ideals. Therefore, $C_{\mathbb{Q}}(\infty)$ is still trivial.

- Finally, let $\mathfrak{m} = 8\infty$. As before $I_{\mathbb{Q}}(8\infty) \cong \left\{ \frac{a}{b} \mathbb{Z} \mid a/b \equiv 1, 3, 5, 7 \pmod{8} \right\}$. Now in this case:

$$P_{\mathbb{Q}}(8\infty) \cong \left\{ \frac{a}{b} \mathbb{Z} \mid a/b \equiv 1 \pmod{8} \text{ and } a/b > 0 \right\}.$$

This time, we really do have $-7\mathbb{Z} \notin P_{\mathbb{Q}}(8\infty)$: we have $7 \not\equiv 1 \pmod{8}$ and also $-8 < 0$. So neither of the generators of $7\mathbb{Z}$ are in $P_{\mathbb{Q}}(8\infty)$. Thus we finally obtain

$$C_{\mathbb{Q}}(8\infty) \cong \{1, 3, 5, 7 \pmod{8}\} / \{1 \pmod{8}\} \cong (\mathbb{Z}/8\mathbb{Z})^\times$$

with the bijection $C_{\mathbb{Q}}(8\infty) \rightarrow (\mathbb{Z}/8\mathbb{Z})^\times$ given by $a\mathbb{Z} \mapsto |a| \pmod{8}$.

More generally,

$$C_{\mathbb{Q}}(m\infty) = (\mathbb{Z}/m\mathbb{Z})^\times.$$

§45.3 Infinite primes in extensions

I want to emphasize that everything above is *intrinsic* to a particular number field K . After this point we are going to consider extensions L/K but it is important to keep in mind the distinction that the concept of modulus and ray class group are objects defined solely from K rather than the above L .

Now take a *Galois* extension L/K of degree m . We already know prime ideals $\mathfrak{p}$ of K break into a product of prime ideals $\mathfrak{P}$ of L in a nice way, so we want to do the same thing with infinite primes. This is straightforward: each of the n infinite primes $\sigma : K \rightarrow \mathbb{C}$ lifts to m infinite primes $\tau : L \rightarrow \mathbb{C}$, by which I mean the diagram

$$\begin{array}{ccc} L & \xrightarrow{\tau} & \mathbb{C} \\ \uparrow & \nearrow \sigma & \\ K & & \end{array}$$

commutes. Hence like before, each infinite prime σ of K has m infinite primes τ of L which lie above it.

For a real prime σ of K , any of the resulting τ above it are complex, we say that the prime σ **ramifies** in the extension L/K . Otherwise it is **unramified** in L/K . An infinite prime of K is always unramified in L/K . In this way, we can talk about an unramified Galois extension L/K : it is one where all primes (finite or infinite) are unramified.

Example 45.3.1 (Ramification of ∞)

Let ∞ be the real infinite prime of $\mathbb{Q}$.

- ∞ is ramified in $\mathbb{Q}(\sqrt{-5})/\mathbb{Q}$.
- ∞ is unramified in $\mathbb{Q}(\sqrt{5})/\mathbb{Q}$.

Note also that if K is totally complex then any extension L/K is unramified.

§45.4 Frobenius element and Artin symbol

Recall the key result:

Theorem 45.4.1 (Frobenius element)

Let L/K be a *Galois* extension. If $\mathfrak{p}$ is a prime unramified in K , and $\mathfrak{P}$ a prime above it in L . Then there is a unique element of $\text{Gal}(L/K)$, denoted $\text{Frob}_{\mathfrak{P}}$, obeying

$$\text{Frob}_{\mathfrak{P}}(\alpha) \equiv \alpha^{N_{\mathfrak{p}}} \pmod{\mathfrak{P}} \quad \forall \alpha \in \mathcal{O}_L.$$

Example 45.4.2 (Example of Frobenius elements)

Let $L = \mathbb{Q}(i)$, $K = \mathbb{Q}$. We have $\text{Gal}(L/K) \cong \mathbb{Z}/2\mathbb{Z}$.

If p is an odd prime with $\mathfrak{P}$ above it, then $\text{Frob}_{\mathfrak{P}}$ is the unique element such that

$$(a + bi)^p \equiv \text{Frob}_{\mathfrak{P}}(a + bi) \pmod{\mathfrak{P}}$$

in $\mathbb{Z}[i]$. In particular,

$$\text{Frob}_{\mathfrak{P}}(i) = i^p = \begin{cases} i & p \equiv 1 \pmod{4} \\ -i & p \equiv 3 \pmod{4}. \end{cases}$$

From this we see that $\text{Frob}_{\mathfrak{P}}$ is the identity when $p \equiv 1 \pmod{4}$ and $\text{Frob}_{\mathfrak{P}}$ is complex conjugation when $p \equiv 3 \pmod{4}$.

Example 45.4.3 (Cyclotomic Frobenius element)

Generalizing previous example, let $L = \mathbb{Q}(\zeta)$ and $K = \mathbb{Q}$, with ζ an m th root of unity. It's well-known that L/K is unramified outside ∞ and prime factors of m . Moreover, the Galois group $\text{Gal}(L/K)$ is $(\mathbb{Z}/m\mathbb{Z})^\times$: the Galois group consists of elements of the form

$$\sigma_n : \zeta \mapsto \zeta^n$$

and $\text{Gal}(L/K) = \{\sigma_n \mid n \in (\mathbb{Z}/m\mathbb{Z})^\times\}$.

Then it follows just like before that if $p \nmid n$ is prime and $\mathfrak{P}$ is above p

$$\text{Frob}_{\mathfrak{P}}(x) = \sigma_p.$$

An important property of the Frobenius element is its order is related to the decomposition of $\mathfrak{p}$ in the higher field L in the nicest way possible:

Lemma 45.4.4 (Order of the Frobenius element)

The Frobenius element $\text{Frob}_{\mathfrak{P}} \in \text{Gal}(L/K)$ of an extension L/K has order equal to the inertial degree of $\mathfrak{P}$, that is,

$$\text{ord } \text{Frob}_{\mathfrak{P}} = f(\mathfrak{P} | \mathfrak{p}).$$

In particular, $\text{Frob}_{\mathfrak{P}} = \text{id}$ if and only if $\mathfrak{p}$ splits completely in L/K .

Proof. We want to understand the order of the map $T : x \mapsto x^{N_{\mathfrak{P}}}$ on the field $\mathcal{O}_K/\mathfrak{P}$. But the latter is isomorphic to the splitting field of $X^{N_{\mathfrak{P}}} - X$ in $\mathbb{F}_p$, by Galois theory of finite fields. Hence the order is $\log_{N_{\mathfrak{P}}}(N_{\mathfrak{P}}) = f(\mathfrak{P} | \mathfrak{p})$. $\square$

Exercise 45.4.5. Deduce from this that the rational primes which split completely in $\mathbb{Q}(\zeta)$ are exactly those with $p \equiv 1 \pmod{m}$. Here ζ is an m th root of unity.

The Galois group acts transitively among the set of $\mathfrak{P}$ above a given $\mathfrak{p}$, so that we have

$$\text{Frob}_{\sigma(\mathfrak{P})} = \sigma \circ (\text{Frob}_{\mathfrak{P}}) \circ \sigma^{-1}.$$

Thus $\text{Frob}_{\mathfrak{P}}$ is determined by its underlying $\mathfrak{p}$ up to conjugation.

In class field theory, we are interested in **abelian extensions**, i.e. those for which $\text{Gal}(L/K)$ is abelian. Here the theory becomes extra nice: the conjugacy classes have size one.

Definition 45.4.6. Assume L/K is an **abelian** extension. Then for a given unramified prime $\mathfrak{p}$ in K , the element $\text{Frob}_{\mathfrak{P}}$ doesn't depend on the choice of $\mathfrak{P}$. We denote the resulting $\text{Frob}_{\mathfrak{P}}$ by the **Artin symbol**.

$$\left(\frac{L/K}{\mathfrak{p}} \right).$$

The definition of the Artin symbol is written deliberately to look like the Legendre symbol. To see why:

Example 45.4.7 (Legendre symbol subsumed by Artin symbol)

Suppose we want to understand $(2/p) \equiv 2^{\frac{p-1}{2}}$ where $p > 2$ is prime. Consider the element

$$\left(\frac{\mathbb{Q}(\sqrt{2})/\mathbb{Q}}{p\mathbb{Z}} \right) \in \text{Gal}(\mathbb{Q}(\sqrt{2})/\mathbb{Q}).$$

It is uniquely determined by where it sends a . But in fact we have

$$\left(\frac{\mathbb{Q}(\sqrt{2})/\mathbb{Q}}{p\mathbb{Z}} \right) (\sqrt{2}) \equiv (\sqrt{2})^p \equiv 2^{\frac{p-1}{2}} \cdot \sqrt{2} \equiv \left(\frac{2}{p} \right) \sqrt{2} \pmod{\mathfrak{P}}$$

where $\left(\frac{2}{p} \right)$ is the usual Legendre symbol, and $\mathfrak{P}$ is above p in $\mathbb{Q}(\sqrt{2})$. Thus the Artin symbol generalizes the quadratic Legendre symbol.

Example 45.4.8 (Cubic Legendre symbol subsumed by Artin symbol)

Similarly, it also generalizes the cubic Legendre symbol. To see this, assume θ is primary in $K = \mathbb{Q}(\sqrt{-3}) = \mathbb{Q}(\omega)$ (thus $\mathcal{O}_K = \mathbb{Z}[\omega]$ is Eisenstein integers). Then for example

$$\left(\frac{K(\sqrt[3]{2})/K}{\theta \mathcal{O}_K} \right) \left(\sqrt[3]{2} \right) \equiv \left(\sqrt[3]{2} \right)^{N(\theta)} \equiv 2^{\frac{N\theta-1}{3}} \cdot \sqrt{2} \equiv \left(\frac{2}{\theta} \right)_3 \sqrt[3]{2} \pmod{\mathfrak{P}}$$

where $\mathfrak{P}$ is above p in $K(\sqrt[3]{2})$.

§45.5 Artin reciprocity

Now, we further capitalize on the fact that $\text{Gal}(L/K)$ is abelian. For brevity, in what follows let $\text{Ram}(L/K)$ denote the primes of K (either finite or infinite) which ramify in L .

Definition 45.5.1. Let L/K be an abelian extension and let $\mathfrak{m}$ be divisible by every prime in $\text{Ram}(L/K)$. Then since L/K is abelian we can extend the Artin symbol multiplicatively to a map

$$\left(\frac{L/K}{\bullet} \right) : I_K(\mathfrak{m}) \twoheadrightarrow \text{Gal}(L/K).$$

This is called the **Artin map**, and it is surjective (for example by Chebotarev Density). Thus we denote its kernel by

$$H(L/K, \mathfrak{m}) \subseteq I_K(\mathfrak{m}).$$

In particular we have $\text{Gal}(L/K) \cong I_K(\mathfrak{m})/H(\mathfrak{m})$.

We can now present the long-awaited Artin reciprocity theorem.

Theorem 45.5.2 (Artin reciprocity)

Let L/K be an abelian extension. Then there is a modulus $\mathfrak{f} = \mathfrak{f}(L/K)$, divisible by exactly the primes of $\text{Ram}(L/K)$, such that: for any modulus $\mathfrak{m}$ divisible by all primes of $\text{Ram}(L/K)$, we have

$$P_K(\mathfrak{m}) \subseteq H(L/K, \mathfrak{m}) \subseteq I_K(\mathfrak{m}) \quad \text{if and only if} \quad \mathfrak{f} \mid \mathfrak{m}.$$

We call $\mathfrak{f}$ the **conductor** of L/K .

So the conductor $\mathfrak{f}$ plays a similar role to the discriminant (divisible by exactly the primes which ramify), and when $\mathfrak{m}$ is divisible by the conductor, $H(L/K, \mathfrak{m})$ is a *congruence subgroup*.

Here's the reason this is called a "reciprocity" theorem. Recalling that $C_K(\mathfrak{f}) = I_K(\mathfrak{f})/P_K(\mathfrak{f})$, the above theorem tells us we get a sequence of maps

$$\begin{array}{ccccc} I_K(\mathfrak{f}) & \twoheadrightarrow & C_K(\mathfrak{f}) & \xrightarrow{\left(\frac{L/K}{\bullet} \right)} & \text{Gal}(L/K) \\ & & \searrow & \nearrow \cong & \\ & & I_K(\mathfrak{f})/H(L/K, \mathfrak{f}) & & \end{array}$$

Consequently:

For primes $\mathfrak{p} \in I_K(\mathfrak{f})$, $\left(\frac{L/K}{\mathfrak{p}}\right)$ depends only on “ $\mathfrak{p} \pmod{\mathfrak{f}}$ ”.

Let’s see how this result relates to quadratic reciprocity.

Example 45.5.3 (Artin reciprocity implies quadratic reciprocity)

The big miracle of quadratic reciprocity states that: for a fixed (squarefree) a , the Legendre symbol $\left(\frac{a}{p}\right)$ should only depend the residue of p modulo something. Let’s see why Artin reciprocity tells us this *a priori*.

Let $L = \mathbb{Q}(\sqrt{a})$, $K = \mathbb{Q}$. Then we’ve already seen that the Artin symbol

$$\left(\frac{\mathbb{Q}(\sqrt{a})/\mathbb{Q}}{\bullet}\right)$$

is the correct generalization of the Legendre symbol. Thus, Artin reciprocity tells us that there is a conductor $\mathfrak{f} = \mathfrak{f}(\mathbb{Q}(\sqrt{a})/\mathbb{Q})$ such that $\left(\frac{\mathbb{Q}(\sqrt{a})/\mathbb{Q}}{p}\right)$ depends only on the residue of p modulo $\mathfrak{f}$, which is what we wanted.

Here is an example along the same lines.

Example 45.5.4 (Cyclotomic field)

Let ζ be a primitive m th root of unity. For primes p , we know that $\text{Frob}_p \in \text{Gal}(\mathbb{Q}(\zeta)/\mathbb{Q})$ is “exactly” $p \pmod{m}$. Let’s translate this idea into the notation of Artin reciprocity.

We are going to prove

$$H(\mathbb{Q}(\zeta)/\mathbb{Q}, m\infty) = P_{\mathbb{Q}}(m\infty) = \left\{ \frac{a}{b}\mathbb{Z} \mid a/b \equiv 1 \pmod{m} \right\}.$$

This is the generic example of achieving the lower bound in Artin reciprocity. It also implies that $\mathfrak{f}(\mathbb{Q}(\zeta)/\mathbb{Q}) \mid m\infty$.

It’s well-known $\mathbb{Q}(\zeta)/\mathbb{Q}$ is unramified outside finite primes dividing m , so that the Artin symbol is defined on $I_K(\mathfrak{m})$. Now the Artin map is given by

$$\begin{aligned} I_{\mathbb{Q}}(\mathfrak{m}) &\xrightarrow{\left(\frac{\mathbb{Q}(\zeta)/\mathbb{Q}}{\bullet}\right)} \text{Gal}(\mathbb{Q}(\zeta)/\mathbb{Q}) \xrightarrow{\cong} (\mathbb{Z}/m\mathbb{Z})^{\times} \\ p &\longmapsto (x \mapsto x^p) \longmapsto p \pmod{m}. \end{aligned}$$

So we see that the kernel of this map is trivial, i.e. it is given by the identity of the Galois group, corresponding to $1 \pmod{m}$. On the other hand, we’ve also computed $P_{\mathbb{Q}}(m\infty)$ already, so we have the desired equality.

In fact, we also have the following “existence theorem”: every congruence subgroup appears uniquely once we fix $\mathfrak{m}$.

Theorem 45.5.5 (Tagaki existence theorem)

Fix K and let $\mathfrak{m}$ be a modulus. Consider any congruence subgroup H , i.e.

$$P_K(\mathfrak{m}) \subseteq H \subseteq I_K(\mathfrak{m}).$$

Then $H = H(L/K, \mathfrak{m})$ for a *unique* abelian extension L/K .

Finally, such subgroups reverse inclusion in the best way possible:

Lemma 45.5.6 (Inclusion-reversing congruence subgroups)

Fix a modulus $\mathfrak{m}$. Let L/K and M/K be abelian extensions and suppose $\mathfrak{m}$ is divisible by the conductors of L/K and M/K . Then

$$L \subseteq M \quad \text{if and only if} \quad H(M/K, \mathfrak{m}) \subseteq H(L/K, \mathfrak{m}).$$

Here by $L \subseteq M$ we mean that L is isomorphic to some subfield of M .

Sketch of proof. Let us first prove the equivalence with $\mathfrak{m}$ fixed. In one direction, assume $L \subseteq M$; one can check from the definitions that the diagram

$$\begin{array}{ccc} I_K(\mathfrak{m}) & \xrightarrow{\left(\frac{M/K}{\bullet}\right)} & \text{Gal}(M/K) \\ & \searrow \left(\frac{L/K}{\bullet}\right) & \downarrow \\ & & \text{Gal}(L/K) \end{array}$$

commutes, because it suffices to verify this for prime powers, which is just saying that Frobenius elements behave well with respect to restriction. Then the inclusion of kernels follows directly. The reverse direction is essentially the Takagi existence theorem. $\square$

Note that we can always take $\mathfrak{m}$ to be the product of conductors here.

To finish, here is a quote from Emil Artin on his reciprocity law:

I will tell you a story about the Reciprocity Law. After my thesis, I had the idea to define L -series for non-abelian extensions. But for them to agree with the L -series for abelian extensions, a certain isomorphism had to be true. I could show it implied all the standard reciprocity laws. So I called it the General Reciprocity Law and tried to prove it but couldn't, even after many tries. Then I showed it to the other number theorists, but they all laughed at it, and I remember Hasse in particular telling me it couldn't possibly be true. Still, I kept at it, but nothing I tried worked. Not a week went by — *for three years!* — that I did not try to prove the Reciprocity Law. It was discouraging, and meanwhile I turned to other things. Then one afternoon I had nothing special to do, so I said, 'Well, I try to prove the Reciprocity Law again.' So I went out and sat down in the garden. You see, from the very beginning I had the idea to use the cyclotomic fields, but they never worked, and now I suddenly saw that all this time I had been using them in the wrong way — and in half an hour I had it.

§45.6 Problems to think about

Problem 45A[†] (Kronecker-Weber theorem). Let L be an abelian extension of $\mathbb{Q}$. Then L is contained in a cyclic extension $\mathbb{Q}(\zeta)$ where ζ is an m th root of unity (for some m).

Problem 45B[†] (Hilbert class field). Let K be any number field. Then there exists a unique abelian extension E/K which is unramified at all primes (finite or infinite) and such that

- E/K is the maximal such extension by inclusion.
- $\text{Gal}(E/K)$ is isomorphic to the class group of E .
- A prime $\mathfrak{p}$ of K splits completely in E if and only if it is a principal ideal of $\mathcal{O}_K$.

We call E the **Hilbert class field** of K .

DRAFT (Evan Chen)
Updated February 15, 2017

XII

Representation Theory

46	Representations of algebras	431
46.1	Algebras	431
46.2	Representations	432
46.3	Direct sums	434
46.4	Irreducible and indecomposable representations	435
46.5	Morphisms of representations	436
46.6	The representations of $\text{Mat}_d(k)$	438
46.7	Problems to think about	439
47	Semisimple algebras	441
47.1	Schur's lemma continued	441
47.2	Density theorem	442
47.3	Semisimple algebras	444
47.4	Maschke's theorem	445
47.5	Example: the representations of $\mathbb{C}[S_3]$	445
47.6	Problems to think about	446
48	Characters	447
48.1	Definitions	447
48.2	The dual space modulo the commutator	448
48.3	Orthogonality of characters	449
48.4	Examples of character tables	451
48.5	Problems to think about	452
49	Some applications	454
49.1	Frobenius divisibility	454
49.2	Burnside's theorem	455
49.3	Frobenius determinant	456

46 Representations of algebras

In the 19th century, the word “group” hadn’t been invented yet; all work was done with subsets of $GL(n)$ or S_n . Only much later was the abstract definition of a group was given, an abstract set G which was an object in its own right.

While this abstraction is good for some reasons, it is often also useful to work with concrete representations. This is the subject of representation theory. Linear algebra is easier than abstract algebra, so if we can take a group G and represent it concretely as a set of matrices in $GL(n)$, this makes them easier to study. This is the *representation theory of groups*: how can we take a group and represent its elements as matrices?

§46.1 Algebras

Prototypical example for this section: $k[x_1, \dots, x_n]$ and $k[G]$.

Rather than working directly with groups from the beginning, it will be more convenient to deal with so-called k -algebras. This setting is more natural and general than that of groups, so once we develop the theory of algebras well enough, it will be fairly painless to specialize to the case of groups.

Colloquially,

An associative k -algebra is a possibly noncommutative ring with a copy of k inside it. It is thus a k -vector space.

I’ll present examples before the definition:

Example 46.1.1 (Examples of k -Algebras)

Let k be any field. The following are examples of k -algebras:

- (a) The field k itself.
- (b) The polynomial ring $k[x_1, \dots, x_n]$.
- (c) The set of $n \times n$ matrices with entries in k , which we denote by $\text{Mat}_n(k)$. Note the multiplication here is not commutative.
- (d) The set $\text{Mat}(V)$ of linear operators $T : V \rightarrow V$, with multiplication given by the composition of operators. (Here V is some vector space over k .) This is really the same as the previous example.

Definition 46.1.2. Let k be a field. A **k -algebra** A is a *possibly noncommutative* ring, equipped with an injective ring homomorphism $k \hookrightarrow A$ (whose image is the “copy of k ”). In particular, $1_k \mapsto 1_A$.

Thus we can consider k as a subset of A , and we then additionally require $\lambda \cdot a = a \cdot \lambda$ for each $\lambda \in k$ and $a \in A$.

If the multiplication operation is also commutative, then we say A is a **commutative algebra**.

Definition 46.1.3. Equivalently, a **k -algebra** A is a k -vector space which also has an associative, bilinear multiplication operation (with an identity 1_A). The “copy of k ” is obtained by considering elements $\lambda 1_A$ for each $\lambda \in k$ (i.e. scaling the identity by the elements of k , taking advantage of the vector space structure).

Abuse of Notation 46.1.4. Some other authors don’t require A to be associative or to have an identity, so to them what we have just defined is an “associative algebra with 1”. However, this is needlessly wordy for our purposes.

Example 46.1.5 (Group algebra)

The **group algebra** $k[G]$ is the k -vector space whose *basis elements* are the elements of a group G , and where the product of two basis elements is the group multiplication. For example, suppose $G = \mathbb{Z}_2 = \{1_G, x\}$. Then

$$k[G] = \{a1_G + bx \mid a, b \in k\}$$

with multiplication given by

$$(a1_G + bx)(c1_G + dx) = (ac + bd)1_G + (bc + ad)x.$$

Question 46.1.6. When is $k[G]$ commutative?

The example $k[G]$ is very important, because (as we will soon see) a representation of the algebra $k[G]$ amounts to a representation of the group G itself.

It is worth mentioning at this point that:

Definition 46.1.7. A **homomorphism** of k -algebras A, B is a linear map $T : A \rightarrow B$ which respects multiplication (i.e. $T(xy) = T(x)T(y)$) and which sends 1_A to 1_B . In other words, T is both a homomorphism as a ring and as a vector space.

Definition 46.1.8. Given k -algebras A and B , the **direct sum** $A \oplus B$ is defined as pairs $a + b$, where addition is done in the obvious way, but we declare $ab = 0$ for any $a \in A$ and $b \in B$.

Question 46.1.9. Show that $1_A + 1_B$ is the multiplicative identity of $A \oplus B$.

§46.2 Representations

Prototypical example for this section: $k[S_3]$ acting on $k^{\oplus 3}$ is my favorite.

Definition 46.2.1. A **representation** of a k -algebra A (also a **left A -module**) is:

- (i) A k -vector space V , and
- (ii) An *action* $\cdot$ of A on V : thus, for every $a \in A$ we can take $v \in V$ and act on it to get $a \cdot v$. This satisfies the usual axioms:
 - $(a + b) \cdot v = a \cdot v + b \cdot v$, $a \cdot (v + w) = a \cdot v + a \cdot w$, and $(ab) \cdot v = a \cdot (b \cdot v)$.
 - $\lambda \cdot v = \lambda v$ for $\lambda \in k$. In particular, $1_A \cdot v = v$.

Definition 46.2.2. The action of A can be more succinctly described as saying that there is a k -algebra homomorphism $\rho : A \rightarrow \text{Mat}(V)$. (So $a \cdot v = \rho(a)(v)$.) Thus we can also define a **representation** of A as a pair

$$(V, \rho : A \rightarrow \text{Mat}(V)).$$

This is completely analogous to how a group action G on a set X with n elements just amounts to a group homomorphism $G \rightarrow S_n$. From this perspective, what we are really trying to do is:

If A is an algebra, we are trying to *represent* the elements of A as matrices.

Abuse of Notation 46.2.3. While a representation is a pair (V, ρ) of *both* the vector space V and the action ρ , we frequently will just abbreviate it to “ V ”. This is probably one of the worst abuses I will commit, but everyone else does it and I fear the mob.

Abuse of Notation 46.2.4. Rather than $\rho(a)(v)$ we will just write $\rho(a)v$.

Example 46.2.5 (Representations of $\text{Mat}(V)$)

(a) Let $A = \text{Mat}_2(\mathbb{R})$. Then there is a representation $(\mathbb{R}^{\oplus 2}, \rho)$ where a matrix $a \in A$ just acts by $a \cdot v = \rho(a)(v) = a(v)$.

(b) More generally, given a vector space V over any field k , there is an obvious representation of $A = \text{Mat}(V)$ by $a \cdot v = \rho(a)(v) = a(v)$ (since $a \in \text{Mat}(V)$).

From the matrix perspective: if $A = \text{Mat}(V)$, then we can just represent A as matrices over V .

(c) There are other representations of $A = \text{Mat}_2(\mathbb{R})$. A silly example is the representation $(\mathbb{R}^{\oplus 4}, \rho)$ given by

$$\rho : \begin{bmatrix} a & b \\ c & d \end{bmatrix} \mapsto \begin{bmatrix} a & b & 0 & 0 \\ c & d & 0 & 0 \\ 0 & 0 & a & b \\ 0 & 0 & c & d \end{bmatrix}.$$

More abstractly, viewing $\mathbb{R}^{\oplus 4}$ as $(\mathbb{R}^{\oplus 2}) \oplus (\mathbb{R}^{\oplus 2})$, this is $a \cdot (v_1, v_2) = (a \cdot v_1, a \cdot v_2)$.

Example 46.2.6 (Representations of polynomial algebras)

(a) Let $A = k$. Then a representation of k is just any k -vector space V .

(b) If $A = k[x]$, then a representation (V, ρ) of A amounts to a vector space V plus the choice of a linear operator $T \in \text{Mat}(V)$ (by $T = \rho(x)$).

(c) If $A = k[x]/(x^2)$ then a representation (V, ρ) of A amounts to a vector space V plus the choice of a linear operator $T \in \text{Mat}(V)$ satisfying $T^2 = 0$.

(d) We can create arbitrary “functional equations” with this pattern. For example, if $A = k[x, y]/(x^2 - x + y, y^4)$ then representing A by V amounts to finding operators $S, T \in \text{Mat}(V)$ satisfying $S^2 = S - T$ and $T^4 = 0$.

Example 46.2.7 (Representations of groups)

(a) Let $A = \mathbb{R}[S_3]$. Then let

$$V = \mathbb{R}^{\oplus 3} = \{(x, y, z) \mid x, y, z \in \mathbb{R}\}.$$

We can let A act on V as follows: given a permutation $\pi \in S_3$, we permute the corresponding coordinates in V . So for example, if

$$\text{If } \pi = (1\ 2) \text{ then } \pi \cdot (x, y, z) = (y, x, z).$$

This extends linearly to let A act on V , by permuting the coordinates.

From the matrix perspective, what we are doing is representing the permutations in S_3 as permutation matrices on $k^{\oplus 3}$, like

$$(1\ 2) \mapsto \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

(b) More generally, let $A = k[G]$. Then a representation (V, ρ) of A amounts to a group homomorphism $\psi : G \rightarrow \text{GL}(V)$. (In particular, $\rho(1_G) = \text{id}_V$.) We call this a **group representation** of G .

Example 46.2.8 (Regular representation)

Any k -algebra A is a representation (A, ρ) over itself, with $a \cdot b = \rho(a)(b) = ab$ (i.e. multiplication given by A). This is called the **regular representation**, denoted $\text{Reg}(A)$.

§46.3 Direct sums

Prototypical example for this section: The example with $\mathbb{R}[S_3]$ seems best.

Definition 46.3.1. Let A be k -algebra and let $V = (V, \rho_V)$ and $W = (W, \rho_W)$ be two representations of A . Then $V \oplus W$ is a representation, with action ρ given by

$$a \cdot (v, w) = (a \cdot v, a \cdot w).$$

This representation is called the **direct sum** of V and W .

Example 46.3.2

Earlier we let $\text{Mat}_2(\mathbb{R})$ act on $\mathbb{R}^{\oplus 4}$ by

$$\rho : \begin{bmatrix} a & b \\ c & d \end{bmatrix} \mapsto \begin{bmatrix} a & b & 0 & 0 \\ c & d & 0 & 0 \\ 0 & 0 & a & b \\ 0 & 0 & c & d \end{bmatrix}.$$

So this is just a direct sum of two two-dimensional representations.

More generally, given representations (V, ρ_V) and (W, ρ_W) the representation ρ of $V \oplus W$ looks like

$$\rho(a) = \begin{bmatrix} \rho_V(a) & 0 \\ 0 & \rho_W(a) \end{bmatrix}.$$

Example 46.3.3 (Representation of S_n decomposes)

Let $A = \mathbb{R}[S_3]$ again, acting via permutation of coordinates on

$$V = \mathbb{R}^{\oplus 3} = \{(x, y, z) \mid x, y, z \in \mathbb{R}\}.$$

Consider the two subspaces

$$W_1 = \{(t, t, t) \mid t \in \mathbb{R}\}$$

$$W_2 = \{(x, y, z) \mid x + y + z = 0\}.$$

Note $V = W_1 \oplus W_2$ as vector spaces. But each of W_1 and W_2 is a subrepresentation (since the action of A keeps each W_i in place), so $V = W_1 \oplus W_2$ as representations too.

Direct sums also come up when we play with algebras.

Proposition 46.3.4 (Representations of $A \oplus B$ are $V_A \oplus V_B$)

Let A and B be k -algebras. Then every representation of $A \oplus B$ is of the form

$$V_A \oplus V_B$$

where V_A and V_B are representations of A and B , respectively.

Sketch of Proof. Let (V, ρ) be a representation of $A \oplus B$. For any $v \in V$, $\rho(1_A + 1_B)v = \rho(1_A)v + \rho(1_B)v$. One can then set $V_A = \{\rho(1_A)v \mid v \in V\}$ and $V_B = \{\rho(1_B)v \mid v \in V\}$. These are disjoint, since if $\rho(1_A)v = \rho(1_B)v'$, we have $\rho(1_A)v = \rho(1_A 1_B)v = \rho(1_A 1_B)v' = 0_V$, and similarly for the other side. $\square$

§46.4 Irreducible and indecomposable representations

Prototypical example for this section: $k[S_3]$ decomposes as the sum of two spaces.

One of the goals of representation theory will be to classify all possible representations of an algebra A . If we want to have a hope of doing this, then we want to discard “silly” representations such as

$$\rho : \begin{bmatrix} a & b \\ c & d \end{bmatrix} \mapsto \begin{bmatrix} a & b & 0 & 0 \\ c & d & 0 & 0 \\ 0 & 0 & a & b \\ 0 & 0 & c & d \end{bmatrix}$$

and focus our attention instead on “irreducible” representations. This motivates:

Definition 46.4.1. Let V be a representation of A . A **subrepresentation** $W \subseteq V$ is a subspace W with the property that for any $a \in A$ and $w \in W$, $a \cdot w = w$. In other words, this subspace is invariant under actions by A .

Thus if $V \neq W_1 \oplus W_2$ for representations W_1, W_2 then W_1 and W_2 are subrepresentations.

Definition 46.4.2. If V has no proper subrepresentations then it is **irreducible**. If $V \neq W_1 \oplus W_2$ for proper subrepresentations W_1, W_2 , then we say it is **indecomposable**.

Definition 46.4.3. For brevity, an **irrep** of an algebra/group is a *finite-dimensional* irreducible representation.

Example 46.4.4 (Representation of S_n decomposes)

Let $A = \mathbb{R}[S_3]$ again, acting via permutation of coordinates on

$$V = \mathbb{R}^{\oplus 3} = \{(x, y, z) \mid x, y, z \in \mathbb{R}\}.$$

Consider again the two subspaces

$$\begin{aligned} W_1 &= \{(t, t, t) \mid t \in \mathbb{R}\} \\ W_2 &= \{(x, y, z) \mid x + y + z = 0\}. \end{aligned}$$

As we've seen, $V = W_1 \oplus W_2$, and thus V is not irreducible. But one can show that W_1 and W_2 are irreducible (and hence indecomposable) as follows.

- For W_1 it's obvious, since W_1 is one-dimensional.
- For W_2 , consider any vector $w = (a, b, c)$ with $a + b + c = 0$ and not all zero. Then WLOG we can assume $a \neq b$ (since not all three coordinates are equal). In that case, $(1\ 2)$ sends w to $w' = (b, a, c)$. Then w and w' span W_2 .

Thus V breaks down completely into irreps.

Unfortunately, if W is a subrepresentation of V , then it is not necessarily the case that we can find a supplementary vector space W' such that $V = W \oplus W'$. Put another way, if V is reducible, we know that it has a subrepresentation, but a decomposition requires *two* subrepresentations. Here is a standard counterexample:

Exercise 46.4.5. Let $A = \mathbb{R}[x]$, and $V = \mathbb{R}^{\oplus 2}$ be the representation with action

$$\rho(x) = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}.$$

Show that the only subrepresentation is $W = \{(t, 0) \mid t \in \mathbb{R}\}$. So V is not irreducible, but it is indecomposable.

Here is a slightly more optimistic example, and the “prototypical example” that you should keep in mind.

Exercise 46.4.6. Let $A = \text{Mat}_d(k)$ and consider the obvious representation $k^{\oplus d}$ of A that we described earlier. Show that it is irreducible. (This is obvious if you understand the definitions well enough.)

§46.5 Morphisms of representations

We now proceed to define the morphisms between representations.

Definition 46.5.1. Let (V, ρ_V) and (W, ρ_W) be representations of A . An **intertwining operator**, or **morphism**, is a linear map $T : V \rightarrow W$ such that

$$T(a \cdot v) = a \cdot T(v)$$

for any $a \in A$, $v \in V$. (Note that the first $\cdot$ is the action of ρ_V and the second $\cdot$ is the action of ρ_W .) This is exactly what you expect if you think that V and W are “left A -modules”. If T is invertible, then it is an **isomorphism** of representations and we say $V \cong W$.

Remark 46.5.2 (For commutative diagram lovers). The condition $T(a \cdot v) = a \cdot T(v)$ can be read as saying that

$$\begin{array}{ccc} V & \xrightarrow{\rho_1(a)} & V \\ \downarrow T & & \downarrow T \\ W & \xrightarrow{\rho_2(a)} & W \end{array}$$

commutes for any $a \in A$.

Remark 46.5.3 (For category lovers). A representation is just a “bilinear” functor from an abelian one-object category $\{*\}$ (so $\text{Hom}(*, *) \cong A$) to the abelian category Vect_k . Then an intertwining operator is just a *natural transformation*.

Here are some examples of intertwining operators.

Example 46.5.4 (Intertwining operators)

- (a) For any $\lambda \in k$, the scalar map $T(v) = \lambda v$ is intertwining.
- (b) If $W \subseteq V$ is a subrepresentation, then the inclusion $W \hookrightarrow V$ is an intertwining operator.
- (c) The projection map $V_1 \oplus V_2 \rightarrow V_1$ is an intertwining operator.
- (d) Let $V = \mathbb{R}^{\oplus 2}$ and represent $A = k[x]$ by (V, ρ) where

$$\rho(x) = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}.$$

Thus $\rho(x)$ is rotation by 90° around the origin. Let T be rotation by 30° . Then $T : V \rightarrow V$ is intertwining (the rotations commute).

Exercise 46.5.5 (Kernel and image are subrepresentations). Let $T : V \rightarrow W$ be an intertwining operator.

- (a) Show that $\ker T \subseteq V$ is a subrepresentation of V .
- (b) Show that $\text{im } T \subseteq W$ is a subrepresentation of W .

The previous lemma gives us the famous Schur’s lemma.

Theorem 46.5.6 (Schur's lemma)

Let V and W be representations of a k -algebra A . Let $T : V \rightarrow W$ be a *nonzero* intertwining operator. Then

- (a) If V is irreducible, then T is injective.
- (b) If W is irreducible, then T is surjective.

In particular if both V and W are irreducible then T is an isomorphism.

An important special case is if k is algebraically closed: then the only intertwining operators $T : V \rightarrow V$ are multiplication by a constant.

Theorem 46.5.7 (Schur's lemma for algebraically closed fields)

Let k be an algebraically closed field. Let V be an irrep of a k -algebra A . Then any intertwining operator $T : V \rightarrow V$ is multiplication by a scalar.

Exercise 46.5.8. Use the fact that T has an eigenvalue λ to deduce this from Schur's lemma. (Consider $T - \lambda \cdot \text{id}_V$, and use Schur to deduce it's zero.)

We have already seen the counterexample of rotation by 90° for $k = \mathbb{R}$; this was the same counterexample we gave to the assertion that all linear maps have eigenvalues.

§46.6 The representations of $\text{Mat}_d(k)$

To give an example of the kind of progress already possible, we prove:

Theorem 46.6.1 (Representations of $\text{Mat}_d(k)$)

Let k be any field, d be a positive integer and let $W = k^{\oplus d}$ be the obvious representation of $A = \text{Mat}_d(k)$. Then the only finite-dimensional representations of $\text{Mat}_d(k)$ are $W^{\oplus n}$ for some positive integer n (up to isomorphism). In particular, it is irreducible if and only if $n = 1$.

For concreteness, I'll just sketch the case $d = 2$, since the same proof applies verbatim to other situations. This shows that the examples of representations of $\text{Mat}_2(\mathbb{R})$ we gave earlier are the only ones.

As we've said this is essentially a functional equation. The algebra $A = \text{Mat}_2(k)$ has basis given by four matrices

$$E_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \quad E_2 = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, \quad E_3 = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \quad E_4 = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$$

satisfying relations like $E_1 + E_2 = \text{id}_A$, $E_i^2 = E_i$, $E_1 E_2 = 0$, etc. So let V be a representation of A , and let $M_i = \rho(E_i)$ for each i ; we want to classify the possible matrices M_i on V satisfying the same functional equations. This is because, for example,

$$\text{id}_V = \rho(\text{id}_A) = \rho(E_1 + E_2) = M_1 + M_2.$$

By the same token $M_1 M_3 = M_3$. Proceeding in a similar way, we can obtain the following multiplication table:

$\times$	M_1	M_2	M_3	M_4	
M_1	M_1	0	M_3	0	
M_2	0	M_2	0	M_4	and $M_1 + M_2 = \text{id}_V$
M_3	0	M_3	0	M_1	
M_4	M_4	0	M_2	0	

Note that each M_i is a linear operator $V \rightarrow V$; for all we know, it could have hundreds of entries. Nonetheless, given the multiplication table of the basis E_i we get the corresponding table for the M_i .

So, in short, the problem is as follows:

Find all vector spaces V and quadruples of matrices M_i satisfying the multiplication table above.

Let $W_1 = M_1(V)$ and $W_2 = M_2(V)$ be the images of M_1 and M_2 .

Claim 46.6.2. $V = W_1 \oplus W_2$.

Proof. First, note that for any $v \in V$ we have

$$v = \rho(\text{id})(v) = (M_1 + M_2)v = M_1 v + M_2 v.$$

Moreover, we have that $W_1 \cap W_2 = \{0\}$, because if $M_1 v_1 = M_2 v_2$ then $M_1 v_1 = M_1(M_1 v_1) = M_1(M_2 v_2) = 0$. $\square$

Claim 46.6.3. $W_1 \cong W_2$.

Proof. Check that the maps

$$W_1 \xrightarrow{\times M_4} W_2 \quad \text{and} \quad W_2 \xrightarrow{\times M_3} W_1$$

are well-defined and mutually inverse. $\square$

Now, let $e_1, \dots, e_n$ be basis elements of W_1 ; thus $M_4 e_1, \dots, M_4 e_n$ are basis elements of W_2 . However, each $\{e_j, M_4 e_j\}$ forms a basis of a subrepresentation isomorphic to $W = k^{\oplus 2}$ (what's the isomorphism?).

This finally implies that all representations of A are of the form $W^{\oplus n}$. In particular, W is irreducible because there are no representations of smaller dimension at all!

§46.7 Problems to think about

Problem 46A[†]. Suppose we have *one-dimensional* representations $V_1 = (V_1, \rho_1)$ and $V_2 = (V_2, \rho_2)$ of A . Show that $V_1 \cong V_2$ if and only if $\rho_1(a)$ and $\rho_2(a)$ are multiplication by the same constant for every $a \in A$.

Problem 46B[†] (Schur's lemma for commutative algebras). Let A be a *commutative* algebra over an algebraically closed field k . Prove that any irrep of A is one-dimensional.

Problem 46C^{*}. Let (V, ρ) be a representation of A . Then $\text{Mat}(V)$ is a representation of A with action given by

$$a \cdot T = \rho(a) \circ T$$

for $T \in \text{Mat}(V)$.

- (a) Show that $\rho : \text{Reg}(A) \rightarrow \text{Mat}(V)$ is an intertwining operator.
- (b) If V is d -dimensional, show that $\text{Mat}(V) \cong V^{\oplus d}$ as representations of A .

Problem 46D*. Fix an algebra A . Find all intertwining operators

$$T : \text{Reg}(A) \rightarrow \text{Reg}(A).$$



Problem 46E. Let (V, ρ) be an *indecomposable* (not irreducible) representation of an algebra A . Prove that any intertwining operator $T : V \rightarrow V$ is either nilpotent or an isomorphism.

(Note that Theorem 46.5.7 doesn't apply, since the field k may not be algebraically closed.)

DRAFT (Evan Chen)
Updated February 15, 2017

47 Semisimple algebras

In what follows, assume the field k is algebraically closed.

Fix an algebra A and suppose you want to study its representations. We have a “direct sum” operation already. So, much like we pay special attention to prime numbers, we’re motivated to study irreducible representations and then build all the representations of A from there.

Unfortunately, we have seen (Exercise 46.4.5) that there exists a representation which is not irreducible, and yet cannot be broken down as a direct sum (indecomposable). This is *weird and bad*, so we want to give a name to representations which are more well-behaved. We say that a representation is **completely reducible** if it doesn’t exhibit this bad behavior.

Even better, we say a finite-dimensional algebra A is **semisimple** if all its finite-dimensional representations are completely reducible. So when we study finite-dimensional representations of semisimple algebras A , they exhibit the good property that we just have to figure out what the irreps are, and then piecing them together will give all the representations of A .

In fact, semisimple algebras A have even nicer properties. The culminating point of the chapter is when we prove that A is semisimple if and only if $A \cong \bigoplus_i \text{Mat}(V_i)$, where the V_i are the irreps of A (yes, there are only finitely many!).

§47.1 Schur’s lemma continued

Prototypical example for this section: For V irreducible, $\text{Hom}_{\text{rep}}(V^{\oplus 2}, V^{\oplus 2}) \cong k^{\oplus 4}$.

Definition 47.1.1. For an algebra A and representations V and W , we let $\text{Hom}_{\text{rep}}(V, W)$ be the set of intertwining operators between them. (It is also a k -algebra.)

By Schur’s lemma, we already know that if V and W are irreps, then

$$\text{Hom}_{\text{rep}}(V, W) \cong \begin{cases} k & \text{if } V \cong W \\ 0 & \text{if } V \not\cong W. \end{cases}$$

Can we say anything more? For example, it also tells us that

$$\text{Hom}_{\text{rep}}(V, V^{\oplus 2}) = k^{\oplus 2}.$$

The possible maps are $v \mapsto (c_1 v_1, c_2 v_2)$ for some choice of $c_1, c_2 \in k$.

More generally, suppose V is an irrep and consider $\text{Hom}_{\text{rep}}(V^{\oplus m}, V^{\oplus n})$. Intertwining operators are determined completely $T : V^{\oplus m} \rightarrow V^{\oplus n}$ by the mn choices of compositions

$$V \hookrightarrow V^{\oplus m} \xrightarrow{T} V^{\oplus n} \twoheadrightarrow V$$

where the first arrow is inclusion to the i th component of $V^{\oplus m}$ (for $1 \leq i \leq m$) and the second arrow is inclusion to the j th component of $V^{\oplus n}$ (for $1 \leq j \leq n$). However, by Schur’s lemma on each of these compositions, we know they must be constant.

Thus, $\text{Hom}_{\text{rep}}(V^{\oplus n}, V^{\oplus m})$ consist of $n \times m$ “matrices” of constants, and the map is provided by

$$\begin{bmatrix} c_{11} & c_{12} & \cdots & c_{1(n-1)} & c_{1n} \\ c_{21} & c_{22} & \cdots & c_{2(n-1)} & c_{2n} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ c_{m1} & c_{m2} & \cdots & c_{m(n-1)} & c_{mn} \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \\ \vdots \\ v_n \end{bmatrix} \in V^{\oplus m}$$

where the $c_{ij} \in k$ but $v_i \in V$; note the type mismatch! This is *not* just a linear map $V^{\oplus n} \rightarrow V^{\oplus m}$; rather, the outputs are m linear combinations of the inputs.

More generally, we have:

Theorem 47.1.2 (Schur’s lemma for completely reducible representations)

Let V and W be completely reducible representations, and set $V = \bigoplus V_i^{\oplus n_i}$, $W = \bigoplus V_i^{\oplus m_i}$ for integers $n_i, m_i \geq 0$, where each V_i is an irrep. Then

$$\text{Hom}_{\text{rep}}(V, W) \cong \bigoplus_i \text{Mat}_{n_i \times m_i}(k)$$

meaning that an intertwining operator $T : V \rightarrow W$ amounts to, for each i , an $n_i \times m_i$ matrix of constants which gives a map $V_i^{\oplus n_i} \rightarrow V_i^{\oplus m_i}$.

Corollary 47.1.3 (Subrepresentations of completely reducible representations)

Let $V = \bigoplus V_i^{\oplus n_i}$ be completely reducible. Then any subrepresentation W of V is isomorphic to $\bigoplus V_i^{\oplus m_i}$ where $m_i \leq n_i$ for each i , and the inclusion $W \hookrightarrow V$ is given by the direct sum of inclusion $V_i^{\oplus m_i} \hookrightarrow V_i^{\oplus n_i}$, which are $n_i \times m_i$ matrices.

Proof. Apply Schur’s lemma to the inclusion $W \hookrightarrow V$. □

§47.2 Density theorem

We are going to take advantage of the previous result to prove that finite-dimensional algebras have finitely many irreps.

Theorem 47.2.1 (Jacobson density theorem)

Let $(V_1, \rho_1), \dots, (V_r, \rho_r)$ be pairwise nonisomorphic finite-dimensional representations of A . Then there is a surjective map of vector spaces


$$\bigoplus_{i=1}^r \rho_i : A \rightarrow \bigoplus_{i=1}^r \text{Mat}(V_i).$$

The right way to think about this theorem is that

Density is the “Chinese remainder theorem” for irreps of A .

Recall that in number theory, the Chinese remainder theorem tells us that given lots of “unrelated” congruences, we can find a single N which simultaneously satisfies them all.

Similarly, given lots of different nonisomorphic representations of A , this means that we can select a single $a \in A$ which induces any tuple $(\rho_1(a), \dots, \rho_r(a))$ of actions we want — a surprising result, since even the $r = 1$ case is not obvious at all!



This also gives us the non-obvious corollary

Corollary 47.2.2 (Finiteness of number of representations)

Any finite-dimensional algebra A has at most $\dim A$ irreps.

Proof. If V_i are such irreps then $A \twoheadrightarrow \bigoplus_i V_i^{\oplus \dim V_i}$, hence we have the inequality $\sum (\dim V_i)^2 \leq \dim A$. $\square$

Proof of density theorem. Let $V = V_1 \oplus \dots \oplus V_r$, so A acts on $V = (V, \rho)$ by $\rho = \bigoplus_i \rho_i$. Thus by Problem 46C*, we can instead consider ρ as an *intertwining operator*

$$\rho : \text{Reg}(A) \rightarrow \bigoplus_{i=1}^r \text{Mat}(V_i) \cong \bigoplus_{i=1}^r V_i^{\oplus d_i}.$$

We will use this instead as it will be easier to work with.

First, we handle the case $r = 1$. Fix a basis $e_1, \dots, e_n$ of $V = V_1$. Assuming for contradiction that the map is not surjective. Then there is a map of representations (by ρ and the isomorphism) $\text{Reg}(A) \rightarrow V^{\oplus n}$ given by $a \mapsto (a \cdot e_1, \dots, a \cdot e_n)$. By hypothesis is not surjective: its image is a *proper* subrepresentation of $V^{\oplus n}$. Assume its image is isomorphic to $V^{\oplus m}$ for $m < n$, so by Theorem 47.1.2 there is a matrix of constants X with

$$\text{Reg}(A) \longrightarrow V^{\oplus n} \xleftarrow{X \cdot -} V^{\oplus m}$$

$$a \longmapsto (a \cdot e_1, \dots, a \cdot e_n)$$

$$1_A \longmapsto (e_1, \dots, e_n) \longleftarrow (v_1, \dots, v_m)$$

where the two arrows in the top row have the same image; hence the pre-image $(v_1, \dots, v_m)$ of $(e_1, \dots, e_n)$ can be found. But since $r < n$ we can find constants $c_1, \dots, c_n$ not all zero such that X applied to the column vector $(c_1, \dots, c_n)$ is zero:

$$\sum_{i=1}^n c_i e_i = [c_1 \ \dots \ c_n] \begin{bmatrix} e_1 \\ \vdots \\ e_n \end{bmatrix} = [c_1 \ \dots \ c_n] X \begin{bmatrix} v_1 \\ \vdots \\ v_m \end{bmatrix} = 0$$

contradicting the fact that e_i are linearly independent. Hence we conclude the theorem for $r = 1$.

As for $r \geq 2$, the image $\rho(A)$ is necessarily of the form $\bigoplus_i V_i^{\oplus r_i}$ (by Corollary 47.1.3) and by the above $r_i = \dim V_i$ for each i . $\square$

§47.3 Semisimple algebras

Definition 47.3.1. A finite-dimensional algebra A is a **semisimple** if every finite-dimensional representation of A is completely reducible.

Theorem 47.3.2 (Semisimple algebras)

Let A be a finite-dimensional algebra. Then the following are equivalent:

- (i) $A \cong \bigoplus_i \text{Mat}_{d_i}(k)$ for some d_i .
- (ii) A is semisimple.
- (iii) $\text{Reg}(A)$ is completely reducible.

Proof. (i) $\implies$ (ii) follows from Theorem 46.6.1 and Proposition 46.3.4. (ii) $\implies$ (iii) is tautological.

To see (iii) $\implies$ (i), we use the following clever trick. Consider

$$\text{Hom}_{\text{rep}}(\text{Reg}(A), \text{Reg}(A)).$$

On one hand, by Problem 46D*, it is isomorphic to A^{op} (A with opposite multiplication), because the only intertwining operators $\text{Reg}(A) \rightarrow \text{Reg}(A)$ are those of the form $-\cdot a$. On the other hand, suppose that we have set $\text{Reg}(A) = \bigoplus_i V_i^{\oplus n_i}$. By Theorem 47.1.2, we have

$$A^{\text{op}} \cong \text{Hom}_{\text{rep}}(\text{Reg}(A), \text{Reg}(A)) = \bigoplus_i \text{Mat}_{n_i \times n_i}(k).$$

But $\text{Mat}_n(k)^{\text{op}} \cong \text{Mat}_n(k)$ (just by transposing), so we recover the desired conclusion. $\square$

In fact, if we combine the above result with the density theorem (and Corollary 47.2.2), we obtain:

Theorem 47.3.3 (Sum of squares formula)

For a finite-dimensional algebra A we have

$$\sum_i \dim(V_i)^2 \leq \dim A$$

where the V_i are the irreps of A ; equality holds exactly when A is semisimple, in which case

$$\text{Reg}(A) \cong \bigoplus_i \text{Mat}(V_i) \cong \bigoplus_I V_i^{\oplus \dim V_i}.$$

Proof. The inequality was already mentioned in Corollary 47.2.2. It is equality if and only if the map $\rho : A \rightarrow \bigoplus_i \text{Mat}(V_i)$ is an isomorphism; this means all V_i are present. $\square$

Remark 47.3.4 (Digression). For any finite-dimensional A , the kernel of the map $\rho : A \rightarrow \bigoplus_i \text{Mat}(V_i)$ is denoted $\text{Rad}(A)$ and is the so-called **Jacobson radical** of A ; it's the set of all $a \in A$ which act by zero in all irreps of A . The usual definition of “semisimple” given in books is that this Jacobson radical is trivial.

§47.4 Maschke's theorem

We now prove that the representation theory of groups is as nice as possible.

Theorem 47.4.1 (Maschke's theorem)

Let G be a finite group, and k an algebraically closed field whose characteristic does not divide $|G|$. Then $k[G]$ is semisimple.

This tells us that when studying representations of groups, all representations are completely reducible.

Proof. Consider any finite-dimensional representation (V, ρ) of $k[G]$. Given a proper subrepresentation $W \subseteq V$, our goal is to construct a supplementary G -invariant subspace W' which satisfies

$$V = W \oplus W'.$$

This will show that indecomposable $\iff$ irreducible, which is enough to show $k[G]$ is semisimple.

Let $\pi : V \rightarrow W$ be any projection of V onto W , meaning $\pi(v) = v \iff v \in W$. We consider the *averaging* map $P : V \rightarrow V$ by

$$P(v) = \frac{1}{|G|} \sum_{g \in G} \rho(g^{-1}) \circ \pi \circ \rho(g).$$

We'll use the following properties of the map:

Exercise 47.4.2. Show that the map P satisfies:

- For any $w \in W$, $P(w) = w$.
- For any $v \in V$, $P(w) \in W$.
- The map $P : V \rightarrow V$ is an intertwining operator.

Thus P is idempotent (it is the identity on its image W), so by Problem 6C* we have $V = \ker P \oplus \operatorname{im} P$, but both $\ker P$ and $\operatorname{im} P$ are subrepresentations as desired. $\square$

Remark 47.4.3. In the case where $k = \mathbb{C}$, there is a shorter proof. Suppose $B : V \times V \rightarrow \mathbb{C}$ is an arbitrary bilinear form. Then we can “average” it to obtain a new bilinear form

$$\langle v, w \rangle \stackrel{\text{def}}{=} \frac{1}{|G|} \sum_{g \in G} B(g \cdot v, g \cdot w).$$

The averaged form $\langle -, - \rangle$ is G -invariant, in the sense that $\langle v, w \rangle = \langle g \cdot v, g \cdot w \rangle$. Then, one sees that if $W \subseteq V$ is a subrepresentation, so is its orthogonal complement $W^\perp$. This implies the result.

§47.5 Example: the representations of $\mathbb{C}[S_3]$

We compute all irreps of $\mathbb{C}[S_3]$. I'll take for granted right now there are exactly three such representations (which will be immediate by the first theorem in the next chapter: we'll in fact see that the number of representations of G is exactly equal to the number of conjugacy classes of G).

Given that, if the three representations have dimension d_1, d_2, d_3 , then we ought to have

$$d_1^2 + d_2^2 + d_3^2 = |G| = 6.$$

From this, combined with some deep arithmetic, we deduce that we should have $d_1 = d_2 = 1$ and $d_3 = 2$ or some permutation.

In fact, we can describe these representations explicitly. First, we define:

Definition 47.5.1. Let G be a group. The complex **trivial group representation** of a group G is the one-dimensional representation $\mathbb{C}_{\text{triv}} = (\mathbb{C}, \rho)$ where $g \cdot v = v$ for all $g \in G$ and $v \in \mathbb{C}$ (i.e. $\rho(g) = \text{id}$ for all $g \in G$).

Remark 47.5.2 (Warning). The trivial representation of an *algebra* A doesn't make sense for us: we might want to set $a \cdot v = v$ but this isn't linear in A . (You *could* try to force it to work by deleting the condition $1_A \cdot v = v$ from our definition; then one can just set $a \cdot v = 0$. But even then $\mathbb{C}_{\text{triv}}$ would not be the trivial representation of $k[G]$.)

Then the representations are:

- The one-dimensional $\mathbb{C}_{\text{triv}}$; each $\sigma \in S_3$ acts by the identity.
- There is a nontrivial one-dimensional representation $\mathbb{C}_{\text{sign}}$ where the map $S_3 \rightarrow \mathbb{C}^\times$ is given by sending σ to the sign of σ . Thus in $\mathbb{C}_{\text{sign}}$ every $\sigma \in S_3$ acts as ± 1 . Of course, $\mathbb{C}_{\text{triv}}$ and $\mathbb{C}_{\text{sign}}$ are not isomorphic (as one-dimensional representations are never isomorphic unless the constants they act on coincide for all a , as we saw in Problem 46A[†]).
- Finally, we have already seen the two-dimensional representation, but now we give it a name. Define refl_0 to be the representation whose vector space is $\{(x, y, z) \mid x + y + z = 0\}$, and whose action of S_3 on it is permutation of coordinates.

Exercise 47.5.3. Show that refl_0 is irreducible, for example by showing directly that no subspace is invariant under the action of S_3 .

Thus V is also not isomorphic to the previous two representations.

This implies that these are all the irreps of S_3 . Note that, if we take the representation V of S_3 on $k^{\oplus 3}$, we just get that $V = \text{refl}_0 \oplus \mathbb{C}_{\text{triv}}$.

§47.6 Problems to think about

Problem 47A. Find all the irreps of $\mathbb{C}[\mathbb{Z}_n]$.

Problem 47B (Maschke requires $|G|$ finite). Consider the representation of the group $\mathbb{R}$ on $\mathbb{C}^{\oplus 2}$ under addition by a homomorphism

$$\mathbb{R} \rightarrow \text{Mat}_2(\mathbb{C}) \quad \text{by} \quad t \mapsto \begin{bmatrix} 1 & t \\ 0 & 1 \end{bmatrix}.$$

Show that this representation is not irreducible, but it is indecomposable.

Problem 47C. Prove that all irreducible representations of a finite group are finite-dimensional.



Problem 47D. Determine all the complex irreps of D_{10} .

48 Characters

Characters are basically the best thing ever. To every representation V of A we will attach a so-called character $\chi_V : A \rightarrow k$. It will turn out that the characters of irreps of V will determine the representation V completely. Thus an irrep is just specified by a set of $\dim A$ numbers.

§48.1 Definitions

Definition 48.1.1. Let $V = (V, \rho)$ be a finite-dimensional representation of A . The **character** $\chi_V : A \rightarrow k$ attached to A is defined $\chi_V = \text{Tr} \circ \rho$, i.e.

$$\chi_V(a) \stackrel{\text{def}}{=} \text{Tr}(\rho(a) : V \rightarrow V).$$

Since Tr and ρ are additive, this is a k -linear map (but it is not multiplicative). Note also that $\chi_{V \oplus W} = \chi_V + \chi_W$ for any representations V and W .

We are especially interested in the case $A = k[G]$, of course. As usual, we just have to specify $\chi_V(g)$ for each $g \in S_3$ to get the whole map $k[G] \rightarrow k$. Thus we often think of χ_V as a function $G \rightarrow k$, called a character of the group G . Here is the case $G = S_3$:

Example 48.1.2 (Character table of S_3)

Let's consider the three irreps of $G = S_3$ from before. For $\mathbb{C}_{\text{triv}}$ all traces are 1; for $\mathbb{C}_{\text{sign}}$ the traces are ± 1 depending on sign (obviously, for one-dimensional maps $k \rightarrow k$ the trace "is" just the map itself). For refl_0 we take a basis $(1, 0, -1)$ and $(0, 1, -1)$, say, and compute the traces directly in this basis.

$\chi_V(g)$	id	(1 2)	(2 3)	(3 1)	(1 2 3)	(3 2 1)
$\mathbb{C}_{\text{triv}}$	1	1	1	1	1	1
$\mathbb{C}_{\text{sign}}$	1	-1	-1	-1	1	1
refl_0	2	0	0	0	-1	-1

The above table is called the **character table** of the group G . The table above has certain mysterious properties, which we will prove as the chapter progresses.

- (I) The value of $\chi_V(g)$ only depends on the conjugacy class of g .
- (II) The number of rows equals the number of conjugacy classes.
- (III) The sum of the squares of any row is 6 again!
- (IV) The "dot product" of any two rows is zero.

Abuse of Notation 48.1.3. The name "character" for $\chi_V : G \rightarrow k$ is a bit of a misnomer. This χ_V is not multiplicative in any way, as the above example shows: one can almost think of it as an element of $k^{\oplus |G|}$.

Question 48.1.4. Show that $\chi_V(1_A) = \dim V$, so one can read the dimensions of the representations from the leftmost column of a character table.

§48.2 The dual space modulo the commutator

For any algebra, we first observe that since $\text{Tr}(TS) = \text{Tr}(ST)$, we have for any V that

$$\chi_V(ab) = \chi_V(ba).$$

This explains observation (I) from earlier:

Question 48.2.1. Deduce that if g and h are in the same conjugacy class of a group G , and V is a representation of $\mathbb{C}[G]$, then $\chi(g) = \chi(h)$.

Now, given our algebra A we define the **commutator** $[A, A]$ to be the (two-sided) ideal¹ generated by elements of the form $xy - yx$. Thus $[A, A]$ is contained in the kernel of each χ_V .

Definition 48.2.2. The space $A/[A, A]$ is called the **abelianization** of A ; for brevity we denote it as A^{ab} . We think of this as “ A modulo the relation $ab = ba$ for each $a, b \in A$.”

So we can think of each character χ_V as an element of $(A^{\text{ab}})^\vee$.

Example 48.2.3 (Examples of abelianizations)

- (a) If A is commutative, then $[A, A] = \{0\}$ and $A^{\text{ab}} = A$.
- (b) If $A = \text{Mat}_k(d)$, then $[A, A]$ consists exactly of the $d \times d$ matrices of trace zero. (Proof: harmless exercise.) Consequently, A^{ab} is one-dimensional.
- (c) Suppose $A = k[G]$. We claim that $\dim A^{\text{ab}}$ is equal to the number of conjugacy classes of A . Indeed, an element of A can be thought of as just an arbitrary function $\xi : G \rightarrow k$. So an element of A^{ab} is a function $\xi : G \rightarrow k$ such that $\xi(gh) = \xi(hg)$ for every $g, h \in G$. This is equivalent to functions from conjugacy classes of G to k .

Theorem 48.2.4 (Character of representations of algebras)

Let A be an algebra over an algebraically closed field. Then

- (a) Characters of pairwise non-isomorphic irreps are linearly independent as elements of A^{ab} .
- (b) If A is finite-dimensional and semisimple, then the characters attached to irreps form a basis of A^{ab} .

In particular, in (b) the number of irreps of A equals $\dim A^{\text{ab}}$.

Proof. Part (a) is more or less obvious by the density theorem. Suppose there is a linear dependence, so that for every a we have

$$c_1\chi_{V_1}(a) + c_2\chi_{V_2}(a) + \cdots + c_r\chi_{V_r}(a) = 0$$

for some integer r .

¹This means the ideal consists of sums elements of the form $a(xy - yx)b$ for $a, b \in A$.

Question 48.2.5. Deduce that $c_1 = \cdots = c_r = 0$ from the density theorem.

For part (b), assume there are r irreps we may assume that

$$A = \bigoplus_{i=1}^r \text{Mat}(V_i)$$

where $V_1, \dots, V_r$ are the irreps of A . Since we have already showed the characters are linearly independent we need only show that $\dim(A/[A, A]) = r$, which follows from the observation earlier that each $\text{Mat}(V_i)$ has a one-dimensional abelianization. $\square$

Since G has $\dim \mathbb{C}[G]^{\text{ab}}$ conjugacy classes, this completes the proof of (II).

§48.3 Orthogonality of characters

Now we specialize to the case of finite groups G , represented over $\mathbb{C}$.

Definition 48.3.1. Let $\text{Classes}(G)$ denote the set conjugacy classes of G .

If G has r conjugacy classes, then it has r irreps. Each (finite-dimensional) representation V , irreducible or not, gives a character χ_V .

Abuse of Notation 48.3.2. From now on, we will often regard χ_V as a function $G \rightarrow \mathbb{C}$ or as a function $\text{Classes}(G) \rightarrow \mathbb{C}$. So for example, we will write both $\chi_V(g)$ (for $g \in G$) and $\chi_V(C)$ (for a conjugacy class C); the latter just means $\chi_V(g_C)$ for any representative $g_C \in C$.

Definition 48.3.3. Let $\text{Fun}_{\text{class}}(G)$ denote the set of functions $\text{Classes}(G) \rightarrow \mathbb{C}$ viewed as a vector space over $\mathbb{C}$. We endow it with the inner form

$$\langle f_1, f_2 \rangle = \frac{1}{|G|} \sum_{g \in G} f_1(g) \overline{f_2(g)}.$$

This is the same “dot product” that we mentioned at the beginning, when we looked at the character table of S_3 . We now aim to prove the following orthogonality theorem, which will imply (III) and (IV) from earlier.

Theorem 48.3.4 (Orthogonality)

For any finite-dimensional complex representations V and W of G we have

$$\langle \chi_V, \chi_W \rangle = \dim \text{Hom}_{\text{rep}}(W, V).$$

In particular, if V and W are irreps then

$$\langle \chi_V, \chi_W \rangle = \begin{cases} 1 & V \cong W \\ 0 & \text{otherwise.} \end{cases}$$

Corollary 48.3.5 (Irreps give an orthonormal basis)

The characters associated to irreps form an *orthonormal* basis of $\text{Fun}_{\text{class}}(G)$.

In order to prove this theorem, we have to define the dual representation and the tensor representation, which give a natural way to deal with the quantity $\chi_V(g)\overline{\chi_W(g)}$.

Definition 48.3.6. Let $V = (V, \rho)$ be a representation of G . The **dual representation** $V^\vee$ is the representation on $V^\vee$ with the action of G given as follows: for each $\xi \in V^\vee$, the action of g gives a $g \cdot \xi \in V^\vee$ specified by

$$v \mapsto g \cdot \xi (\rho(g^{-1})(v)).$$

Definition 48.3.7. Let $V = (V, \rho_V)$ and $W = (W, \rho_W)$ be group representations of G . The **tensor product** of V and W is the group representation on $V \otimes W$ with the action of G given on pure tensors by

$$g \cdot (v \otimes w) = (\rho_V(g)(v)) \otimes (\rho_W(g)(w))$$

which extends linearly to define the action of G on all of $V \otimes W$.

Remark 48.3.8. Warning: the definition for tensors does *not* extend to algebras. We might hope that $a \cdot (v \otimes w) = (a \cdot v) \otimes (a \cdot w)$ would work, but this is not even linear in $a \in A$ (what happens if we take $a = 2$, for example?).

Theorem 48.3.9 (Character traces)

If V and W are finite-dimensional representations of G , then for any $g \in G$.

- (a) $\chi_{V \oplus W}(g) = \chi_V(g) + \chi_W(g)$.
- (b) $\chi_{V \otimes W}(g) = \chi_V(g) \cdot \chi_W(g)$.
- (c) $\chi_{V^\vee}(g) = \overline{\chi_V(g)}$.

Proof. Parts (a) and (b) follow from the identities $\text{Tr}(S \oplus T) = \text{Tr}(S) + \text{Tr}(T)$ and $\text{Tr}(S \otimes T) = \text{Tr}(S) \text{Tr}(T)$. However, part (c) is trickier. As $(\rho(g))^{|G|} = \rho(g^{|G|}) = \rho(1_G) = \text{id}_V$ by Lagrange's theorem, we can diagonalize $\rho(g)$, say with eigenvalues $\lambda_1, \dots, \lambda_n$ which are $|G|$ th roots of unity, corresponding to eigenvectors $e_1, \dots, e_n$. Then we see that in the basis $e_1^\vee, \dots, e_n^\vee$, the action of g on $V^\vee$ has eigenvalues $\lambda_1^{-1}, \lambda_2^{-1}, \dots, \lambda_n^{-1}$. So

$$\chi_V(g) = \sum_{i=1}^n \lambda_i \quad \text{and} \quad \chi_{V^\vee}(g) = \sum_{i=1}^n \lambda_i^{-1} = \sum_{i=1}^n \overline{\lambda_i}$$

where the last step follows from the identity $|z| = 1 \iff z^{-1} = \overline{z}$. $\square$

Remark 48.3.10 (Warning). The identities (b) and (c) do not extend linearly to $\mathbb{C}[G]$, i.e. it is not true for example that $\chi_V(a) = \overline{\chi_V(a)}$ if we think of χ_V as a map $\mathbb{C}[G] \rightarrow \mathbb{C}$.

Proof of orthogonality relation. The key point is that we can now reduce the sums of products to just a single character by

$$\chi_V(g)\overline{\chi_W(g)} = \chi_{V \otimes W^\vee}(g).$$

So we can rewrite the sum in question as just

$$\langle \chi_V, \chi_W \rangle = \frac{1}{|G|} \sum_{g \in G} \chi_{V \otimes W^\vee}(g) = \chi_{V \otimes W^\vee} \left(\frac{1}{|G|} \sum_{g \in G} g \right).$$

Let $P : V \otimes W^\vee \rightarrow V \otimes W^\vee$ be the action of $\frac{1}{|G|} \sum_{g \in G} g$, so we wish to find $\text{Tr } P$.

Exercise 48.3.11. Show that P is idempotent. (Compute $P \circ P$ directly.)

Hence $V \otimes W^\vee = \ker P \oplus \operatorname{im} P$ (by Problem 6C*) and $\operatorname{im} P$ is subspace of elements which are fixed under G . From this we deduce that

$$\operatorname{Tr} P = \dim \operatorname{im} P = \dim \{x \in V \otimes W^\vee \mid g \cdot x = x \forall g \in G\}.$$

Now, consider the usual isomorphism $V \otimes W^\vee \rightarrow \operatorname{Hom}(W, V)$.

Exercise 48.3.12. Let $g \in G$. Show that under this isomorphism, $T \in \operatorname{Hom}(W, V)$ satisfies $g \cdot T = T$ if and only if $T(g \cdot w) = g \cdot T(w)$ for each $w \in W$. (This is just unwinding three or four definitions.)

Consequently, $\chi_{V \otimes W^\vee}(P) = \operatorname{Tr} P = \dim \operatorname{Hom}_{\operatorname{rep}}(W, V)$ as desired. $\square$

The orthogonality relation gives us a fast and mechanical way to check whether a finite-dimensional representation V is irreducible. Namely, compute the traces $\chi_V(g)$ for each $g \in G$, and then check whether $\langle \chi_V, \chi_V \rangle = 1$. So, for example, we could have seen the three representations of S_3 that we found were irreps directly from the character table. Thus, we can now efficiently verify any time we have a complete set of irreps.

§48.4 Examples of character tables

Example 48.4.1 (Dihedral group on 10 elements)

Let $D_{10} = \langle r, s \mid r^5 = s^2 = 1, rs = sr^{-1} \rangle$. Let $\omega = \exp(\frac{2\pi i}{5})$. We write four representations of D_{10} :

- $\mathbb{C}_{\operatorname{triv}}$, all elements of D_{10} act as the identity.
- $\mathbb{C}_{\operatorname{sign}}$, all elements of D_{10} act as the identity.
- V_1 , which is two-dimensional and given by $r \mapsto \begin{bmatrix} \omega & 0 \\ 0 & \omega^4 \end{bmatrix}$ and $s \mapsto \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$.
- V_2 , which is two-dimensional and given by $r \mapsto \begin{bmatrix} \omega^2 & 0 \\ 0 & \omega^3 \end{bmatrix}$ and $s \mapsto \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$.

We claim that these four representations are irreducible and pairwise non-isomorphic. We do so by writing the character table:

D_{10}	1	r, r^4	r^2, r^3	sr^k
$\mathbb{C}_{\operatorname{triv}}$	1	1	1	1
$\mathbb{C}_{\operatorname{sign}}$	1	1	1	-1
V_1	2	$\omega + \omega^4$	$\omega^2 + \omega^3$	0
V_2	2	$\omega^2 + \omega^3$	$\omega + \omega^4$	0

Then a direct computation shows the orthogonality relations, hence we indeed have an orthonormal basis. For example, $\langle \mathbb{C}_{\operatorname{triv}}, \mathbb{C}_{\operatorname{sign}} \rangle = 1 + 2 \cdot 1 + 2 \cdot 1 + 5 \cdot (-1) = 0$.

Example 48.4.2 (Character table of S_4)

We now have enough machinery to compute the character table of S_4 , which has five conjugacy classes (corresponding to cycle types id, 2, 3, 4 and 2 + 2). First of all, we note that it has two one-dimensional representations, $\mathbb{C}_{\text{triv}}$ and $\mathbb{C}_{\text{sign}}$, and these are the only ones (because there are only two homomorphisms $S_4 \rightarrow \mathbb{C}^\times$). So thus far we have the table

S_4	1	(• •)	(• • •)	(• • • •)	(• •)(• •)
$\mathbb{C}_{\text{triv}}$	1	1	1	1	1
$\mathbb{C}_{\text{sign}}$	1	-1	1	-1	1
$\vdots$			$\vdots$		

Note the columns represent $1 + 6 + 8 + 6 + 3 = 24$ elements.

Now, the remaining three representations have dimensions d_1, d_2, d_3 with

$$d_1^2 + d_2^2 + d_3^2 = 4! - 2 = 22$$

which has only $(d_1, d_2, d_3) = (2, 3, 3)$ and permutations. Now, we can take the refl₀ representation

$$\{(w, x, y, z) \mid w + x + y + z = 0\}$$

with basis $(1, 0, 0, -1)$, $(0, 1, 0, -1)$ and $(0, 0, 1, -1)$. This can be geometrically checked to be irreducible, but we can also do this numerically by computing the character directly (this is tedious): it comes out to have 3, -1, 0, 1, -1 which indeed gives norm

$$\langle \chi_{\text{refl}_0}, \chi_{\text{refl}_0} \rangle = \frac{1}{4!} \left(\underbrace{3^2}_{\text{id}} + \underbrace{6 \cdot (-1)^2}_{(\bullet \bullet)} + \underbrace{8 \cdot (0)^2}_{(\bullet \bullet \bullet)} + \underbrace{6 \cdot (1)^2}_{(\bullet \bullet \bullet \bullet)} + \underbrace{3 \cdot (-1)^2}_{(\bullet \bullet)(\bullet \bullet)} \right) = 1.$$

Note that we can also tensor this with the sign representation, to get another irreducible representation (since $\mathbb{C}_{\text{sign}}$ has all traces ± 1 , the norm doesn't change). Finally, we recover the final row using orthogonality (which we name $\mathbb{C}^2$, for lack of a better name); hence the completed table is as follows.

S_4	1	(• •)	(• • •)	(• • • •)	(• •)(• •)
$\mathbb{C}_{\text{triv}}$	1	1	1	1	1
$\mathbb{C}_{\text{sign}}$	1	-1	1	-1	1
$\mathbb{C}^2$	2	0	2	-1	0
refl ₀	3	-1	0	1	-1
refl ₀ $\otimes$ $\mathbb{C}_{\text{sign}}$	3	1	0	-1	-1

§48.5 Problems to think about

Problem 48A[†] (Reading decompositions from characters). Let W be a complex representation of a finite group G . Let $V_1, \dots, V_r$ be the complex irreps of G and set $n_i = \langle \chi_W, \chi_{V_i} \rangle$. Prove that each n_i is a positive integer and

$$W = \bigoplus_{i=1}^r V_i^{\oplus n_i}$$

Problem 48B. Consider complex representations of $G = S_4$. The representation $\text{refl}_0 \otimes \text{refl}_0$ is 9-dimensional, so it is clearly reducible. Compute its decomposition in terms of the five irreducible representations.

Problem 48C (Tensoring by one-dimensional irreps). Let V and W be irreps of G , with $\dim W = 1$. Show that $V \otimes W$ is irreducible.

Problem 48D (Quaternions). Compute the character table of the quaternion group Q_8 .



Problem 48E* (Second orthogonality formula). Let g and h be elements of a finite group G , and let $V_1, \dots, V_r$ be the irreps of G . Prove that

$$\sum_{i=1}^r \chi_{V_i}(g) \overline{\chi_{V_i}(h)} = \begin{cases} |Z(g)| & \text{if } g \text{ and } h \text{ are conjugates} \\ 0 & \text{otherwise.} \end{cases}$$

Here, $Z(g) = \{x \in G : xg = gx\}$ is the center of G .

DRAFT (Evan Chen)
Updated February 15, 2017

49 Some applications

With all this setup, we now take the time to develop some nice results which are of independent interest.

§49.1 Frobenius divisibility

Theorem 49.1.1 (Frobenius divisibility)

Let V be a complex irrep of a finite group G . Then $\dim V$ divides $|G|$.

The proof of this will require algebraic integers (developed in the algebraic number theory chapter). Recall that an *algebraic integer* is a complex number which is the root of a polynomial with integer coefficients, and that these algebraic integers form a ring $\bar{\mathbb{Z}}$ under addition and multiplication, and that $\bar{\mathbb{Z}} \cap \mathbb{Q} = \mathbb{Z}$.

First, we prove:

Lemma 49.1.2 (Elements of $\mathbb{Z}[G]$ are integral)

Let $\alpha \in \mathbb{Z}[G]$. Then there exists a monic polynomial P with integer coefficients such that $P(\alpha) = 0$.

Proof. Let A_k be the $\mathbb{Z}$ -span of $1, \alpha^1, \dots, \alpha^k$. Since $\mathbb{Z}[G]$ is Noetherian, the inclusions $A_0 \subseteq A_1 \subseteq A_2 \subseteq \dots$ cannot all be strict, hence $A_k = A_{k+1}$ for some k , which means α^{k+1} can be expressed in terms of lower powers of α . $\square$

Proof of Frobenius divisibility. Let $C_1, \dots, C_m$ denote the conjugacy classes of G . Then consider the rational number

$$\frac{|G|}{\dim V};$$

we will show it is an algebraic integer, which will prove the theorem. Observe that we can rewrite it as

$$\frac{|G|}{\dim V} = \frac{|G|}{\dim V} \langle \chi_V, \chi_V \rangle = \sum_{g \in G} \frac{\chi_V(g) \overline{\chi_V(g)}}{\dim V}.$$

We split the sum over conjugacy classes, so

$$\frac{|G|}{\dim V} = \sum_{i=1}^m \overline{\chi_V(C_i)} \cdot \frac{|C_i| \chi_V(C_i)}{\dim V}.$$

We claim that for every i ,

$$\frac{|C_i| \chi_V(C_i)}{\dim V} = \frac{1}{\dim V} \operatorname{Tr} T_i$$

is an algebraic integer, where

$$T_i \stackrel{\text{def}}{=} \rho \left(\sum_{h \in C_i} h \right).$$

To see this, note that T_i commutes with elements of G , and hence is an intertwining operator $T_i : V \rightarrow V$. Thus by Schur's lemma, $T_i = \lambda_i \cdot \text{id}_V$ and $\text{Tr } T = \lambda_i \dim V$. By Lemma 49.1.2, $\lambda_i \in \overline{\mathbb{Z}}$, as desired.

Now we are done, since $\overline{\chi_V(C_i)} \in \overline{\mathbb{Z}}$ too (it is the sum of conjugates of roots of unity), so $\frac{|G|}{\dim V}$ is the sum of products of algebraic integers, hence itself an algebraic integer. $\square$

§49.2 Burnside's theorem

We now prove a group-theoretic result. This is the famous poster child for representation theory (in the same way that RSA is the poster child of number theory) because the result is purely group theoretic.

Recall that a group is **simple** if it has no normal subgroups. In fact, we will prove:

Theorem 49.2.1 (Burnside)

Let G be a nonabelian group of order $p^a q^b$ (where $a, b \geq 0$). Then G is not simple.

In what follows p and q will always denote prime numbers.

Lemma 49.2.2 (On $\gcd(|C|, \dim V) = 1$)

Let $V = (V, \rho)$ be a complex irrep of G . Assume C is a conjugacy class of G with $\gcd(|C|, \dim V) = 1$. Then for any $g \in C$, either

- $\rho(g)$ is multiplication by a scalar, or
- $\chi_V(g) = \text{Tr } \rho(g) = 0$.

Proof. If ε_i are the n eigenvalues of $\rho(g)$ (which are roots of unity), then from the proof of Frobenius divisibility we know $\frac{|C|}{n} \chi_V(g) \in \overline{\mathbb{Z}}$, thus from $\gcd(|C|, n) = 1$ we get

$$\frac{1}{n} \chi_V(g) = \frac{1}{n} (\varepsilon_1 + \cdots + \varepsilon_n) \in \overline{\mathbb{Z}}.$$

So this follows readily from a fact from algebraic number theory, namely Problem 36E*: either $\varepsilon_1 = \cdots = \varepsilon_n$ (first case) or $\varepsilon_1 + \cdots + \varepsilon_n = 0$ (second case). $\square$

Lemma 49.2.3 (Simple groups don't have prime power conjugacy classes)

Let G be a finite simple group. Then G cannot have a conjugacy class of order p^k (where $k > 0$).

Proof. By contradiction. Assume C is such a conjugacy class, and fix any $g \in C$. By the second orthogonality formula (Problem 48E*) applied g and 1_G (which are not conjugate since $g \neq 1_G$) we have

$$\sum_{i=1}^r \dim V_i \chi_{V_i}(g) = 0$$

where V_i are as usual all irreps of G .

Exercise 49.2.4. Show that there exists a nontrivial irrep V such that $p \nmid \dim V$ and $\chi_V(g) \neq 0$. (Proceed by contradiction to show that $-\frac{1}{p} \in \mathbb{Z}$ if not.)

Let $V = (V, \rho)$ be the irrep mentioned. By the previous lemma, we now know that $\rho(g)$ acts as a scalar in V .

Now consider the subgroup

$$H = \langle ab^{-1} \mid a, b \in C \rangle \subseteq G.$$

We claim this is a nontrivial normal subgroup of G . It is easy to check H is normal, and since $|C| > 1$ we have that H is nontrivial. Each element of H acts trivially in G , so since V is nontrivial and irreducible $H \neq G$. This contradicts the assumption that G was simple. $\square$

With this lemma, Burnside's theorem follows by partitioning the $|G|$ elements of our group into conjugacy classes. Assume for contradiction G is simple. Each conjugacy class must have order either 1 (of which there are $|Z(G)|$) or divisible by pq , but on the other hand the sum equals $|G| = p^a q^b$. Consequently, we must have $|Z(G)| > 1$. But G is not abelian, hence $Z(G) \neq G$, thus the center $Z(G)$ is a nontrivial normal subgroup, contradicting the assumption that G was simple.

§49.3 Frobenius determinant

We finish with the following result, the problem that started the branch of representation theory. Given a finite group G , we create n variables $\{x_g\}_{g \in G}$, and an $n \times n$ matrix X_G whose (g, h) th entry is x_{gh} .

Example 49.3.1 (Frobenius determinants)

(a) If $G = \mathbb{Z}_2 = \langle T \mid T^2 = 1 \rangle$ then the matrix would be

$$X_G = \begin{bmatrix} x_{\text{id}} & x_T \\ x_T & x_{\text{id}} \end{bmatrix}.$$

Then $\det X_G = (x_{\text{id}} - x_T)(x_{\text{id}} + x_T)$.

(b) If $G = S_3$, a long computation gives the irreducible factorization of $\det X_G$ is

$$\left(\sum_{\sigma \in S_3} x_\sigma \right) \left(\sum_{\sigma \in S_3} \text{sign}(\sigma) x_\sigma \right) \left(F(x_{\text{id}}, x_{(123)}, x_{(321)}) - F(x_{(12)}, x_{(23)}, x_{(31)}) \right)^2$$

where $F(a, b, c) = a^2 + b^2 + c^2 - ab - bc - ca$; the latter factor is irreducible.

Theorem 49.3.2 (Frobenius determinant)

The polynomial $\det X_G$ (in $|G|$ variables) factors into a product of irreducible polynomials such that

- (i) The number of polynomials equals the number of conjugacy classes of G , and
- (ii) The multiplicity of each polynomial equals its degree.

You may already be able to guess how the “sum of squares” result is related! (Indeed, look at $\deg \det X_G$.)

Legend has it that Dedekind observed this behavior first in 1896. He didn’t know how to prove it in general, so he sent it in a letter to Frobenius, who created representation theory to solve the problem.

With all the tools we’ve built, it is now fairly straightforward to prove the result.

Proof. Let $V = (V, \rho) = \text{Reg}(\mathbb{C}[G])$ and let $V_1, \dots, V_r$ be the irreps of G . Let’s consider the map $T : \mathbb{C}[G] \rightarrow \mathbb{C}[G]$ which has matrix X_G in the usual basis of $\mathbb{C}[G]$, namely

$$T = T(\{x_g\}) = \sum_{g \in G} x_g \rho(g) \in \text{Mat}(V).$$

Thus we want to examine $\det T$.

But we know that $V = \bigoplus_{i=1}^r V_i^{\oplus \dim V_i}$ as before, and so breaking down T over its subspaces we know

$$\det T = \prod_{i=1}^r (\det(T|_{V_i}))^{\dim V_i}.$$

So we only have to show two things: the polynomials $\det T|_{V_i}$ are irreducible, and they are pairwise different for different i .

Let $V_i = (V_i, \rho_i)$, and pick $k = \dim V_i$.

- *Irreducible:* By the density theorem, for any $M \in \text{Mat}(V_i)$ there exists a *particular* choice of complex numbers $x_g \in G$ such that

$$M = \sum_{g \in G} x_g \cdot \rho_i(g) = (T|_{V_i})(\{x_g\}).$$

View $\rho_i(g)$ as a $k \times k$ matrix with complex coefficients. Thus the “generic” $(T|_{V_i})(\{x_g\})$, viewed as a matrix with polynomial entries, must have linearly independent entries (or there would be some matrix in $\text{Mat}(V_i)$ that we can’t achieve).

Then, the assertion follows (by a linear variable change) from the simple fact that the polynomial $\det(y_{ij})_{1 \leq i, j \leq m}$ in m^2 variables is always irreducible.

- *Pairwise distinct:* We show that from $\det T|_{V_i}(\{x_g\})$ we can read off the character χ_{V_i} , which proves the claim. In fact

Exercise 49.3.3. Pick *any* basis for V_i . If $\dim V_i = k$, and $1_G \neq g \in G$, then

$$\chi_{V_i}(g) \text{ is the coefficient of } x_g^{k-1} x_{1_G}.$$

Thus, we are done. □

XIII

Algebraic Geometry I: Varieties

50 Affine varieties	459
50.1 Affine varieties	459
50.2 Naming affine varieties via ideals	460
50.3 Radical ideals and Hilbert's Nullstellensatz	461
50.4 Pictures of varieties in $\mathbb{A}^1$	462
50.5 Prime ideals correspond to irreducible affine varieties	463
50.6 Pictures in $\mathbb{A}^2$ and $\mathbb{A}^3$	464
50.7 Maximal ideals	465
50.8 Why schemes?	466
50.9 Problems to think about	466
51 Affine varieties as ringed spaces	467
51.1 Synopsis	467
51.2 The Zariski topology on $\mathbb{A}^n$	467
51.3 The Zariski topology on affine varieties	469
51.4 Coordinate rings	470
51.5 The sheaf of regular functions	471
51.6 Regular functions on distinguished open sets	472
51.7 Baby ringed spaces	474
51.8 Problems to think about	474
52 Projective varieties	475
52.1 Graded rings	475
52.2 The ambient space	476
52.3 Homogeneous ideals	477
52.4 As ringed spaces	478
52.5 Examples of regular functions	480
52.6 Problems to think about	481
53 Bonus: Bézout's theorem	482
53.1 Non-radical ideals	482
53.2 Hilbert functions of finitely many points	483
53.3 Hilbert polynomials	485
53.4 Bézout's theorem	487
53.5 Applications	487
53.6 Problems to think about	488

50 Affine varieties

In this chapter we introduce affine varieties. We introduce them in the context of coordinates, but over the course of the other chapters we'll gradually move away from this perspective to viewing varieties as “intrinsic objects”, rather than embedded in coordinates.


For simplicity, we'll do almost everything over the field of complex numbers, but the discussion generalizes to any algebraically closed field.

§50.1 Affine varieties

Prototypical example for this section: $\mathcal{V}(y - x^2)$ is a parabola.

Definition 50.1.1. Given a set of polynomials $S \subseteq \mathbb{C}[x_1, \dots, x_n]$ (not necessarily finite or even countable), we let $\mathcal{V}(S)$ denote the set of points vanishing on *all* the polynomials in S . Such a set is called an **affine variety**. It lives in **n -dimensional affine space**, denoted $\mathbb{A}^n$ (to distinguish it from projective space later).

For example, a parabola is the zero locus of the polynomial $\mathcal{V}(y - x^2)$. Picture:



Example 50.1.2 (Examples of affine varieties)

These examples are in two-dimensional space $\mathbb{A}^2$, whose points are pairs (x, y) .

- (a) A straight line can be thought of as $\mathcal{V}(Ax + By + C)$.
- (b) A parabola as above can be pictured as $\mathcal{V}(y - x^2)$.
- (c) A hyperbola might be the zero locus of the polynomial $\mathcal{V}(xy - 1)$.
- (d) The two axes can be thought of as $\mathcal{V}(xy)$; this is the set of points such that $x = 0$ or $y = 0$.
- (e) A point (x_0, y_0) can be thought of as $\mathcal{V}(x - x_0, y - y_0)$.
- (f) The entire space $\mathbb{A}^2$ can be thought of as $\mathcal{V}(0)$.
- (g) The empty set is the zero locus of the constant polynomial 1, that is $\mathcal{V}(1)$.

§50.2 Naming affine varieties via ideals

Prototypical example for this section: $\mathcal{V}(I)$ is a parabola, where $I = (y - x^2)$.

As you might have already noticed, a variety can be named by $\mathcal{V}(-)$ in multiple ways. For example, the set of solutions to

$$x = 3 \text{ and } y = 4$$

is just the point $(3, 4)$. But this is also the set of solutions to

$$x = 3 \text{ and } y = x + 1.$$

So, for example

$$\{(3, 4)\} = \mathcal{V}(x - 3, y - 4) = \mathcal{V}(x - 3, y - x - 1).$$

That's a little annoying, because in an ideal world we would have *one* name for every variety. Let's see if we can achieve this.

A partial solution is to use *ideals* rather than small sets. That is, consider the ideal

$$I = (x - 3, y - 4) = \{p(x, y) \cdot (x - 3) + q(x, y) \cdot (y - 4) \mid p, q \in \mathbb{C}[x, y]\}$$

and look at $\mathcal{V}(I)$.

Question 50.2.1. Convince yourself that $\mathcal{V}(I) = \{(3, 4)\}$.

So rather than writing $\mathcal{V}(x - 3, y - 4)$ it makes sense to think about this as $\mathcal{V}(I)$, where $I = (x - 3, y - 4)$ is the *ideal* generated by the two polynomials $x - 3$ and $y - 4$. This is an improvement because

Question 50.2.2. Check that $(x - 3, y - x - 1) = (x - 3, y - 4)$.

Needless to say, this pattern holds in general.

Question 50.2.3. Let $\{f_i\}$ be a set of polynomials, and consider the ideal I generated by these $\{f_i\}$. Show that $\mathcal{V}(\{f_i\}) = \mathcal{V}(I)$.

Thus we will only consider $\mathcal{V}(I)$ when I is an ideal. Of course, frequently our ideals are generated by one or two polynomials, which leads to:

Abuse of Notation 50.2.4. Given a set of polynomials $f_1, \dots, f_m$ we let $\mathcal{V}(f_1, \dots, f_m)$ be shorthand for $\mathcal{V}((f_1, \dots, f_m))$. In other words we let $\mathcal{V}(f_1, \dots, f_m)$ abbreviate $\mathcal{V}(I)$, where I is the *ideal* $I = (f_1, \dots, f_m)$.

This is where the Noetherian condition really shines: it guarantees that every ideal $I \subseteq \mathbb{C}[x_1, \dots, x_n]$ can be written in the form above with *finitely* many polynomials, because it is *finitely generated*. (The fact that $\mathbb{C}[x_1, \dots, x_n]$ is Noetherian follows from the Hilbert basis theorem, which is Theorem 12.6.6). This is a relief, because dealing with infinite sets of polynomials is not much fun.

§50.3 Radical ideals and Hilbert's Nullstellensatz

Prototypical example for this section: $\sqrt{(x^2)} = (x)$ in $\mathbb{C}[x]$, $\sqrt{(12)} = (6)$ in $\mathbb{Z}$.

You might ask whether the name is unique now: that is, if $\mathcal{V}(I) = \mathcal{V}(J)$, does it follow that $I = J$? The answer is unfortunately no: the counterexample can be found in just $\mathbb{A}^1$. It is

$$\mathcal{V}(x) = \mathcal{V}(x^2).$$

In other words, the set of solutions to $x = 0$ is the same as the set of solutions to $x^2 = 0$.

Well, that's stupid. We want an operation which takes the ideal (x^2) and makes it into the ideal (x) . The way to do so is using the radical of an ideal.

Definition 50.3.1. Let R be a ring. The **radical** of an ideal $I \subseteq R$, denoted $\sqrt{I}$, is defined by

$$\sqrt{I} = \{r \in R \mid r^m \in I \text{ for some integer } m \geq 1\}.$$

For example, $\sqrt{(x^2)} = (x)$. You may like to take the time to verify that $\sqrt{I}$ is actually an ideal.

This is actually the same as the notion of “radical” in number theory. In $\mathbb{Z}$, the radical of an ideal (n) corresponds to just removing all the duplicate prime factors, so for example

$$\sqrt{(12)} = (6).$$

In particular, if you try to take $\sqrt{(6)}$, you just get (6) back; you don't squeeze out any new prime factors. And that's true in general.

Definition 50.3.2. An ideal I is called **radical** if $\sqrt{I} = I$.

Question 50.3.3. Show that $\sqrt{I}$ is always radical.

Question 50.3.4. Prime ideals are radical.

Useful equivalent definition (more in line with the number theory example), which we quote without proof:

Theorem 50.3.5 (Radical is intersection of prime ideals)

For any ideal I , we have

$$\sqrt{I} = \bigcap_{I \subseteq \mathfrak{p} \text{ prime}} \mathfrak{p}.$$

In any case, we have:

Question 50.3.6. Verify that $\mathcal{V}(I) = \mathcal{V}(\sqrt{I})$.

This makes sense: you would never refer to $x = 0$ as $x^2 = 0$. With this, we obtain a theorem called Hilbert's Nullstellensatz.

Theorem 50.3.7 (Hilbert's Nullstellensatz)

Given an affine variety $V = \mathcal{V}(I)$, the set of polynomials which vanish on all points of V is precisely $\sqrt{I}$. Thus if I and J are ideals in $\mathbb{C}[x_1, \dots, x_n]$, then

$$\mathcal{V}(I) = \mathcal{V}(J) \text{ if and only if } \sqrt{I} = \sqrt{J}.$$

In other words

Radical ideals in $\mathbb{C}[x_1, \dots, x_n]$ correspond exactly to n -dimensional affine varieties.

The proof of Hilbert's Nullstellensatz will be given in Problem 50B; for now it is worth remarking that it relies essentially on the fact that $\mathbb{C}$ is *algebraically closed*. For example, it is false in $\mathbb{R}[x]$, with $(x^2 + 1)$ being a maximal ideal with empty vanishing set.

§50.4 Pictures of varieties in $\mathbb{A}^1$

Prototypical example for this section: Finite sets of points (in fact these are the only nontrivial examples).

Let's first draw some pictures. In what follows I'll draw $\mathbb{C}$ as a straight line... sorry.

First of all, let's look at just the complex line $\mathbb{A}^1$. What are the various varieties on it? For starters, we have a single point $9 \in \mathbb{C}$, generated by $(x - 9)$.

$$\begin{array}{c} \mathbb{A}^1 \qquad \qquad \qquad 9 \\ \leftarrow \text{-----} \rightarrow \\ \mathcal{V}(x - 9) \end{array}$$

Another example is the point 4. And in fact, if we like we can get an ideal consisting of just these two points; consider $\mathcal{V}((x - 4)(x - 9))$.

$$\begin{array}{c} \mathbb{A}^1 \qquad \qquad \qquad 4 \qquad \qquad 9 \\ \leftarrow \text{-----} \rightarrow \\ \mathcal{V}((x - 4)(x - 9)) \end{array}$$

In general, in $\mathbb{A}^1$ you can get finitely many points $\{a_1, \dots, a_n\}$ by just taking

$$\mathcal{V}((x - a_1)(x - a_2) \dots (x - a_n)).$$

On the other hand, you can't get the set $\{0, 1, 2, \dots\}$ as an affine variety; the only polynomial vanishing on all those points is the zero polynomial. In fact, you can convince yourself that these are the only affine varieties, with two exceptions:

- The entire line $\mathbb{A}^1$ is given by $\mathcal{V}(0)$, and
- The empty set is given by $\mathcal{V}(1)$.

Question 50.4.1. Show that these are the only varieties of $\mathbb{A}^1$. (Let $\mathcal{V}(I)$ be the variety and pick a $0 \neq f \in I$.)

As you might correctly guess, we have:

Theorem 50.4.2 (Intersections and unions of varieties)

- (a) The intersection of affine varieties (even infinitely many) is an affine variety.
 (b) The union of finitely many affine varieties is an affine variety.

In fact we have

$$\bigcap_{\alpha} \mathcal{V}(I_{\alpha}) = \mathcal{V}\left(\sum_{\alpha} I_{\alpha}\right) \quad \text{and} \quad \bigcup_{k=1}^n \mathcal{V}(I_k) = \mathcal{V}\left(\bigcap_{k=1}^n I_k\right).$$

You are welcome to prove this easy result yourself.

Remark 50.4.3. Part (a) is a little misleading in that the sum $I + J$ need not be radical: take for example $I = (y - x^2)$ and $J = (y)$ in $\mathbb{C}[x, y]$, where $x \in \sqrt{I + J}$ and $x \notin I + J$. But in part (b) for radical ideals I and J , the intersection $I \cap J$ is radical.

Now note that most of the affine varieties of $\mathbb{A}^1$, like $\{4, 9\}$, are just unions of the simplest “one-point” ideals. To ease our classification, we can restrict our attention to the case of *irreducible* varieties;

Definition 50.4.4. A variety V is **irreducible** if it cannot be written as the nontrivial union of two varieties $V = V_1 \cup V_2$. (Here nontrivial means that $V_1, V_2 \subsetneq V$.)

Hence, the irreducible varieties of $\mathbb{A}^1$ are

- (i) The empty set,
- (ii) a single point, and
- (iii) the entire line $\mathbb{A}^1$.

In a moment, we’ll draw pictures in $\mathbb{A}^2$, but first let me point out something about irreducible affine varieties.

§50.5 Prime ideals correspond to irreducible affine varieties

Prototypical example for this section: (xy) corresponds to the union of two lines in $\mathbb{A}^2$.

We have seen that the radical ideals are in one-to-one correspondence with affine varieties. In the next sections we answer the two questions:

- What property of $\mathcal{V}(I)$ corresponds to I being prime?
- What property of $\mathcal{V}(I)$ corresponds to I being maximal?

The first question is easier to answer. Let’s take a standard non-prime ideal in $\mathbb{C}[x, y]$, such as $I = (xy)$. Its vanishing set $\mathcal{V}(I)$ is the union of two lines $x = 0$ and $y = 0$. So $\mathcal{V}(I)$ is reducible.

Let’s make this work in general.

Theorem 50.5.1 (Prime $\iff$ irreducible)

Let I be a radical ideal, and $V = \mathcal{V}(I)$ a nonempty variety. Then I is prime if and only if V is irreducible.

Proof. First, assume V is irreducible; we'll show I is prime. Let $f, g \in \mathbb{C}[x_1, \dots, x_n]$ so that $fg \in I$. Then V is a subset of the union $\mathcal{V}(f) \cup \mathcal{V}(g)$; actually, $V = (V \cap \mathcal{V}(f)) \cup (V \cap \mathcal{V}(g))$. Since V is irreducible, we may assume $V = V \cap \mathcal{V}(f)$, hence f vanishes on all of V . So $f \in I$.

The reverse direction is similar. $\square$

Remark 50.5.2. The above proof illustrates the following principle: Let V be an irreducible variety. Suppose that $V \subseteq V_1 \cup V_2$; this implies $V = (V_1 \cap V) \cup (V_2 \cap V)$. Recall that the intersection of two varieties is a variety. Thus an irreducible variety can't even be *contained* in a nontrivial union of two varieties.

§50.6 Pictures in $\mathbb{A}^2$ and $\mathbb{A}^3$


Prototypical example for this section: Various curves and hypersurfaces.

Next, we turn our attention to the “complex affine plane”, $\mathbb{A}^2$. What are the irreducible affine varieties in it?

As we saw in the previous discussion, naming irreducible affine varieties in $\mathbb{A}^2$ amounts to naming the prime ideals of $\mathbb{C}[x, y]$. Here are a few.

- (i) (0) is prime. $\mathcal{V}(0)$ as usual corresponds to the entire plane.
- (ii) $(x - a, y - b)$ is a prime ideal, since $\mathbb{C}[x, y]/(x - a, y - b) \cong \mathbb{C}$ is an integral domain. (In fact, since $\mathbb{C}$ is a field, the ideal $(x - a, y - b)$ is *maximal*). The vanishing set of this is $\mathcal{V}(x - a, y - b) = \{(a, b)\} \in \mathbb{C}^2$, so these ideals correspond to a single point.
- (iii) Let $f(x, y)$ be an irreducible polynomial, like $y - x^2$. Then (f) is a prime ideal! Here $\mathcal{V}(I)$ is a “degree one curve”.

By using some polynomial algebra (again you're welcome to check this; Euclidean algorithm), these are in fact the only prime ideals of $\mathbb{C}[x, y]$. Here's a picture.



As usual, you can make varieties which are just unions of these irreducible ones. For example, if you wanted the variety consisting of a parabola $y = x^2$ plus the point $(20, 15)$ you would write

$$\mathcal{V}((y - x^2)(x - 20), (y - x^2)(y - 15)).$$

The picture in $\mathbb{A}^3$ is harder to describe. Again, you have points $\mathcal{V}(x - a, y - b, z - c)$ corresponding to be zero-dimensional points (a, b, c) , and two-dimensional surfaces $\mathcal{V}(f)$ for each irreducible polynomial f (for example, $x + y + z = 0$ is a plane). But there are more prime ideals, like $\mathcal{V}(x, y)$, which corresponds to the intersection of the planes $x = 0$

and $y = 0$: this is the one-dimensional z -axis. It turns out there is no reasonable way to classify the “one-dimensional” varieties; they correspond to “irreducible curves”.

Thus, as Ravi Vakil [Va15] says,

The purely algebraic question of determining the prime ideals of $\mathbb{C}[x, y, z]$ has a fundamentally geometric answer.

§50.7 Maximal ideals

Prototypical example for this section: All maximal ideals are $(x_1 - a_1, \dots, x_n - a_n)$.

Question 50.7.1. Show that if $I \subseteq J$ then $\mathcal{V}(I) \supseteq \mathcal{V}(J)$. Thus $\mathcal{V}(-)$ is *inclusion-reversing*.

Thus, bigger ideals correspond to smaller varieties. As the above pictures might have indicated, the smallest varieties are *single points*. Moreover, as you might guess from the name, the biggest ideals are the *maximal ideals*. As an example, all ideals of the form

$$(x_1 - a_1, \dots, x_n - a_n)$$

are maximal, since the quotient

$$\mathbb{C}[x_1, \dots, x_n] / (x_1 - a_1, \dots, x_n - a_n) \cong \mathbb{C}$$

is a field. The question is: are all maximal ideals of this form?

The answer is in the affirmative. It’s equivalent to:

Theorem 50.7.2 (Weak Nullstellensatz)

Let $I \subsetneq \mathbb{C}[x_1, \dots, x_n]$ be a proper ideal. Then the variety $\mathcal{V}(I) \neq \emptyset$.

The proof of this is surprisingly pernicious, so we won’t include it here; see [Va15, §7.4.3].

From this we can deduce that all maximal ideals are of the above form.

Theorem 50.7.3 (Maximal ideals)

Every maximal ideal of $\mathbb{C}[x_1, \dots, x_n]$ is of the form $\mathcal{V}(x_1 - a_1, \dots, x_n - a_n)$.

WN implies MI. Let J be a maximal ideal, and consider the corresponding variety $V = \mathcal{V}(J)$. By WN, it contains some point $p = (a_1, \dots, a_n)$. Now, define $I = (x_1 - a_1, \dots, x_n - a_n)$; this ideal contains all polynomials vanishing at p , so necessarily $J \subseteq I \subsetneq \mathbb{C}[x_1, \dots, x_n]$. Then by maximality of J we have $J = I$. $\square$

Again this uses the fact that $\mathbb{C}$ is algebraically closed. Thus:

Over $\mathbb{C}$, maximal ideals correspond to single points.

Consequently, our various ideals over $\mathbb{C}$ correspond to various flavors of affine varieties:

Algebraic flavor	Geometric flavor
radical ideal	affine variety
prime ideal	irreducible variety
maximal ideal	single point
any ideal	scheme


There’s one thing I haven’t talked about: what’s the last entry?

§50.8 Why schemes?

One of the most elementary motivations for the scheme is that we would like to use them to count multiplicity. That is, consider the intersection

$$\mathcal{V}(y - x^2) \cap \mathcal{V}(y) \subseteq \mathbb{A}^2$$

This is the intersection of the parabola with the tangent x -axis, this is the green dot below.



Unfortunately, as a variety, it is just a single point! However, we want to think of this as a “double point”: after all, in the obvious sense it has multiplicity 2. You can detect this when you look at the ideals:

$$(y - x^2) + (y) = (x^2, y)$$

and thus, if we blithely ignore taking the radical, we get

$$\mathbb{C}[x, y]/(x^2, y) \cong \mathbb{C}[\varepsilon]/(\varepsilon^2).$$

So the ideals in question are noticing the presence of a double point.

In order to encapsulate this, we need a more refined object than a variety, which (at the end of the day) is really a bunch of points. This refined object is the *scheme*.

§50.9 Problems to think about

Problem 50A. Show that a *real* affine variety $V \subseteq \mathbb{A}_{\mathbb{R}}^n$ can always be written in the form $\mathcal{V}(f)$.



Problem 50B. Show that Hilbert’s Nullstellensatz in n dimensions follows from the Weak Nullstellensatz. (This solution is called the **Rabinowitsch Trick**.)

51 Affine varieties as ringed spaces


As in the previous chapter, we are working only over affine varieties in $\mathbb{C}$ for simplicity.

§51.1 Synopsis

Group theory was a strange creature in the early 19th century. During the 19th century, a group was literally defined as a subset of $\mathrm{GL}(n)$ or of S_n . Indeed, the word “group” hadn’t been invented yet. This may sound ludicrous, but it was true – Sylow developed his theorems without this notion. Only much later was the abstract definition of a group given, an abstract set G which was *independent* of any embedding into S_n , and an object in its own right.

We are about to make the same type of change for our affine varieties. Rather than thinking of them as an object locked into an ambient space $\mathbb{A}^n$ we are instead going to try to make them into an object in their own right. Specifically, for us an affine variety will become a *topological space* equipped with a *ring of functions* for each of its open sets: this is why we call it a **ringed space**.

The bit about the topological space is not too drastic. The key insight is the addition of the ring of functions. For example, consider the double point from last chapter.



As a set, it is a single point, and thus it can have only one possible topology. But the addition of the function ring will let us tell it apart from just a single point.

This construction is quite involved, so we’ll proceed as follows: we’ll define the structure bit by bit onto our existing affine varieties in $\mathbb{A}^n$, until we have all the data of a ringed space. In later chapters, these ideas will grow up to become the core of modern algebraic geometry: the *scheme*.

§51.2 The Zariski topology on $\mathbb{A}^n$

Prototypical example for this section: In $\mathbb{A}^1$, closed sets are finite collections of points. In $\mathbb{A}^2$, a nonempty open set is the whole space minus some finite collection of curves/points.

We begin by endowing a topological structure on every variety V . Since our affine varieties (for now) all live in $\mathbb{A}^n$, all we have to do is put a suitable topology on $\mathbb{A}^n$, and then just view V as a subspace.

However, rather than putting the standard Euclidean topology on $\mathbb{A}^n$, we put a much more bizarre topology.

Definition 51.2.1. In the **Zariski topology** on $\mathbb{A}^n$, the *closed sets* are those of the form

$$\mathcal{V}(I) \quad \text{where} \quad I \subseteq \mathbb{C}[x_1, \dots, x_n].$$

Of course, the open sets are complements of such sets.

Example 51.2.2 (Zariski topology on $\mathbb{A}^1$)

Let us determine the open sets of $\mathbb{A}^1$, which as usual we picture as a straight line (ignoring the fact that $\mathbb{C}$ is two-dimensional).

Since $\mathbb{C}[x]$ is a principal ideal domain, rather than looking at $\mathcal{V}(I)$ for every $I \subseteq \mathbb{C}[x]$, we just have to look at $\mathcal{V}(f)$ for a single f . There are a few flavors of polynomials f :

- The zero polynomial 0 which vanishes everywhere: this implies that the entire space $\mathbb{A}^1$ is a closed set.
- The constant polynomial 1 which vanishes nowhere. This implies that $\emptyset$ is a closed set.
- A polynomial $c(x - t_1)(x - t_2) \dots (x - t_n)$ of degree n . It has n roots, and so $\{t_1, \dots, t_n\}$ is a closed set.

Hence the closed sets of $\mathbb{A}^1$ are exactly all of $\mathbb{A}^1$ and finite sets of points (including $\emptyset$). Consequently, the *open* sets of $\mathbb{A}^1$ are

- $\emptyset$, and
- $\mathbb{A}^1$ minus a finite collection (possibly empty) of points.

Thus, the picture of a “typical” open set $\mathbb{A}^1$ might be



It’s everything except a few marked points!

Example 51.2.3 (Zariski topology on $\mathbb{A}^2$)

Similarly, in $\mathbb{A}^2$, the interesting closed sets are going to consist of finite unions (possibly empty) of

- Closed curves, like $\mathcal{V}(y - x^2)$ (which is a parabola), and
- Single points, like $\mathcal{V}(x - 3, y - 4)$ (which is the point $(3, 4)$).

Of course, the entire space $\mathbb{A}^2 = \mathcal{V}(0)$ and the empty set $\emptyset = \mathcal{V}(1)$ are closed sets.

Thus the nonempty open sets in $\mathbb{A}^2$ consist of the *entire* plane, minus a finite collection of points and one-dimensional curves.

Question 51.2.4. Draw a picture (to the best of your artistic ability) of a “typical” open set in $\mathbb{A}^2$.

All this is to say

The nonempty Zariski open sets are *huge*.

This is an important difference than what you're used to in topology. To be very clear:

- In the past, if I said something like “has so-and-so property in a neighborhood of point p ”, one thought of this as saying “is true in a small region around p ”.
- In the Zariski topology, “has so-and-so property in a neighborhood of point p ” should be thought of as saying “is true for virtually all points, other than those on certain curves”.

Indeed, “neighborhood” is no longer a very accurate description.

It remains to verify that as I've stated it, the closed sets actually form a topology. That is, I need to verify briefly that

- $\emptyset$ and $\mathbb{A}^n$ are both closed.
- Intersections of closed sets (even infinite) are still closed.
- Finite unions of closed sets are still closed.

Well, closed sets are the same as varieties, so we already know this!


§51.3 The Zariski topology on affine varieties

Prototypical example for this section: If $V = \mathcal{V}(y - x^2)$ is a parabola, then V minus $(1, 1)$ is open in V . Also, the plane minus the origin is $D(x) \cup D(y)$.

As we said before, by considering a variety V as a subspace of $\mathbb{A}^n$ it inherits the Zariski topology. One should think of an open subset of V as “ V minus a few Zariski-closed sets”. For example:

Example 51.3.1 (Open set of a variety)

Let $V = \mathcal{V}(y - x^2) \subseteq \mathbb{A}^2$ be a parabola, and let $U = V \setminus \{(1, 1)\}$. We claim U is open in V .



Indeed, $\tilde{U} = \mathbb{A}^2 \setminus \{(1, 1)\}$ is open in $\mathbb{A}^2$ (since it is the complement of the closed set $\mathcal{V}(x - 1, y - 1)$), so $U = \tilde{U} \cap V$ is open in V . Note that on the other hand the set U is *not* open in $\mathbb{A}^2$.

We'll finish by giving a nice basis of the Zariski topology.

Definition 51.3.2. Given $V \subseteq \mathbb{A}^n$ an affine variety and $f \in \mathbb{C}[x_1, \dots, x_n]$, we define the **distinguished open set** $D(f)$ to be the open set in V of points not vanishing on f :

$$D(f) = \{p \in V \mid f(p) \neq 0\} = V \setminus \mathcal{V}(f).$$

[Va15] suggests remembering the notation $D(f)$ as “doesn’t-vanish set”.

Example 51.3.3 (Examples of (unions of) distinguished open sets)

- (a) If $V = \mathbb{A}^1$ then $D(x)$ corresponds to a line minus a point.
- (b) If $V = \mathcal{V}(y - x^2) \subseteq \mathbb{A}^2$, then $D(x - 1)$ corresponds to the parabola minus $(1, 1)$.
- (c) If $V = \mathbb{A}^2$, then $D(x) \cup D(y) = \mathbb{A}^2 \setminus \{(0, 0)\}$ is the punctured plane.

Question 51.3.4. Show that every open set of a $V \subseteq \mathbb{A}^n$ can be written as a finite union of distinguished open sets.

Question 51.3.5. Give an example of an open set of $\mathbb{A}^2$ which is not a distinguished open set. (There was one above already.)

§51.4 Coordinate rings

Prototypical example for this section: If $V = \mathcal{V}(y - x^2)$ then $\mathbb{C}[V] = \mathbb{C}[x, y]/(y - x^2)$.

The next thing we do is consider the functions from V to the base field $\mathbb{C}$. We restrict our attention to algebraic (polynomial) functions on a variety V : they should take every point $(a_1, \dots, a_n)$ on V to some complex number $P(a_1, \dots, a_n) \in \mathbb{C}$. For example, a valid function on a three-dimensional affine variety might be $(a, b, c) \mapsto a$; we just call this projection “ x ”. Similarly we have a canonical projection y and z , and we can create polynomials by combining them, say $x^2y + 2xyz$.

Definition 51.4.1. The **coordinate ring** $\mathbb{C}[V]$ of a variety V is the ring of polynomial functions on V . (Notation explained next section.)

At first glance, we might think this is just $\mathbb{C}[x_1, \dots, x_n]$. But on closer inspection we realize that *on a given variety*, some of these functions are the same. For example, consider in $\mathbb{A}^2$ the parabola $V = \mathcal{V}(y - x^2)$. Then the two functions

$$\begin{aligned} V &\rightarrow \mathbb{C} \\ (x, y) &\mapsto x^2 \\ (x, y) &\mapsto y \end{aligned}$$

are actually the same function! We have to “mod out” by the ideal I which generates V . This leads us naturally to:

Theorem 51.4.2 (Coordinate rings correspond to ideal)

Let I be a radical ideal, and $V = \mathcal{V}(I) \subseteq \mathbb{A}^n$. Then

$$\mathbb{C}[V] \cong \mathbb{C}[x_1, \dots, x_n]/I.$$

Proof. There's a natural surjection as above

$$\mathbb{C}[x_1, \dots, x_n] \twoheadrightarrow \mathbb{C}[V]$$

and the kernel is I . □

Thus properties of a variety V correspond to properties of the ring $\mathbb{C}[V]$.

§51.5 The sheaf of regular functions

Prototypical example for this section: Let $V = \mathbb{A}^1$, $U = V \setminus \{0\}$. Then $1/x \in \mathcal{O}_V(U)$ is regular on U .

Let V be an affine variety and let $\mathbb{C}[V]$ be its coordinate ring. We want to define a notion of $\mathcal{O}_V(U)$ for any open set U : the “nice” functions on any open subset. Obviously, any function in $\mathbb{C}[V]$ will work as a function on $\mathcal{O}_V(U)$. However, to capture more of the structure we want to loosen our definition of “nice” function slightly by allowing *rational* functions.

The chief example is that $1/x$ should be a regular function on $\mathbb{A}^1 \setminus \{0\}$. The first natural guess is:

Definition 51.5.1. Let $U \subseteq V$ be an open set of the variety V . A **rational function** on U is a quotient $f(x)/g(x)$ of two elements f and g in $\mathbb{C}[V]$, where we require that $g(x) \neq 0$ for $x \in U$.

However, the definition is slightly too restrictive; we have to allow for multiple representations:

Definition 51.5.2. Let $U \subseteq V$ be open. We say a function $\phi : U \rightarrow \mathbb{C}$ is a **regular function** if for every point $p \in U$, we can find an open set $U_p \subseteq U$ containing p and a rational function f_p/g_p on U_p such that

$$\phi(x) = \frac{f_p(x)}{g_p(x)} \quad \forall x \in U_p.$$

In particular, we require $g_p(x) \neq 0$ on the set U_p . We denote the set of all regular functions on U by $\mathcal{O}_V(U)$.

Thus,

ϕ is regular on U if it is locally a rational function.

This definition is misleadingly complicated, and the examples should illuminate it significantly. Firstly, in practice, most of the time we will be able to find a “global” representation of a regular function as a quotient, and we will not need to fuss with the p ’s. For example:

Example 51.5.3 (Regular functions)

- (a) Any function in $f \in \mathbb{C}[V]$ is clearly regular, since we can take $g_p = 1$, $f_p = f$ for every p . So $\mathbb{C}[V] \subseteq \mathcal{O}_V(U)$ for any open set U .
- (b) Let $V = \mathbb{A}^1$, $U_0 = V \setminus \{0\}$. Then $1/x \in \mathcal{O}_V(U_0)$ is regular on U .
- (c) Let $V = \mathbb{A}^1$, $U_{12} = V \setminus \{1, 2\}$. Then

$$\frac{1}{(x-1)(x-2)} \in \mathcal{O}_V(U_{12})$$

is regular on U .

The “local” clause with p ’s is still necessary, though.

Example 51.5.4 (Requiring local representations)

Consider the variety

$$V = \mathcal{V}(ab - cd) \subseteq \mathbb{A}^4$$

and the open set $U = V \setminus \mathcal{V}(b, d)$. There is a regular function on U given by

$$(a, b, c, d) \mapsto \begin{cases} a/d & d \neq 0 \\ c/b & b \neq 0. \end{cases}$$

Clearly these are the “same function” (since $ab = cd$), but we cannot write “ a/d ” or “ c/b ” to express it because we run into divide-by-zero issues. That’s why in the definition of a regular function, we have to allow multiple representations.

In fact, we will see later on that the definition of a regular function is a special case of a more general construction called *sheafification*, in which “sheaves of functions which are P ” are transformed into “sheaves of functions which are *locally* P ”.

§51.6 Regular functions on distinguished open sets

Prototypical example for this section: Regular functions on $\mathbb{A}^1 \setminus \{0\}$ are $P(x)/x^n$.

The division-by-zero, as one would expect, essentially prohibits regular functions on the entire space V ; i.e. there are no regular functions in $\mathcal{O}_V(V)$ that were not already in $\mathbb{C}[V]$. Actually, we have a more general result which computes the regular functions on distinguished open sets:

Theorem 51.6.1 (Regular functions on distinguished open sets)

Let $V \subseteq \mathbb{A}^n$ be an affine variety and $D(g)$ a distinguished open subset of it. Then

$$\mathcal{O}_V(D(g)) = \left\{ \frac{f}{g^n} \mid f \in \mathbb{C}[V] \text{ and } n \in \mathbb{Z} \right\}.$$

In particular, $\mathcal{O}_V(V) = \mathcal{O}_V(D(1)) \cong \mathbb{C}[V]$.

The proof of this theorem requires the Nullstellensatz, so it relies on $\mathbb{C}$ being algebraically closed. In fact, a counter-example is easy to find if we replace $\mathbb{C}$ by $\mathbb{R}$: consider $\frac{1}{x^2+1}$.

Proof. Obviously, every function of the form f/g^n works, so we want the reverse direction. This is long, and perhaps should be omitted on a first reading.

Here's the situation. Let $U = D(g)$. We're given a regular function ϕ , meaning at every point $p \in D(g)$, there is a neighborhood U_p on which ϕ can be expressed as f_p/g_p (where $f_p, g_p \in \mathbb{C}[V]$). Then, we want to construct an $f \in \mathbb{C}[V]$ and an integer n such that $\phi = f/g^n$.

First, look at a particular U_p and f_p/g_p . Shrink U_p to a distinguished open set $D(h_p)$. Then, let $\tilde{f}_p = f_p h_p$ and $\tilde{g}_p = g_p h_p$. Thus we have that

$$\frac{\tilde{f}_p}{\tilde{g}_p} \text{ is correct on } D(h_p) \subseteq U \subseteq X.$$

The upshot of using the modified f_p and g_p is that:

$$\tilde{f}_p \tilde{g}_q = \tilde{f}_q \tilde{g}_p \quad \forall p, q \in U.$$

Indeed, it is correct on $D(h_p) \cap D(h_q)$ by definition, and outside this set both the left-hand side and right-hand side are zero.

Now, we know that $D(g) = \bigcup_{p \in U} D(\tilde{g}_p)$, i.e.

$$\mathcal{V}(g) = \bigcap_{p \in U} \mathcal{V}(\tilde{g}_p).$$

So by the Nullstellensatz we know that

$$g \in \sqrt{(\tilde{g}_p : p \in U)} \implies \exists n : g^n \in (\tilde{g}_p : p \in U).$$

In other words, for some n and $k_p \in \mathbb{C}[V]$ we have

$$g^n = \sum_p k_p \tilde{g}_p$$

where only finitely many k_p are not zero. Now, we claim that

$$f \stackrel{\text{def}}{=} \sum_p k_p \tilde{f}_p$$

works. This just observes by noting that for any $q \in U$, we have

$$f \tilde{g}_q - g^n \tilde{f}_q = \sum_p k_p (\tilde{f}_p \tilde{g}_q - \tilde{g}_p \tilde{f}_q) = 0. \quad \square$$

This means that the *global* regular functions are just the same as those in the coordinate ring: you don't gain anything new by allowing it to be locally a quotient. (The same goes for distinguished open sets.)

Example 51.6.2 (Regular functions on distinguished open sets)

(a) As said already, taking $g = 1$ we recover $\mathcal{O}_V(V) \cong \mathbb{C}[V]$ for any affine variety V .

(b) Let $V = \mathbb{A}^1$, $U_0 = V \setminus \{0\}$. Then

$$\mathcal{O}_V(U_0) = \left\{ \frac{P(x)}{x^n} \mid P \in \mathbb{C}[x], \quad n \in \mathbb{Z} \right\}.$$

So more examples are $1/x$ and $(x+1)/x^3$.

Question 51.6.3. Why doesn't our theorem on regular functions apply to Example 51.5.4?

The regular functions will become of crucial importance once we define a scheme in the next chapter.

§51.7 Baby ringed spaces

In summary, given an affine variety V we have:

- A structure of a set of points,
- A structure of a topological space V on these points, and
- For every open set $U \subseteq V$, a ring $\mathcal{O}_V(U)$. Elements of the rings are functions $U \rightarrow \mathbb{C}$.

Let us agree that:

Definition 51.7.1. A **baby ringed space** is a topological space X equipped with a ring $\mathcal{O}_X(U)$ for every open set U . It is required that elements of the ring $\mathcal{O}_X(U)$ are functions $f : U \rightarrow \mathbb{C}$; we call these the *regular functions* of X on U .

Therefore, affine varieties are baby ringed spaces.

Remark 51.7.2. This is not a standard definition. Hehe.

The reason this is called a “baby ringed space” is that in a *ringed space*, the rings $\mathcal{O}_V(U)$ can actually be *any rings*, but they have to satisfy a set of fairly technical conditions. When this happens, it's the $\mathcal{O}_V$ that does all the work; we think of $\mathcal{O}_V$ as a type of functor called a *sheaf*.

Since we are only studying affine/projective/quasi-projective varieties for the next chapters, we will just refer to these as baby ringed spaces so that we don't have to deal with the entire definition. The key concept is that we want to think of these varieties as *intrinsic objects*, free of any embedding. A baby ringed space is philosophically the correct thing to do.

Anyways, affine varieties are baby ringed spaces $(V, \mathcal{O}_V)$. In the next chapter we'll meet projective and quasi-projective varieties, which give more such examples of (baby) ringed spaces. With these examples in mind, we will finally lay down the complete definition of a ringed space, and use this to define a scheme.

§51.8 Problems to think about

Problem 51A[†]. Show that for any $n \geq 1$ the Zariski topology of $\mathbb{A}^n$ is *not* Hausdorff.

Problem 51B[†]. Let V be an affine variety, and consider its Zariski topology.

- Show that the Zariski topology is **Noetherian**, meaning there is no infinite descending chain $Z_1 \supsetneq Z_2 \supsetneq Z_3 \supsetneq \dots$ of closed subsets.
- Prove that a Noetherian topological space is compact. Hence varieties are topologically compact.

Problem 51C^{*} (Punctured Plane). Let $V = \mathbb{A}^2$ and let $X = \mathbb{A}^2 \setminus \{(0,0)\}$ be the punctured plane (which is an open set of V). Compute $\mathcal{O}_V(X)$.

52 Projective varieties

Having studied affine varieties in $\mathbb{A}^n$, we now consider $\mathbb{CP}^n$. We will also make it into a baby ringed space in the same way as with $\mathbb{A}^n$.

§52.1 Graded rings

Prototypical example for this section: $\mathbb{C}[x_0, \dots, x_n]$ is a graded ring.

We first take the time to state what a graded ring is, just so that we have this language to use (now and later).

Definition 52.1.1. A **graded ring** R is a ring with the following additional structure: as an abelian group, it decomposes as

$$R = \bigoplus_{d \geq 0} R^d$$

where $R^0, R^1, \dots$, are abelian groups. The ring multiplication has the property that if $r \in R^d$ and $s \in R^e$, we have $rs \in R^{d+e}$. Elements of an R^d are called **homogeneous elements**; we write “ $d = \deg r$ ” to mean “ $r \in R^d$ ”.

We denote by R^+ the ideal $R \setminus R_0$ generated by the homogeneous elements of nonzero degree, and call it the **irrelevant ideal**.

Remark 52.1.2. For experts: all our graded rings are commutative with 1.

Example 52.1.3 (Examples of graded rings)

- (a) The ring $\mathbb{C}[x]$ is graded by degree: as abelian groups, $\mathbb{C}[x] \cong \mathbb{C} \oplus x\mathbb{C} \oplus x^2\mathbb{C} \oplus \dots$
- (b) More generally, the polynomial ring $\mathbb{C}[x_0, \dots, x_n]$ is graded by degree.

Abuse of Notation 52.1.4. The notation $\deg r$ is abusive in the case $r = 0$; note that $0 \in R^d$ for every d . So it makes sense to talk about “the” degree of r except when $r = 0$.

We will frequently refer to homogeneous ideals:

Definition 52.1.5. An ideal $I \subseteq \mathbb{C}[x_0, \dots, x_n]$ is **homogeneous** if it can be written as $I = (f_1, \dots, f_m)$ where each f_i is a homogeneous polynomial.

Remark 52.1.6. If I and J are homogeneous, then so are $I + J$, IJ , $I \cap J$, $\sqrt{I}$.

Lemma 52.1.7 (Graded quotients are graded too)

Let I be a homogeneous ideal of a graded ring R . Then

$$R/I = \bigoplus_{d \geq 0} R^d / (R^d \cap I)$$

realizes R/I as a graded ring.

Since these assertions are just algebra, we omit their proofs here.

Example 52.1.8 (Example of a graded quotient ring)

Let $R = \mathbb{C}[x, y]$ and set $I = (x^3, y^2)$. Then

$$\begin{aligned} R^0 &= \mathbb{C} \\ R^1 &= \mathbb{C}x \oplus \mathbb{C}y \\ R^2 &= \mathbb{C}x^2 \oplus \mathbb{C}xy \\ R^3 &= \mathbb{C}x^2y \\ R^d &= 0 \quad \forall d \geq 4. \end{aligned}$$

So in fact R/I is graded, and is a six-dimensional $\mathbb{C}$ -vector space.

§52.2 The ambient space

Prototypical example for this section: Perhaps $\mathcal{V}_{pr}(x^2 + y^2 - z^2)$.

The set of points we choose to work with is $\mathbb{CP}^n$ this time, which for us can be thought of as the set of n -tuples

$$(x_0 : x_1 : \cdots : x_n)$$

not all zero, up to scaling. Equivalently, it is the set of lines through the origin in $\mathbb{C}^{n+1}$. Projective space is defined in full in Section 20.6, and you should refer there if you aren't familiar with projective space.

The right way to think about it is “ $\mathbb{A}^n$ plus points at infinity”:

Definition 52.2.1. We define the set

$$U_i = \{(x_0 : \cdots : x_n) \mid x_i \neq 0\} \subseteq \mathbb{CP}^n.$$

These are called the **standard affine charts**.

The name comes from:

Exercise 52.2.2. (Mandatory) Give a natural bijection from U_i to $\mathbb{A}^n$. Thus we can think of $\mathbb{CP}^n$ as the affine set U_i plus “points at infinity”.

Remark 52.2.3. In fact, these charts U_i make $\mathbb{CP}^n$ with its usual topology into a complex manifold with holomorphic transition functions.

However, like before we want to impose a Zariski topology on it. For concreteness, let's consider $\mathbb{CP}^2 = \{(x_0 : x_1 : x_2)\}$. We wish to consider zero loci in $\mathbb{CP}^2$, just like we did in affine space, and hence obtain a notion of a projective variety.

But this isn't so easy: for example, the function “ x_0 ” is not a well-defined function on points in $\mathbb{CP}^2$ because $(x_0 : x_1 : x_2) = (5x_0 : 5x_1 : 5x_2)$! So we'd love to consider these “pseudo-functions” that still have zero loci. These are just the homogeneous polynomials f , because f is homogeneous of degree d if and only if

$$f(\lambda x_0, \dots, \lambda x_n) = \lambda^d f(x_0, \dots, x_n).$$

In particular, the relation “ $f(x_0, \dots, x_n) = 0$ ” is well-defined if F is homogeneous. Thus, we can say:

Definition 52.2.4. If f is homogeneous, we can then define its **vanishing locus** as

$$\mathcal{V}_{\text{pr}}(f) = \{(x_0 : \cdots : x_n) \mid f(x_0, \dots, x_n) = 0\}.$$

The homogeneous condition is really necessary. For example, to require “ $x_0 - 1 = 0$ ” makes no sense, since the points $(1 : 1 : 1)$ and $(2015 : 2015 : 2015)$ are the same.

It’s trivial to verify that homogeneous polynomials do exactly what we want; hence we can now define:


Definition 52.2.5. A **projective variety** in $\mathbb{CP}^n$ is the common zero locus of an arbitrary collection of homogeneous polynomials in $n + 1$ variables.

Example 52.2.6 (A conic in $\mathbb{CP}^2$, or a cone in $\mathbb{C}^3$)

Let’s try to picture the variety

$$\mathcal{V}_{\text{pr}}(x^2 + y^2 - z^2) \subseteq \mathbb{CP}^2$$

which consists of the points $[x : y : z]$ such that $x^2 + y^2 = z^2$. If we view this as subspace of $\mathbb{C}^3$ (i.e. by thinking of $\mathbb{CP}^2$ as the set of lines through the origin), then we get a “cone”:



If we take the standard affine charts now, we obtain:

- At $x = 1$, we get a hyperbola $\mathcal{V}(1 + y^2 - z^2)$.
- At $y = 1$, we get a hyperbola $\mathcal{V}(1 + x^2 - z^2)$.
- At $z = 1$, we get a circle $\mathcal{V}(x^2 + y^2 - 1)$.

That said, over $\mathbb{C}$ a hyperbola and circle are the same thing; I’m cheating a little by drawing $\mathbb{C}$ as one-dimensional, just like last chapter.

Question 52.2.7. Draw the intersection of the cone above with the $z = 1$ plane, and check that you do in fact get a circle. (This geometric picture will be crucial later.)

§52.3 Homogeneous ideals

Now, the next thing we want to do is define $\mathcal{V}_{\text{pr}}(I)$ for an ideal I . Of course, we again run into an issue with things like $x_0 - 1$ not making sense.

The way out of this is to use only *homogeneous* ideals.

Definition 52.3.1. If I is a homogeneous ideal, we define

$$\mathcal{V}_{\text{pr}}(I) = \{x \mid f(x) = 0 \forall f \in I\}.$$

Exercise 52.3.2. Show that the notion “ $f(x) = 0 \forall f \in I$ ” is well-defined for a homogeneous ideal I .

So, we would hope for a Nullstellensatz-like theorem which bijects the homogeneous radical ideals to projective ideals. Unfortunately:

Example 52.3.3 (Irrelevant ideal)

To crush some dreams and hopes, consider the ideal

$$I = (x_0, x_1, \dots, x_n).$$

This is called the **irrelevant ideal**; it is a homogeneous radical yet $\mathcal{V}_{\text{pr}}(I) = \emptyset$.

However, other than the irrelevant ideal:

Theorem 52.3.4 (Homogeneous Nullstellensatz)

Let I and J be homogeneous ideals.

- (a) If $\mathcal{V}_{\text{pr}}(I) = \mathcal{V}_{\text{pr}}(J) \neq \emptyset$ then $\sqrt{I} = \sqrt{J}$.
- (b) If $\mathcal{V}_{\text{pr}}(I) = \emptyset$, then either $I = (1)$ or $\sqrt{I} = (x_0, x_1, \dots, x_n)$.

Thus there is a natural bijection between:

- projective varieties in $\mathbb{CP}^n$, and
- homogeneous radical ideals of $\mathbb{C}[x_1, \dots, x_n]$ except for the irrelevant ideal.

Proof. For the first part, let $V = \mathcal{V}_{\text{pr}}(I)$ and $W = \mathcal{V}_{\text{pr}}(J)$ be projective varieties in $\mathbb{CP}^n$. We can consider them as *affine varieties* in $\mathbb{A}^{n+1}$ by using the interpretation of $\mathbb{CP}^n$ as lines through the origin in $\mathbb{C}^n$.

Algebraically, this is done by taking the homogeneous ideals $I, J \subseteq \mathbb{C}[x_0, \dots, x_n]$ and using the same ideals to cut out *affine varieties* $V_{\text{aff}} = \mathcal{V}(I)$ and $W_{\text{aff}} = \mathcal{V}(J)$ in $\mathbb{A}^{n+1}$. For example, the cone $x^2 + y^2 - z^2 = 0$ is a conic (a one-dimensional curve) in $\mathbb{CP}^2$, but can also be thought of as a cone (which is a two-dimensional surface) in $\mathbb{A}^3$.

Then for (a), we have $V_{\text{aff}} = W_{\text{aff}}$, so $\sqrt{I} = \sqrt{J}$.

For (b), either V_{aff} is empty or it is just the origin of $\mathbb{A}^{n+1}$, so the Nullstellensatz implies either $I = (1)$ or $\sqrt{I} = (x_0, \dots, x_n)$ as desired. $\square$

Projective analogues of Theorem 50.4.2 (on intersections and unions of varieties) hold verbatim for projective varieties as well.

§52.4 As ringed spaces

Prototypical example for this section: The regular functions on $\mathbb{CP}^1$ minus a point are exactly those of the form $P(s/t)$.

Now, let us make every projective variety V into a baby ringed space. We already have the set of points, a subset of $\mathbb{CP}^n$.

The topology is defined as follows.

Definition 52.4.1. We endow $\mathbb{CP}^n$ with the **Zariski topology** by declaring the sets of the form $\mathcal{V}_{\text{pr}}(I)$, where I is a homogeneous ideal, to be the closed sets.

Every projective variety V then inherits the Zariski topology from its parent $\mathbb{CP}^n$. The **distinguished open sets** $D(f)$ are $V \setminus \mathcal{V}_{\text{pr}}(f)$.

Thus every projective variety V is now a topological space. It remains to endow it with a sheaf of regular functions $\mathcal{O}_V$. To do this we have to be a little careful. In the affine case we had a nice little ring of functions, the coordinate ring $\mathbb{C}[x_0, \dots, x_n]/I$, that we could use to provide the numerator and denominators. So, it seems natural to then define:

Definition 52.4.2. The **homogeneous coordinate ring** of a projective variety $V = \mathcal{V}_{\text{pr}}(I) \subseteq \mathbb{CP}^n$, where I is homogeneous radical, is defined as the ring

$$\mathbb{C}[V] = \mathbb{C}[x_0, \dots, x_n]/I.$$

However, when we define a rational function we must impose a new requirement that the numerator and denominator are the same degree.

Definition 52.4.3. Let $U \subseteq V$ be an open set of a projective variety V . A **rational function** ϕ on a projective variety V is a quotient f/g , where $f, g \in \mathbb{C}[V]$, and f and g are homogeneous of the same degree, and $\mathcal{V}_{\text{pr}}(g) \cap U = \emptyset$. In this way we obtain a function $\phi : U \rightarrow \mathbb{C}$.

Example 52.4.4 (Examples of rational functions)

Let $V = \mathbb{CP}^1$ have coordinates $(s : t)$.

- (a) If $U = V$, then constant functions $c/1$ are the only rational functions on U .
- (b) Now let $U_1 = V \setminus \{(1 : 0)\}$. Then, an example of a regular function is

$$\frac{s^2 + 9t^2}{t^2} = \left(\frac{s}{t}\right)^2 + 9.$$

If we think of U_1 as $\mathbb{C}$ (i.e. $\mathbb{CP}^1$ minus an infinity point, hence like $\mathbb{A}^1$) then really this is just the function $x^2 + 9$.

Then we can repeat the same definition as before:

Definition 52.4.5. Let $U \subseteq V$ be an open set of a projective variety V . We say a function $\phi : U \rightarrow \mathbb{C}$ is a **regular function** if for every point p , we can find an open set U_p containing p and a rational function f_p/g_p on U_p such that

$$\phi(x) = \frac{f_p(x)}{g_p(x)} \quad \forall x \in U_p.$$

In particular, we require $U_p \cap \mathcal{V}_{\text{pr}}(g_p) = \emptyset$. We denote the set of all regular functions on U by $\mathcal{O}_V(U)$.

Of course, the rational functions from the previous example are examples of regular functions as well. This completes the definition of a projective variety V as a baby ringed space.

§52.5 Examples of regular functions

Naturally, I ought to tell you what the regular functions on distinguished open sets are; this is an analog to Theorem 51.6.1 from last time.

Theorem 52.5.1 (Regular functions on distinguished open sets for projective varieties)

Let V be a projective variety, and let $g \in \mathbb{C}[V]$ be homogeneous of *positive degree* (thus g is nonconstant). Then

$$\mathcal{O}_V(D(g)) = \left\{ \frac{f}{g^r} \mid f \in \mathbb{C}[V] \text{ homogeneous of degree } r \deg g \right\}.$$

What about the case $g = 1$? A similar result holds, but we need the assumption that V is irreducible. (This has the same definition as the affine case, since “irreducible” is an adjective of topological spaces.)

Theorem 52.5.2 (Only constant regular functions on projective space)

Let V be an *irreducible* projective variety. Then the only regular functions on V are constant, thus we have

$$\mathcal{O}_V(V) \cong \mathbb{C}.$$

This relies on the fact that $\mathbb{C}$ is algebraically closed.

Proofs of these are omitted for now.

The reason we need V irreducible is otherwise we could, for example, take V to be the union of two points; in this case $\mathcal{O}_V(V) \cong \mathbb{C}^{\oplus 2}$.

It might seem strange that $\mathcal{O}_V(D(g))$ behaves so differently when $g = 1$. One vague explanation is that in a projective variety, a distinguished open $D(g)$ looks much like an affine variety if $\deg g > 0$. For example, in $\mathbb{CP}^1$ we have $\mathbb{CP}^1 \setminus \{0\} \cong \mathbb{A}^1$ (where $\cong$ is used in a sense that I haven’t made precise). Thus the claim becomes related to the corresponding affine result. But if $\deg g = 0$ and $g \neq 0$, then $D(g)$ is the entire projective variety, which does not look affine, and thus the analogy breaks down.

Example 52.5.3 (Regular functions on $\mathbb{CP}^1$)

Let $V = \mathbb{CP}^1$, with coordinates $(s : t)$.

- (a) By Theorem 52.5.1, if U_1 is the standard affine chart omitting the point $(1 : 0)$, we have $\mathcal{O}_V(U_1) = \left\{ \frac{f}{t^n} \mid \deg f = n \right\}$. One can write this as

$$\mathcal{O}_V(U_1) \cong \{P(s/t) \mid P \in \mathbb{C}[x]\} \cong \mathcal{O}_{\mathbb{A}^1}(\mathbb{A}^1).$$

This conforms with our knowledge that U_1 “looks very much like $\mathbb{A}^1$ ”.

- (b) As V is irreducible, $\mathcal{O}_V(V) = \mathbb{C}$: there are no nonconstant functions on $\mathbb{CP}^1$.

Example 52.5.4 (Regular functions on $\mathbb{CP}^2$)

Let $\mathbb{CP}^2$ have coordinates $(x : y : z)$ and let $U_0 = \{(x : y : 1) \in \mathbb{CP}^2\}$ be the distinguished open set $D(z)$. Then in the same vein,

$$\mathcal{O}_{\mathbb{CP}^2}(U_0) = \left\{ \frac{P(x, y)}{z^n} \mid \deg P = n \right\} \cong \{P(x/z, y/z) \mid P \in \mathbb{C}[x, y]\}.$$

§52.6 Problems to think about

DRAFT (Evan Chen)
Updated February 15, 2017

53 Bonus: Bézout's theorem

In this chapter we discuss Bézout's theorem. It makes precise the idea that two degree d and e curves in $\mathbb{CP}^2$ should intersect at “exactly” de points. (We work in projective space so e.g. any two lines intersect.)

§53.1 Non-radical ideals


Prototypical example for this section: Tangent to the parabola.

We need to account for multiplicities. So we will whenever possible work with homogeneous ideals I , rather than varieties V , because we want to allow the possibility that I is not radical. Let's see how we might do so.


For a first example, suppose we intersect $y = x^2$ with the line $y = 1$; or more accurately, in projective coordinates of $\mathbb{CP}^2$, the parabola $zy = x^2$ and $y = z$. The intersection of the ideals is

$$(zy - x^2, y - z) = (x^2 - z^2, y - z) \subseteq \mathbb{C}[x, y, z].$$

So this corresponds to having two points; this gives two intersection points: $(1 : 1 : 1)$ and $(-1 : 1 : 1)$. Here is a picture of the two varieties in the affine $z = 1$ chart:



That's fine, but now suppose we intersect $zy = x^2$ with the line $x = 0$ instead. Then we instead get a “double point”:



The corresponding ideal is this time

$$(zy - x^2, y) = (x^2, y) \subseteq \mathbb{C}[x, y, z].$$

This ideal is *not* radical, and when we take $\sqrt{(x^2, y)} = (x, y)$ we get the ideal which corresponds to a single projective point $(0 : 0 : 1)$ of $\mathbb{CP}^2$. This is why we work with ideals rather than varieties: we need to tell the difference between (x^2, y) and (x, y) .

§53.2 Hilbert functions of finitely many points

Prototypical example for this section: The Hilbert function attached to the double point (x^2, y) is eventually the constant 2.

Definition 53.2.1. Given a nonempty projective variety V , there is a unique radical ideal I such that $V = \mathcal{V}_{\text{pr}}(I)$. In this chapter we denote it by $\mathcal{I}_{\text{rad}}(V)$. For an empty variety we set $\mathcal{I}_{\text{rad}}(\emptyset) = (1)$, rather than choosing the irrelevant ideal.

Definition 53.2.2. Let $I \subseteq \mathbb{C}[x_0, \dots, x_n]$ be homogeneous. We define the **Hilbert function** of I , denoted $h_I : \mathbb{Z}_{\geq 0} \rightarrow \mathbb{Z}_{\geq 0}$ by

$$h_I(d) = \dim_{\mathbb{C}} (\mathbb{C}[x_0, \dots, x_n]/I)^d$$

i.e. $h_I(d)$ is the dimension of the d th graded part of $\mathbb{C}[x_0, \dots, x_n]/I$.

Definition 53.2.3. If V is a projective variety, we set $h_V = h_{\mathcal{I}_{\text{rad}}(V)}$, where I is the radical ideal satisfying $V = \mathcal{V}_{\text{pr}}(I)$. If $V = \emptyset$, we choose $I = (1)$.

Example 53.2.4 (Examples of Hilbert functions in zero dimensions)

For concreteness, let us use $\mathbb{CP}^2$.

(a) If V is the single point $(0 : 0 : 1)$, with ideal $\mathcal{I}_{\text{rad}}(V) = (x, y)$, then

$$\mathbb{C}[x, y, z]/(x, y) \cong \mathbb{C}[z] \cong \mathbb{C} \oplus z\mathbb{C} \oplus z^2\mathbb{C} \oplus z^3\mathbb{C} \dots$$

which has dimension 1 in all degrees. Consequently, we have

$$h_I(d) \equiv 1.$$

(b) Now suppose we use the “double point” ideal $I = (x^2, y)$. This time, we have

$$\begin{aligned} \mathbb{C}[x, y, z]/(x^2, y) &\cong \mathbb{C}[z] \oplus x\mathbb{C}[z] \\ &\cong \mathbb{C} \oplus (x\mathbb{C} \oplus z\mathbb{C}) \oplus (xz\mathbb{C} \oplus z^2\mathbb{C}) \oplus (xz^2\mathbb{C} \oplus z^3\mathbb{C}) \oplus \dots \end{aligned}$$

From this we deduce that

$$h_I(d) = \begin{cases} 2 & d = 1, 2, 3, \dots \\ 1 & d = 0. \end{cases}$$

(c) Let's now take the variety $V = \{(1 : 1 : 1), (-1 : 1 : 1)\}$ consisting of two points, with $\mathcal{I}_{\text{rad}}(V) = (x^2 - z^2, y - z)$. Then

$$\begin{aligned} \mathbb{C}[x, y, z]/(x^2 - z^2, y - z) &\cong \mathbb{C}[x, z]/(x^2 - z^2) \\ &\cong \mathbb{C}[z] \oplus x\mathbb{C}[z]. \end{aligned}$$

So this example has the same Hilbert function as the previous one.

Abuse of Notation 53.2.5. I'm abusing the isomorphism symbol $\mathbb{C}[z] \cong \mathbb{C} \oplus z\mathbb{C} \oplus z^2\mathbb{C}$ and similarly in other examples. This is an isomorphism only on the level of $\mathbb{C}$ -vector spaces. However, in computing Hilbert functions of other examples I will continue using this abuse of notation.

Example 53.2.6 (Hilbert functions for empty varieties)

Suppose $I \subsetneq \mathbb{C}[x_0, \dots, x_n]$ is an ideal, possibly not radical but such that

$$\mathcal{V}_{\text{pr}}(I) = \emptyset$$

hence $\sqrt{I} = (x_0, \dots, x_n)$ is the irrelevant ideal. Thus there are integers d_i for $i = 0, \dots, n$ such that $x_i^{d_i} \in I$ for every i ; consequently, $h_I(d) = 0$ for any $d > d_0 + \dots + d_n$. We summarize this by saying that

$$h_I(d) = 0 \text{ for all } d \gg 0.$$

Here the notation $d \gg 0$ means “all sufficiently large d ”.

From these examples we see that if I is an ideal, then the Hilbert function appears to eventually be constant, with the desired constant equal to the size of $\mathcal{V}_{\text{pr}}(I)$, “with multiplicity” in the case that I is not radical.

Let’s prove this. Before proceeding we briefly remind the reader of short exact sequences: a sequence of maps of $0 \rightarrow V \hookrightarrow W \twoheadrightarrow X \rightarrow 0$ is one such that the $\text{im}(V \hookrightarrow W) = \ker(W \twoheadrightarrow X)$ (and of course the maps $V \hookrightarrow W$ and $W \twoheadrightarrow X$ are injective and surjective). If V, W, X are finite-dimensional vector spaces over $\mathbb{C}$ this implies that $\dim W = \dim V + \dim X$.

Proposition 53.2.7 (Hilbert functions of $I \cap J$ and $I + J$)

Let I and J be homogeneous ideals in $\mathbb{C}[x_0, \dots, x_n]$. Then

$$h_{I \cap J} + h_{I+J} = h_I + h_J.$$

Proof. Consider any $d \geq 0$. Let $S = \mathbb{C}[x_0, \dots, x_n]$ for brevity. Then

$$0 \longrightarrow [S/(I \cap J)]^d \hookrightarrow [S/I]^d \oplus [S/J]^d \twoheadrightarrow [S/(I + J)]^d \longrightarrow 0$$

$$f \longmapsto (f, f)$$

$$(f, g) \longmapsto f - g$$

is a short exact sequence of vector spaces. Therefore, for every $d \geq 0$ we have that

$$\dim [S/I]^d \oplus [S/J]^d = \dim [S/(I \cap J)]^d + \dim [S/(I + J)]^d$$

which gives the conclusion. $\square$

Example 53.2.8 (Hilbert function of two points in $\mathbb{CP}^1$)

In $\mathbb{CP}^1$ with coordinate ring $\mathbb{C}[s, t]$, consider $I = (s)$ the ideal corresponding to the point $(0 : 1)$ and $J = (t)$ the ideal corresponding to the point $(1 : 0)$. Then $I \cap J = (st)$ is the ideal corresponding to the disjoint union of these two points, while $I + J = (s, t)$ is the irrelevant ideal. Consequently $h_{I+J}(d) = 0$ for $d \gg 0$. Therefore, we get

$$h_{I \cap J}(d) = h_I(d) + h_J(d) \text{ for } d \gg 0$$

so the Hilbert function of a two-point projective variety is the constant 2 for $d \gg 0$.

This example illustrates the content of the main result:

Theorem 53.2.9 (Hilbert functions of zero-dimensional varieties)

Let V be a projective variety consisting of m points (where $m \geq 0$ is an integer). Then

$$h_V(d) = m \text{ for } d \gg 0.$$

Proof. We already did $m = 0$, so assume $m \geq 1$. Let $I = \mathcal{I}_{\text{rad}}(V)$ and for $k = 1, \dots, m$ let $I_k = \mathcal{I}_{\text{rad}}(k\text{th point of } V)$.

Exercise 53.2.10. Show that $h_{I_k}(d) = 1$ for every d . (Modify Example 53.2.4(a).)

Hence we can proceed by induction on $m \geq 2$, with the base case $m = 1$ already done above. For the inductive step, we use the projective analogues of Theorem 50.4.2. We know that $h_{I_1 \cap \dots \cap I_{m-1}}(d) = m - 1$ for $d \gg 0$ (this is the first $m - 1$ points; note that $I_1 \cap \dots \cap I_{m-1}$ is radical). To add in the m th point we note that

$$h_{I_1 \cap \dots \cap I_m}(d) = h_{I_1 \cap \dots \cap I_{m-1}}(d) + h_{I_m}(d) - h_J(d)$$

where $J = (I_1 \cap \dots \cap I_{m-1}) + I_m$. The ideal J may not be radical, but satisfies $\mathcal{V}_{\text{pr}}(J) = \emptyset$ by an earlier example, hence $h_J = 0$ for $d \gg 0$. This completes the proof. $\square$

In exactly the same way we can prove that:

Corollary 53.2.11 (h_I eventually constant when $\dim \mathcal{V}_{\text{pr}}(I) = 0$)

Let I be an ideal, not necessarily radical, such that $\mathcal{V}_{\text{pr}}(I)$ consists of finitely many points. Then the Hilbert h_I is eventually constant.

Proof. Induction on the number of points, $m \geq 1$. The base case $m = 1$ was essentially done in Example 53.2.4(b) and Exercise 53.2.10. The inductive step is literally the same as in the proof above, except no fuss about radical ideals. $\square$

§53.3 Hilbert polynomials

So far we have only talked about Hilbert functions of zero-dimensional varieties, and showed that they are eventually constant. Let's look at some more examples.

Example 53.3.1 (Hilbert function of $\mathbb{CP}^n$)

The Hilbert function of $\mathbb{CP}^n$ is

$$h_{\mathbb{CP}^n}(d) = \binom{d+n}{n} = \frac{1}{n!}(d+n)(d+n-1)\dots(d+1)$$

by a “balls and urns” argument. This is a polynomial of degree n .

Example 53.3.2 (Hilbert function of the parabola)

Consider the parabola $zy - x^2$ in $\mathbb{CP}^2$ with coordinates $\mathbb{C}[x, y, z]$. Then

$$\mathbb{C}[x, y, z]/(zy - x^2) \cong \mathbb{C}[y, z] \oplus x\mathbb{C}[y, z].$$

A combinatorial computation gives that

$$\begin{array}{ll} h_{(zy-x^2)}(0) = 1 & \text{Basis } 1 \\ h_{(zy-x^2)}(1) = 3 & \text{Basis } x, y, z \\ h_{(zy-x^2)}(2) = 5 & \text{Basis } xy, xz, y^2, yz, z^2. \end{array}$$

We thus in fact see that $h_{(zy-x^2)}(d) = 2d - 1$.

In fact, this behavior of “eventually polynomial” always works.

Theorem 53.3.3 (Hilbert polynomial)

Let $I \subseteq \mathbb{C}[x_0, \dots, x_n]$ be a homogeneous ideal, not necessarily radical. Then

- (a) There exists a polynomial χ_I such that $h_I(d) = \chi_I(d)$ for all $d \gg 0$.
- (b) $\deg \chi_I = \dim \mathcal{V}_{\text{pr}}(I)$ (if $\mathcal{V}_{\text{pr}}(I) = \emptyset$ then $\chi_I = 0$).
- (c) The polynomial $m! \cdot \chi_I$ has integer coefficients.

Proof. The base case was addressed in the previous section.

For the inductive step, consider $\mathcal{V}_{\text{pr}}(I)$ with dimension m . Consider a hyperplane H such that no irreducible component of $\mathcal{V}_{\text{pr}}(I)$ is contained inside H (we quote this fact without proof, as it is geometrically obvious, but the last time I tried to write the proof I messed up). For simplicity, assume WLOG that $H = \mathcal{V}_{\text{pr}}(x_0)$.

Let $S = \mathbb{C}[x_0, \dots, x_n]$ again. Now, consider the short exact sequence

$$0 \longrightarrow [S/I]^{d-1} \xrightarrow{\times x_0} [S/I]^d \longrightarrow [S/(I + (x_0))]^d \longrightarrow 0$$

$$f \longmapsto fx_0$$

$$f \longmapsto f.$$

(The injectivity of the first map follows from the assumption about irreducible components of $\mathcal{V}_{\text{pr}}(I)$.) Now exactness implies that

$$h_I(d) - h_I(d-1) = h_{I+(x_0)}(d).$$

The last term geometrically corresponds to $\mathcal{V}_{\text{pr}}(I) \cap H$; it has dimension $m-1$, so by the inductive hypothesis we know that

$$h_I(d) - h_I(d-1) = \frac{c_0 d^{m-1} + c_1 d^{m-2} + \dots + c_{m-1}}{(m-1)!} \quad d \gg 0$$

for some integers $c_0, \dots, c_{m-1}$. Then we are done by the theory of **finite differences** of polynomials. $\square$

§53.4 Bézout's theorem

Definition 53.4.1. We call χ_I the **Hilbert polynomial** of I . If χ_I is nonzero, we call the leading coefficient of $m!\chi_I$ the **degree** of I , which is an integer, denoted $\deg I$.

Of course for projective varieties V we let $h_V = h_{\mathcal{I}_{\text{rad}}(V)}$.

Example 53.4.2 (Examples of degrees)

- (a) If V is a finite set of $n \geq 1$ points, it has degree n .
- (b) If I corresponds to a double point, it has degree 2.
- (c) $\mathbb{CP}^n$ has degree 1.
- (d) The parabola has degree 2.

Now, you might guess that if f is a homogeneous quadratic polynomial then the degree of the principal ideal (f) is 2, and so on. (Thus for example we expect a circle to have degree 2.) This is true:

Theorem 53.4.3 (Bézout's theorem)

Let I be a homogeneous ideal of $\mathbb{C}[x_0, \dots, x_n]$, such that $\dim \mathcal{V}_{\text{pr}}(I) \geq 1$. Let $f \in \mathbb{C}[x_0, \dots, x_n]$ be a homogeneous polynomial of degree k which does not vanish on any irreducible component of $\mathcal{V}_{\text{pr}}(I)$. Then

$$\deg(I + (f)) = k \deg I.$$

Proof. Let $S = \mathbb{C}[x_0, \dots, x_n]$ again. This time the exact sequence is

$$0 \longrightarrow [S/I]^{d-k} \xrightarrow{\times f} [S/I]^d \longrightarrow [S/(I + (f))]^d \longrightarrow 0$$

We leave this olympiad-esque exercise as Problem 53A. □

§53.5 Applications

First, we show that the notion of degree is what we expect.

Corollary 53.5.1 (Hypersurfaces: the degree deserves its name)

Let V be a hypersurface, i.e. $\mathcal{I}_{\text{rad}}(V) = (f)$ for f a homogeneous polynomial of degree k . Then $\deg V = k$.

Proof. Recall $\deg(0) = \deg \mathbb{CP}^n = 1$. Take $I = (0)$ in Bézout's theorem. □

The common special case in $\mathbb{CP}^2$ is:

Corollary 53.5.2 (Bézout's theorem for curves)

For any two curves X and Y in $\mathbb{CP}^2$ without a common irreducible component,


$$|X \cap Y| \leq \deg X \cdot \deg Y.$$

Now, we use this to prove Pascal's theorem.

Theorem 53.5.3 (Pascal's theorem)

Let A, B, C, D, E, F be six distinct points which lie on a conic $\mathcal{C}$ in $\mathbb{CP}^2$. Then the points $AB \cap DE$, $BC \cap EF$, $CD \cap FA$ are collinear.

Proof. Let X be the variety equal to the union of the three lines AB, CD, EF , hence $X = \mathcal{V}_{\text{pr}}(f)$ for some cubic polynomial f (which is the product of three linear ones). Similarly, let $Y = \mathcal{V}_{\text{pr}}(g)$ be the variety equal to the union of the three lines BC, DE, FA .



Now let P be an arbitrary point on the conic $\mathcal{C}$, distinct from the six points A, B, C, D, E, F . Consider the projective variety

$$V = \mathcal{V}_{\text{pr}}(\alpha f + \beta g)$$

where the constants α and β are chosen such that $P \in V$.

Question 53.5.4. Show that V also contains the six points A, B, C, D, E, F as well as the three points $AB \cap DE$, $BC \cap EF$, $CD \cap FA$ regardless of which α and β are chosen.

Now, note that $|V \cap \mathcal{C}| \geq 7$. But $\dim V = 3$ and $\dim \mathcal{C} = 2$. This contradicts Bézout's theorem unless V and $\mathcal{C}$ share an irreducible component. This can only happen if V is the union of a line and conic, for degree reasons; i.e. we must have that

$$V = \mathcal{C} \cup \text{line}.$$

Finally note that the three intersection points $AB \cap DE$, $BC \cap EF$ and $CD \cap FA$ do not lie on $\mathcal{C}$, so they must lie on this line. $\square$

§53.6 Problems to think about

Problem 53A. Complete the proof of Bézout's theorem from before.



Problem 53B (January TST for IMO 2016). Let ABC be an acute scalene triangle and let P be a point in its interior. Let A_1, B_1, C_1 be projections of P onto triangle sides BC, CA, AB , respectively. Find the locus of points P such that AA_1, BB_1, CC_1 are concurrent and $\angle PAB + \angle PBC + \angle PCA = 90^\circ$.

XIV

Algebraic Geometry II: Schemes

54	Morphisms of varieties	490
54.1	Defining morphisms of baby ringed spaces	490
54.2	Examples	491
54.3	Quasi-projective varieties	493
54.4	Some applications	494
54.5	Problems to think about	495
55	Sheaves and ringed spaces	496
55.1	Pre-sheaves	496
55.2	Sheaves	497
55.3	Stalks	498
55.4	Sections “are” sequences of germs	501
55.5	Sheafification	502
55.6	Morphisms of sheaves	503
55.7	Local rings, and locally ringed spaces	504
55.8	Morphisms of (locally) ringed spaces	505
55.9	Problems to think about	507
56	Schemes	508
56.1	The set of points	508
56.2	The Zariski topology of the spectrum	509
56.3	The structure sheaf	510
56.4	Example: fat points	512
56.5	Properties of affine schemes	513
56.6	Schemes	515
56.7	Projective scheme	515
56.8	Where to go from here	517
56.9	Problems to think about	517

54 Morphisms of varieties

Now that we've taken all the time to define affine and projective varieties, we want to take the time to define the *morphisms* between these objects. Fortunately, we know both affine and projective varieties are special cases of baby ringed spaces, so in fact we will just define a morphism between *any* two baby ringed spaces.

§54.1 Defining morphisms of baby ringed spaces

Prototypical example for this section: See next section.

Let $(X, \mathcal{O}_X)$ and $(Y, \mathcal{O}_Y)$ be baby ringed spaces, and think about how to define a morphism between them.

The guiding principle in algebra is that we want morphisms to be functions on underlying structure, but also respect the enriched additional data on top. To give some examples from the very beginning of time:

Example 54.1.1 (How to define a morphism)

- Consider groups. A group G has an underlying set (of elements), which we then enrich with a multiplication operation. So a homomorphism is a map of the underlying sets, plus it has to respect the group multiplication.
- Consider R -modules. Each R -module has an underlying abelian group, which we then enrich with scalar multiplication. So we require that a linear map respects the scalar multiplication as well, in addition to being a homomorphism of abelian groups.
- Consider topological spaces. A space X has an underlying set (of points), which we then enrich with a topology of open sets. So we consider maps of the set of points which respect the topology (pre-images of open sets are open).

This time, the ringed spaces $(X, \mathcal{O}_X)$ have an underlying *topological space*, which we have enriched with a structure sheaf. So, we want a continuous map $f : X \rightarrow Y$ of these topological spaces, which we then need to respect the sheaf of regular functions.

How might we do this? Well, if we let $\psi : Y \rightarrow \mathbb{C}$ be a regular function, then there's a natural way to get then composition gives us a way to write a map $X \rightarrow Y \rightarrow \mathbb{C}$. We then want to require that this is also a regular function.

More generally, we can take any regular function on Y and obtain some function on X , which we call a pullback. We then require that all the pullbacks are regular on X .

Definition 54.1.2. Let $(X, \mathcal{O}_X)$ and $(Y, \mathcal{O}_Y)$ be baby ringed spaces. Given a map $f : X \rightarrow Y$ and a regular function $\phi \in \mathcal{O}_Y(U)$, we define the **pullback** of ϕ , denoted $f^*\phi$, to be the composed function

$$f^{\text{pre}}(U) \xrightarrow{f} U \xrightarrow{\phi} \mathbb{C}.$$

The use of the word “pullback” is the same as in our study of differential forms.


Definition 54.1.3. Let $(X, \mathcal{O}_X)$ and $(Y, \mathcal{O}_Y)$ be baby ringed spaces. A continuous map of topological spaces $f : X \rightarrow Y$ is a **morphism** if every pullback of a regular function on Y is a regular function on X .

Two baby ringed spaces are **isomorphic** if there are mutually inverse morphisms between them, which we then call **isomorphisms**.

In particular, the pullback gives us a (reversed) *ring homomorphism*

$$f^* : \mathcal{O}_Y(U) \rightarrow \mathcal{O}_X(f^{\text{pre}}(U))$$

for *every* U ; thus our morphisms package a lot of information. Here's a picture of a morphism f , and the pullback of $\phi : U \rightarrow \mathbb{C}$ (where $U \subseteq Y$).



On a more philosophical point, we like this definition because it adheres to our philosophy of treating our varieties as intrinsic objects, rather than embedded ones. However, it is somewhat of a nuisance to actually verify it, and in what follows I will wave my hands a lot in claiming that something is a morphism, since doing so is mostly detail checking. The theorems which follow will give us alternative definitions of morphism which are more coordinate-based and easier to use for actual computations.

§54.2 Examples

Example 54.2.1 (The pullback of $\frac{1}{x-25}$ under $t \mapsto t^2$)

The map

$$f : X = \mathbb{A}^1 \rightarrow Y = \mathbb{A}^1 \quad \text{by} \quad t \mapsto t^2$$

is a morphism of varieties. For example, consider the regular function $\varphi = \frac{1}{x-25}$ on the open set $Y \setminus \{25\} \subseteq Y$. The f -inverse image is $X \setminus \{\pm 5\}$. Thus the pullback is

$$f^*\varphi : X \setminus \{\pm 5\} \rightarrow Y \setminus \{25\} \quad \text{by} \quad x \mapsto \frac{1}{x^2 - 25}$$

which is regular on $X \setminus \{\pm 5\}$.

More generally, given any polynomial $P(t)$, the map $t \mapsto P(t)$ will work.

Exercise 54.2.2. Let $X = \mathbb{A}^1$, $Y = \mathbb{A}^1$. By considering $\text{id} \in \mathcal{O}_Y(Y)$, show that no other regular functions exist.

In fact, let's generalize the previous exercise:

Theorem 54.2.3 (Regular maps of affine varieties are globally polynomials)

Let $X \subseteq \mathbb{A}^m$ and $Y \subseteq \mathbb{A}^n$ be affine varieties. Every morphism $f : X \rightarrow Y$ of varieties is given by

$$x = (x_1, \dots, x_m) \mapsto (P_1(x), \dots, P_n(x))$$

where $P_1, \dots, P_n$ are polynomials.

Proof. It's not too hard to see that all such functions work, so let's go the other way. Let $f : X \rightarrow Y$ be a morphism.

First, remark that $f^{\text{pre}}(Y) = X$. Now consider the regular function $\pi_1 \in \mathcal{O}_Y(Y)$, given by the projection $(y_1, \dots, y_n) \mapsto y_1$. Thus we need $f \circ \pi_1$ to be regular on X .

But for affine varieties $\mathcal{O}_X(X)$ is just the coordinate ring $\mathbb{C}[X]$ and so we know there is a polynomial P_1 such that $f \circ \pi_1 = P_1$. Similarly for the other coordinates. $\square$

Unfortunately, the situation is a little weirder in the projective setting. If $X \subseteq \mathbb{CP}^m$ and $Y \subseteq \mathbb{CP}^n$ are projective varieties, then every function

$$x = (x_0 : x_1 : \dots : x_m) \mapsto (P_0(x) : P_1(x) : \dots : P_n(x))$$

is a valid morphism, provided the P_i are homogeneous of the same degree and don't all vanish simultaneously. However if we try to repeat the proof for affine varieties we run into an issue: there is no π_1 morphism. (Would we send $(1 : 1) = (2 : 2)$ to 1 or 2?)

And unfortunately, there is no way to repair this. Counterexample:

Example 54.2.4 (Projective map which is not globally polynomial)

Let $V = \mathcal{V}_{\text{pr}}(xy - z^2) \subseteq \mathbb{CP}^2$. Then the map

$$V \rightarrow \mathbb{CP}^1 \quad \text{by} \quad (x : y : z) \mapsto \begin{cases} (x : z) & x \neq 0 \\ (z : y) & y \neq 0 \end{cases}$$

turns out to be a morphism of projective varieties. This is well defined just because $(x : z) = (z : y)$ if $x, y \neq 0$; this should feel reminiscent of the definition of regular function.

The good news is that “local” issues are the only limiting factor.

Theorem 54.2.5 (Regular maps of projective varieties are locally polynomials)

Let $X \subseteq \mathbb{CP}^m$ and $Y \subseteq \mathbb{CP}^n$ be projective varieties and let $f : X \rightarrow Y$ be a morphism. Then at every point $p \in X$ there exists a neighborhood $U_p \ni p$ and polynomials $P_0, P_1, \dots, P_n$ (which depend on U) so that

$$f(x) = (P_0(x) : P_1(x) : \dots : P_n(x)) \quad \forall x = (x_0 : \dots : x_m) \in U_p.$$

Of course the polynomials P_i must be homogeneous of the same degree and cannot vanish simultaneously on any point of U_p .

Example 54.2.6 (Example of an isomorphism)

In fact, the map $V = \mathcal{V}_{\text{pr}}(xy - z)^2 \rightarrow \mathbb{CP}^1$ is an isomorphism. The inverse map $\mathbb{CP}^1 \rightarrow V$ is given by

$$(s : t) \mapsto (s^2 : st : t^2).$$

Thus actually $V \cong \mathbb{CP}^1$.

Remark 54.2.7. Note also that if $X \cong Y$ as quasi-projective varieties, then we get a lot of information. For one thing, we know that X and Y are homeomorphic topological spaces. But moreover, we get bijections between all the rings $\mathcal{O}_X(f^{\text{pre}}(U))$ and $\mathcal{O}_Y(U)$; in particular, we have $\mathcal{O}_X(X) \cong \mathcal{O}_Y(Y)$.

One last useful example: we can have maps from affine spaces to projective spaces.

Example 54.2.8 (Embedding $\mathbb{A}^1 \hookrightarrow \mathbb{CP}^1$)

Consider a morphism

$$f : \mathbb{A}^1 \hookrightarrow \mathbb{CP}^1 \quad \text{by} \quad t \mapsto (t : 1).$$

This is also a morphism of varieties. (Can you see what the pullbacks look like?) This reflects the fact that $\mathbb{CP}^1$ is “ $\mathbb{A}^1$ plus a point at infinity”.

§54.3 Quasi-projective varieties

Prototypical example for this section: $\mathbb{A}^1 \setminus \{0\}$ might be best.

In light of the previous example, we saw $\mathbb{A}^1$ is in some sense a subset of $\mathbb{CP}^1$. However, $\mathbb{A}^1$ is not itself a projective variety, and so we want to loosen our definition a little:

Definition 54.3.1. A **quasi-projective variety** is an open set X of a projective variety V . It is a baby ringed space $(X, \mathcal{O}_X)$ too, because for any open set $U \subseteq X$ we simply define $\mathcal{O}_X(U) = \mathcal{O}_V(U)$.

Every projective variety is quasi-projective by taking $X = V$. But since we have a definition of morphism, we can now show:

Proposition 54.3.2 (Affine $\subseteq$ quasi-projective)

Every affine variety is isomorphic to a quasi-projective one.

Sketch of proof. We'll just do the example $V = \mathcal{V}(y - x^2)$. We take the projective variety $W = \mathcal{V}_{\text{pr}}(zy - x^2)$, and look at the corresponding standard open set $D(z)$. Then there is an isomorphism $V \rightarrow D(z)$ by $(x, y) \mapsto (x : y : 1)$, with inverse map given by $(x : y : z) \mapsto (x/z, y/z)$. $\square$

In summary, we have three different types of varieties:

- Affine varieties,
- Projective varieties, and

- Quasi-projective varieties.

These are baby ringed spaces, and so the notion of isomorphism shows that the last type subsumes the first two.

The reason quasi-projective varieties are interesting is that almost all varieties which occur in real life are quasi-projective. In fact, it is tricky even to exhibit a single example of a variety which is not quasi-projective. I certainly don't know one.

§54.4 Some applications

Natural question: are there quasi-projective varieties which are neither affine nor projective?

Our first guess might be to take the simplest projective variety, say $\mathbb{CP}^1$, and delete a point (to get an open set). This is quasi-projective, but it's isomorphic to $\mathbb{A}^1$. So instead we start with the simplest affine variety, say $\mathbb{A}^1$, and delete a point.

Surprisingly, this doesn't work.

Example 54.4.1 ($\mathbb{A}^1 \setminus \{0\}$ is affine)

Let $X = \mathbb{A}^1 \setminus \{0\}$ be an quasi-projective variety. We claim that in fact we have an isomorphism

$$X \cong V = \mathcal{V}(xy - 1) \subseteq \mathbb{A}^2$$

which shows that X is still isomorphic to an affine variety. The maps are

$$\begin{aligned} X &\leftrightarrow V \\ t &\mapsto (t, 1/t) \\ x &\leftarrow (x, y). \end{aligned}$$

Intuitively, the “hyperbola $y = 1/x$ ” in $\mathbb{A}^2$ can be projected onto the x -axis.

In fact, deleting any number of points from $\mathbb{A}^1$ fails if we delete $\{1, 2, 3\}$, the resulting open set is isomorphic as a baby ringed space to $\mathcal{V}(y(x-1)(x-2)(x-3)-1)$, which colloquially might be called $y = \frac{1}{(x-1)(x-2)(x-3)}$.

More generally, we have:

Theorem 54.4.2 (Distinguished open subsets of affines are affine)

Consider $X = D(f) \subseteq V = \mathcal{V}(f_1, \dots, f_m) \subseteq \mathbb{A}^n$, where V is an affine variety, and the distinguished open set X is thought of as a quasi-projective variety. Define

$$W = \mathcal{V}(f_1, \dots, f_m, y \cdot f - 1) \subseteq \mathbb{A}^{n+1}$$

where y is the $(n+1)$ st coordinate of $\mathbb{A}^{n+1}$.

Then $X \cong W$.

Fortunately, the true example is not too far away. First, check:

Question 54.4.3. Let $f : X \rightarrow Y$ be a morphism of quasi-projective varieties. Take a point $p \in X$ and let $q = f(p) \in Y$. Let φ be a regular function on Y .

Prove that if $\varphi(q) = 0$ then $(f^*\varphi)(p) = 0$.

Example 54.4.4 (Punctured plane is not affine)

We claim that $X = \mathbb{A}^2 \setminus \{(0, 0)\}$ is not affine. This is confusing, so you will have to pay attention.

Assume for contradiction there is an affine variety V and an isomorphism

$$f : X \rightarrow V.$$

Then taking the pullback we get a ring isomorphism

$$f^* : \mathcal{O}_V(V) \rightarrow \mathcal{O}_X(X) = \mathbb{C}[x, y].$$

Now let $\mathcal{O}_V(V) = \mathbb{C}[a, b]$ where $f^*(a) = x$, $f^*(b) = y$. In particular, we actually have to have $V \cong \mathbb{A}^2$.

Now in the *affine* variety V we can take $\mathcal{V}(a)$ and $\mathcal{V}(b)$. These have nonempty intersection since (a, b) is a maximal ideal in $\mathcal{O}_V(V)$. Call this point q . Now let $p = f^{-1}(q)$. Then by the previous question, p should have vanish on both the regular functions x and y in $\mathcal{O}_X(X)$. But there is no point in X with this property, which is a contradiction.

§54.5 Problems to think about

Problem 54A. Consider the map

$$\mathbb{A}^1 \rightarrow \mathcal{V}(y^2 - x^3) \subseteq \mathbb{A}^2 \quad \text{by} \quad t \mapsto (t^2, t^3).$$

Show that it is a morphism of varieties, but it is not an isomorphism.

Problem 54B[†]. Show that every projective variety has a neighborhood which is isomorphic to an affine variety. In this way, “projective varieties are locally affine”.

Problem 54C. Let V be a affine variety and let W be a irreducible projective variety. Prove that $V \cong W$ if and only if V and W are a single point.

55 Sheaves and ringed spaces

Most of the complexity of the affine variety V earlier comes from $\mathcal{O}_V$. This is a type of object called a “sheaf”. The purpose of this chapter is to completely define what this sheaf is, and just what it is doing.

The typical example to keep in mind is a sheaf of “functions with property P ” on a topological space X : for every open set U , $\mathcal{F}(U)$ gives us the ring of functions on X . However, we will work very abstractly and only assume $\mathcal{F}(U)$ is a ring, without an interpretation as “functions”.

The payoff for this abstraction is that it will allow us to define an arbitrary scheme in the next chapter. Varieties use $\mathbb{C}[x_1, x_2, \dots, x_n]/I$ as their “ring of functions”, and by using the fully general sheaf we will be replace this with *any* commutative ring. In particular, we can use the case where I is not radical, such as $\mathbb{C}[x]/(x^2)$; this gives the “multiplicity” behavior that we sought after all along.

§55.1 Pre-sheaves

Prototypical example for this section: The sheaf of holomorphic (or regular, continuous, differentiable, constant, whatever) functions.

The proper generalization of our $\mathcal{O}_V$ is a so-called sheaf of rings. Recall that $\mathcal{O}_V$ took open sets of V to rings, with the interpretation that $\mathcal{O}_V(U)$ was a “ring of functions”.

In light of this, we first define:

Definition 55.1.1. For a topological space X let $\text{OpenSets}(X)$ denote its open sets of X .

Definition 55.1.2. A **pre-sheaf** of rings on a space X is a function

$$\mathcal{F} : \text{OpenSets}(X) \rightarrow \text{Rings}$$

meaning each open set gets associated with a ring $\mathcal{F}(U)$. Each individual element of $\mathcal{F}(U)$ is called a **section**. An element of $\mathcal{F}(X)$ is called a **global section**.

It is also equipped with a **restriction map** for any $U_1 \subseteq U_2$; this is a map

$$\text{res}_{U_1, U_2} : \mathcal{F}(U_2) \rightarrow \mathcal{F}(U_1)$$

such that $\text{res}_{U, U}$ is the identity and whenever $U_1 \subseteq U_2 \subseteq U_3$ the diagram

$$\begin{array}{ccc} \mathcal{F}(U_3) & \xrightarrow{\text{res}_{U_2, U_3}} & \mathcal{F}(U_2) \\ & \searrow \text{res}_{U_1, U_3} & \downarrow \text{res}_{U_1, U_2} \\ & & \mathcal{F}(U_1) \end{array}$$

commutes. (Restricting “big to medium to small” is the same as “big to small”.)

Abuse of Notation 55.1.3. If $s \in \mathcal{F}(U_2)$ is some section and $U_1 \subseteq U_2$, then rather than write $\text{res}_{U_1, U_2}(s)$ I will write $s|_{U_1}$ instead: “ s restricted to U_1 ”. This is abuse of notation because the section s is just an element of some ring, and in the most abstract of cases may not have a natural interpretation as function.

Example 55.1.4 (Examples of pre-sheaves)

- (a) For an affine variety V , $\mathcal{O}_V$ is of course a sheaf, with $\mathcal{O}_V(U)$ being the ring of regular functions on U . The restriction map just says that if $U_1 \subseteq U_2$, then a function $s \in \mathcal{O}_V(U_2)$ can also be thought of as a function $s|_{U_1} \in \mathcal{O}_V(U_1)$, hence the name “restriction”. The commutativity of the diagram then follows.
- (b) Let $X \subseteq \mathbb{R}^n$ be an open set. Then there is a sheaf of smooth/differentiable/etc. functions on X . In fact, one can do the same construction for any manifold M .
- (c) Similarly, if $X \subseteq \mathbb{C}$ is open, we can construct a sheaf of holomorphic functions on X .

In all these examples, the sections $s \in \mathcal{F}(U)$ are really functions on the space, but in general they need not be.

Now, we give a second, equivalent and far shorter definition of pre-sheaf:

Abuse of Notation 55.1.5. By abuse of notation, $\text{OpenSets}(X)$ will also be thought of as a posetal category by inclusion. Thus $\emptyset$ is an initial object and the entire space X is a terminal object.

Definition 55.1.6. A **pre-sheaf** of rings on X is a contravariant functor

$$\mathcal{F} : \text{OpenSets}(X)^{\text{op}} \rightarrow \text{Rings}.$$

Question 55.1.7. Check that these definitions are equivalent.

It is now clear that we can actually replace **Rings** with any (concrete) category we want. So we should think:

Pre-sheaves should be thought of as “returning the ring of functions with a property P ”.

§55.2 Sheaves

Prototypical example for this section: Constant functions aren’t pre-sheaves, but locally constant ones are.

Now, the main idea next is that

Sheaves are pre-sheaves for which P is a local property.

The formal definition doesn’t illuminate this as much as the examples do, but I have to give it first for the examples to make sense.


Definition 55.2.1. A **sheaf** $\mathcal{F}$ is a pre-sheaf obeying two additional axioms: Suppose U is covered by open sets $U_\alpha \subseteq U$. Then:

1. (Identity) If $s, t \in \mathcal{F}(U)$ are sections, and $s|_{U_\alpha} = t|_{U_\alpha}$ for all α , then $s = t$.
2. (Collation) Consider sections $s_\alpha \in \mathcal{F}(U_\alpha)$ for each α . Suppose that

$$s_\alpha|_{U_\alpha \cap U_\beta} = s_\beta|_{U_\alpha \cap U_\beta}$$

for each U_α and U_β . Then we can find $s \in \mathcal{F}(U)$ such that $s|_{U_\alpha} = s_\alpha$.

This is best illustrated by picture: consider an open cover $U_1 \cup U_2$.



Then for a sheaf of functions, the axioms are saying that:

- If s and t are functions (with property P) on the whole space $U = U_1 \cup U_2$, and $s|_{U_1} = t|_{U_1}$, $s|_{U_2} = t|_{U_2}$, then $s = t$ on the entire union. This is clear.
- If s_1 is a function with property P on U_1 and s_2 is a function with property P on U_2 , and the two functions agree on the overlap, then one can collate them to obtain a function s on the whole space: this is obvious, but **the catch is that the collated function needs to have property P as well**. That's why it matters that P is local.

The reason we need these axioms is that in our abstract definition of a sheaf, the output of the sheaf is an abstract ring, which need not actually have a concrete interpretation as “functions on X ”, even though our examples will usually have this property.

Now for the examples, which are more enlightening:

Example 55.2.2 (Examples and non-examples of sheaves)

- Pre-sheaves of arbitrary / continuous / differentiable / smooth / holomorphic functions are still sheaves. This is because to verify a function is continuous, one only needs to look at small neighborhoods at once.
- For a complex variety V , $\mathcal{O}_V$ is a sheaf, precisely because our definition was *locally* quotients of polynomials.
- The pre-sheaf of *constant* real functions on a space X is *not* a sheaf, because it fails the collation axiom. Namely, suppose that $U_1 \cap U_2 = \emptyset$. Then if s_1 is the constant function 1 on U_1 while s_2 is the constant function 2 on U_2 , then we cannot collate these to a constant function on $U_1 \cup U_2$.
- On the other hand, *locally constant* functions do produce a sheaf. (A function is locally constant if for every point it is constant on some neighborhood.)

In fact, the sheaf in (d) is what is called a *sheafification* of the pre-sheaf constant functions, which we define momentarily.

§55.3 Stalks

Prototypical example for this section: Germs of real smooth functions tell you the derivatives, but germs of holomorphic functions determine the entire function.

Let $\mathcal{F}$ be a pre-sheaf. If we have a function $s \in \mathcal{F}(U)$ and a point $p \in U$, then in general it doesn't make sense to ask what $s(p)$ is (even though all our examples look like this), because $\mathcal{F}(U)$ is an arbitrary ring. So, we will replace the notion of $s(p)$ with a so-called *germ*.

Definition 55.3.1. Let $\mathcal{F}$ be a pre-sheaf of rings. For every point p we define the **stalk** $\mathcal{F}_p$ to be the set

$$\{(s, U) \mid s \in \mathcal{F}(U), p \in U\}$$

modulo the relation $\sim$ that

$$(s_1, U_1) \sim (s_2, U_2) \text{ if } s_1|_{U_1 \cap U_2} = s_2|_{U_1 \cap U_2}.$$

The equivalence classes themselves are called **germs**.

Definition 55.3.2. The germ of a given $s \in \mathcal{F}(U)$ at a point p is the equivalence class for $(s, U) \in \mathcal{F}_p$. We denote this by $[s]_p$.

Remark 55.3.3. It is almost never useful to think of a germ as an ordered pair, since the set U can get arbitrarily small. Instead, one should think of a germ as a “shred” of some section near p .

So what’s happening is: We consider functions s defined near p and, since we only care about local behavior, we identify any two functions agreeing on a neighborhood of p , no matter how small.

Notice that the stalk is itself a ring as well: for example, addition is done by

$$(s_1, U_1) + (s_2, U_2) = (s_1|_{U_1 \cap U_2} + s_2|_{U_1 \cap U_2}, U_1 \cap U_2).$$

So the germ $[s]_p$ now plays the role of the “value” $s(p)$. But actually, it carries much more information than that.

Example 55.3.4 (Germs of real smooth functions)

Let $X = \mathbb{R}$ and let $\mathcal{F}$ be the sheaf on X of smooth functions (i.e. $\mathcal{F}(U)$ is the set of smooth real-valued functions on U).

Consider a global section, $s : \mathbb{R} \rightarrow \mathbb{R}$ (thus $s \in \mathcal{F}(X)$) and its germ at 0.

- (a) From the germ we can read off $s(0)$, obviously.
- (b) We can also find $s'(0)$, because the germ carries enough information to compute the limit $\lim_{h \rightarrow 0} \frac{1}{h}[s(h) - s(0)]$.
- (c) Similarly, we can compute the second derivative and so on.
- (d) However, we can’t read off, say, $s(3)$ from the germ. For example, take

$$s(x) = \begin{cases} e^{-\frac{1}{x-1}} & x > 1 \\ 0 & x \leq 1. \end{cases}$$

Note $s(3) = e^{-\frac{1}{2}}$, but $[\text{zero function}]_0 = [s]_0$. So germs can’t distinguish between the zero function and s .

Example 55.3.5 (Germes of holomorphic functions)

Holomorphic functions are very strange in this respect. Consider the sheaf $\mathcal{F}$ on $\mathbb{C}$ of *holomorphic* functions.

Take $s : \mathbb{C} \rightarrow \mathbb{C}$ a global section. Given the germ of s at 0, we can read off $s(0)$, $s'(0)$, et cetera. The miracle of complex analysis is that just knowing the derivatives of s at zero is enough to reconstruct all of s : we can compute the Taylor series of s now. **Thus germs of holomorphic functions determine the entire function**; they “carry much more information” than their real counterparts.

In particular, we can concretely describe the sheaf:

$$\mathcal{F}_p = \left\{ \sum_{k \geq 0} c_k (z - p)^k \text{ convergent near } p \right\}.$$

In particular, this includes germs of meromorphic functions, so long as there is no pole at p itself.

And of course, our algebraic geometry example.

Abuse of Notation 55.3.6. Rather than writing $(\mathcal{O}_X)_p$ we will write $\mathcal{O}_{X,p}$.

Example 55.3.7 (Stalks of the sheaf on an affine variety)

Let $V \subseteq \mathbb{A}^n$ be a variety, and assume $p \in V$. Then, a regular function φ on $U \subseteq V$ is supposed to be a function on U that “locally” is a quotient of two functions in $\mathbb{C}[V]$.

However, **as far as the germ is concerned, we only care about whichever quotient applies near the point p** . In light of this, we only think about the representation at p , and ignore the local clause completely: thus the entire stalk can be thought of as

$$\mathcal{O}_{V,p} = \left\{ \left(\frac{f}{g}, U \right) \mid U \ni p, f, g \in \mathbb{C}[V], g \neq 0 \text{ on } U \right\}$$

modulo the usual relations.

Now, since we happen to be working with complex polynomials, we know that a rational function is determined by its behavior on any neighborhood of p (complex analysis forever!). Thus

$$\mathcal{O}_{V,p} = \left\{ \frac{f}{g} \mid f, g \in \mathbb{C}[V], g(p) \neq 0 \right\}.$$

which don’t vanish on p .

Remark 55.3.8 (For category lovers). You may notice that $\mathcal{F}_p$ seems to be “all the $\mathcal{F}_p(U)$ coming together”, where $p \in U$. And in fact, $\mathcal{F}_p(U)$ is the categorical *limit* of the diagram formed by all the $\mathcal{F}(U)$ such that $p \in U$. This is often written

$$\mathcal{F}_p = \varinjlim_{U \ni p} \mathcal{F}(U)$$

Thus we can define stalks in any category with limits, though to be able to talk about germs the category needs to be concrete.

§55.4 Sections “are” sequences of germs

Prototypical example for this section: A real function on U is a sequence of real numbers $f(p)$ for each $p \in U$ satisfying some local condition. Analogously, a section $s \in \mathcal{F}(U)$ is a sequence of germs satisfying some local compatibility condition.

Let $\mathcal{F}$ be a sheaf on a space X now. The purpose of this section is to convince you that the correct way to think about a section $s \in \mathcal{F}(U)$ is as a sequence of germs above every point $p \in U$.

How do we do this? Given $s \in \mathcal{F}(U)$ we consider its germ $[s]_p$ above every $p \in U$. To visualize, picture an open set $U \subseteq X$, and above every point $p \in U$ imagine there is a little speck $[s]_p$, which bundles the information of s near p .

Example 55.4.1 (Real functions vs. germs)

Let X be a space and let $\mathcal{F}$ be the sheaf of smooth functions. Given a section $f \in \mathcal{F}(U)$,

- As a function, f is just a choice of value $f(p) \in \mathbb{R}$ at every point p , subject to a local “smooth” condition.
- Let’s now think of f as a sequence of germs. At every point p the germ $[f]_p \in \mathcal{F}_p$ gives us the value $f(p)$ as we described above. The germ packages even more data than this: from the germ $[f]_p$ alone we can for example compute $f'(p)$. Nonetheless we stretch the analogy and think of f as a choice of germ $[f]_p \in \mathcal{F}_p$ at each point p .

Thus we can replace the notion of the value $f(p)$ with germ $[f]_p$. This is useful because in a general sheaf $\mathcal{F}$, the notion $s(p)$ is not defined while the notion $[s]_p$ is.

From the above example it’s obvious that if we know each germ $[s]_p$ this should let us reconstruct the entire section s . Let’s check this from the sheaf axioms:

Exercise 55.4.2 (Sections are determined by stalks). Let $\mathcal{F}$ be a sheaf. Consider the natural map $\mathcal{F}(U) \rightarrow \prod_{p \in U} \mathcal{F}_p$ described above. Show that this map is injective, i.e. the germs of s at every point $p \in U$ determine the section s . (You will need the “identity” sheaf axiom, but not “collating”.)

However, this map is clearly not surjective!

Question 55.4.3. Come up with a counterexample to break surjectivity. (This is like asking “come up with a non-smooth function”.)

Nonetheless we can describe the image: we want a sequence of germs $(g_p)_{p \in U}$ such that near every germ g_p , the germs g_q are “compatible” with g_p . We make this precise:

Definition 55.4.4. Let $\mathcal{F}$ be sheaf (or even a pre-sheaf) and U open. A sequence $(g_p)_{p \in U}$ of germs ($g_p \in \mathcal{F}_p$) is said to be **compatible** if they can be “locally collated”:

For any $p \in U$ there exists a neighborhood $U_p \ni p$ and a section $s \in \mathcal{F}(U_p)$ on it such that $[s]_q = g_q$ for each $q \in U_p$.

Intuitively, the germs should “collate together” to some section near each *individual* point q (but not necessarily to a section on all of U).

We let the reader check this definition is what we want:

Exercise 55.4.5. Prove that any choice of compatible germs over U collates together to a section of $\mathcal{F}$. (You will need the “collating” sheaf axiom, but not “identity”.)

Putting together the previous two exercise gives:

Theorem 55.4.6 (Sections “are” just compatible germs)

Let $\mathcal{F}$ be a sheaf. There is a natural bijection between

- sections of $\mathcal{F}(U)$, and
- sequences of compatible germs over U .

This is in exact analogy to the way that e.g. a smooth real-valued function on U is a choice of real number $f(p) \in \mathbb{R}$ at each point $p \in U$ satisfying a local smoothness condition.

Thus the notion of stalks is what lets us recover the viewpoint that sections are “functions”. Therefore for theoretical purposes,

You should usually think of a section $s \in \mathcal{F}(U)$ as a sequence of germs.

In particular, this makes restriction morphisms easy to deal with: just truncate the sequence of germs!

§55.5 Sheafification

Prototypical example for this section: The pre-sheaf of locally constant functions becomes the sheaf of constant functions.

Now that we have the language of germs, we can define the so-called sheafification. The idea is that if $\mathcal{F}$ is the pre-sheaf of “functions with property P ” then we want to associate a sheaf $\mathcal{F}^{\text{sh}}$ of “functions which are locally P ”, which makes them into a sheaf. We have already seen two examples of this:

Example 55.5.1 (Sheafification)

- If X is a topological space, and $\mathcal{F}$ is the pre-sheaf of constant functions on open sets of X , then $\mathcal{F}^{\text{sh}}$ is the sheaf of locally constant functions.
- If V is an affine variety, and $\mathcal{F}$ is the pre-sheaf of rational functions, then $\mathcal{F}^{\text{sh}}$ is the sheaf of regular functions (which are locally rational).

So how do we encode “locally P ”? Answer: we saw last time that for a sheaf $\mathcal{F}$, compatible germs biject to sections. So we *force* this to be true by defining:

Definition 55.5.2. The **sheafification** $\mathcal{F}^{\text{sh}}$ of a pre-sheaf $\mathcal{F}$ is defined by

$$\mathcal{F}^{\text{sh}}(U) = \{\text{sequences of compatible germs } (g_p)_{p \in U}\}.$$

Exercise 55.5.3. Reconcile this definition with the two examples we gave.

Abuse of Notation 55.5.4. Technically I haven't told you what the restriction morphisms of $\mathcal{F}^{\text{sh}}(U)$ are but, they are not hard to guess. I'll usually be equally sloppy in the future: when defining a sheaf $\mathcal{F}$, I'll only say what $\mathcal{F}(U)$ is, with the restriction morphisms $\mathcal{F}(U_2) \rightarrow \mathcal{F}(U_1)$ being implicit.

The construction is contrived so that given a section $(g_p)_{p \in U} \in \mathcal{F}^{\text{sh}}(U)$ the germ at a point p is g_p :

Lemma 55.5.5 (Pre-sheaves and sheaves have the same stalk)

Let $\mathcal{F}$ be a pre-sheaf and $\mathcal{F}^{\text{sh}}$ its sheafification. Then for any point q , there is an isomorphism

$$(\mathcal{F}^{\text{sh}})_q \cong \mathcal{F}_q.$$

Proof. A germ in $(\mathcal{F}^{\text{sh}})_q$ looks like $((g_p)_{p \in U}, U)$, where $g_p = (s_p, U_p)$ are themselves germs of $\mathcal{F}_p$, and $q \in U$. Then the isomorphism is given by

$$((g_p)_{p \in U}, U) \mapsto g_q \in \mathcal{F}_q.$$

The inverse map is given by for each $g = (s, U) \in \mathcal{F}_q$ by

$$g \mapsto ((g)_{p \in U}, U) \in (\mathcal{F}^{\text{sh}})_q$$

i.e. the sequence of germs is the constant sequence. □

We'll later see that if $\mathcal{F}$ is already a sheaf, then $\mathcal{F}^{\text{sh}} = \mathcal{F}$.

§55.6 Morphisms of sheaves

First, recall that a sheaf is a contravariant functor (pre-sheaf) with extra conditions. In light of this, it is not hard to guess the definition of a morphism of pre-sheaves:

Definition 55.6.1. A **morphism of (pre-)sheaves** $\alpha : \mathcal{F} \rightarrow \mathcal{G}$ on the same space X is a **natural transformation** of the underlying functors. Isomorphism of sheaves is defined in the usual way.

Question 55.6.2. Show that this amounts to: for each $U \subseteq X$ we need to specify a morphism $\alpha_U : \mathcal{F}(U) \rightarrow \mathcal{G}(U)$ such that the diagram

$$\begin{array}{ccc} \mathcal{F}(U_2) & \xrightarrow{\alpha_{U_2}} & \mathcal{G}(U_2) \\ \text{res}_{U_1, U_2} \downarrow & & \downarrow \text{res}_{U_1, U_2} \\ \mathcal{F}(U_1) & \xrightarrow{\alpha_{U_1}} & \mathcal{G}(U_1) \end{array}$$

commutes any time that $U_1 \subseteq U_2$.

However, in the sheaf case we like stalks more than sections because they are theoretically easier to think about. And in fact:

Proposition 55.6.3 (Morphisms determined by stalks)

A morphism of sheaves $\alpha : \mathcal{F} \rightarrow \mathcal{G}$ induces a morphism of stalks

$$\alpha_p : \mathcal{F}_p \rightarrow \mathcal{G}_p$$

for every point $p \in X$. Moreover, the sequence $(\alpha_p)_{p \in X}$ determines α uniquely.

Proof. The morphism α_p itself is just $(s, U) \mapsto (\alpha_U(s), U)$.

Question 55.6.4. Show this is well-defined.

Now suppose $\alpha, \beta : \mathcal{F} \rightarrow \mathcal{G}$ satisfy $\alpha_p = \beta_p$ for every p . We want to show $\alpha_U(s) = \beta_U(s)$ for every $s \in \mathcal{F}(U)$.

Question 55.6.5. Verify this using the description of sections as sequences of germs. $\square$

Thus a morphism of sheaves can be instead modelled as a morphism of all the stalks. We will see later on that this viewpoint is quite useful.

§55.7 Local rings, and locally ringed spaces

Prototypical example for this section: Smooth functions f on $X \subseteq \mathbb{R}^n$ have invertible germs at p unless $f(p) = 0$.

The stalks of the examples we produced above are special types of rings, called *local rings*. Algebraically, the definition of these is:

Definition 55.7.1. A **local ring** R is a ring with exactly one maximal ideal.

Exercise 55.7.2. Show that a ring R is a local ring if there exists a proper ideal $\mathfrak{m} \subsetneq R$ such that all elements of $R \setminus \mathfrak{m}$ are units. (Bonus: prove the converse.)

To see why this definition applies to the stalks above, we need to identify what the maximal ideal is. Let's go back to the example of $X = \mathbb{R}$ and $\mathcal{F}(U)$ the smooth functions, and consider the stalk $\mathcal{F}_p$, where $p \in X$. Define the ideal $\mathfrak{m}_p$ to be the set of germs (s, U) for which $s(p) = 0$.

Then $\mathfrak{m}_p$ is maximal: we have an exact sequence

$$0 \rightarrow \mathfrak{m}_p \rightarrow \mathcal{F}_p \xrightarrow{(s, U) \mapsto s(p)} \mathbb{R} \rightarrow 0$$

and so $\mathcal{F}_p / \mathfrak{m}_p \cong \mathbb{R}$, which is a field.

It remains to check there are no nonzero maximal ideals. Now note that if $s \notin \mathfrak{m}_p$, then s is nonzero in some neighborhood of p , then one can construct the function $1/s$ in a neighborhood of p . So **every element of $\mathcal{F}_p \setminus \mathfrak{m}_p$ is a unit**; $\mathfrak{m}_p$ is in fact the only maximal ideal!

More generally,

If $\mathcal{F}$ consists of “field-valued functions”, the stalk $\mathcal{F}_p$ probably has a maximal ideal consisting of the germs vanishing at p .

The discussion above implies, for example:

Proposition 55.7.3 (Stalks are often local rings)

The stalks of each of the following types of sheaves are local rings:

- (a) Sheaves of continuous real/complex functions on a topological space
- (b) Sheaves of smooth functions on any manifold
- (c) Regular functions on an algebraic variety V .

We can now define:

Definition 55.7.4. A **ringed space** is a topological space X equipped with a sheaf $\mathcal{O}_X$ of rings. Suppose that for every point p , the stalk $\mathcal{O}_{X,p}$ is a local ring. Then we say that $(X, \mathcal{O}_X)$ is a **locally ringed space**. We denote the maximal ideals by $\mathfrak{m}_{X,p}$.

In particular, in the previous chapter we showed that every affine variety could be built into a locally ringed space. Hooray!

Abuse of Notation 55.7.5. A ringed space $(X, \mathcal{O}_X)$ is abbreviated to just X , while $p \in X$ means “ p is in the topological space X ”.

§55.8 Morphisms of (locally) ringed spaces

Finally, it remains to define a morphism of locally ringed space. To do this we have to build up in several steps.

Remark 55.8.1. I always secretly felt that one can probably get away with just suspending belief and knowing “there is a reasonable definition of morphisms of locally ringed spaces”. Daring readers are welcome to try this approach.

Morphisms of ringed spaces

We talked about the morphisms of locally ringed spaces when we discussed projective varieties. We want to copy the same construction here.

The idea was as follows: to “preserve” the structure via $f : X \rightarrow Y$, we took every regular function on Y and made it into a function on X by composing it with f . This gave us a ring homomorphism $f^* : \mathcal{O}_Y(U) \rightarrow \mathcal{O}_X(f^{\text{pre}}(U))$ for every open set U in Y .

Unfortunately, in the general case we don’t have a notion of function, so “composing by f ” doesn’t make sense. The solution to this is to *include all the ring homomorphisms* f^* in the data for a morphism $f : X \rightarrow Y$; we package this information using the morphism of sheaves.

Here are the full details. Up above we only defined a morphism of sheaves on the same space, while right now we have a sheaf on X and Y . So we have to first “push” the sheaf on X into a sheaf on Y by

$$U \mapsto \mathcal{O}_X(f^{\text{pre}}(U)).$$

This is important enough to have a name:

Definition 55.8.2. Let $\mathcal{F}$ be a sheaf on X , and $f : X \rightarrow Y$ a continuous map. The **pushforward sheaf** $f_*\mathcal{F}$ on Y is defined by

$$(f_*\mathcal{F})(U) = \mathcal{F}(f^{\text{pre}}(U)) \quad \forall U \subseteq Y.$$

This makes sense, since $f^{\text{pre}}(U)$ is open in X .

Picture:

$$\begin{array}{ccc}
 X & \xrightarrow{f} & Y \\
 \text{OpenSets}(X)^{\text{op}} & \xleftarrow{f^{\text{pre}}} & \text{OpenSets}(Y)^{\text{op}} \\
 & \searrow \mathcal{F} & \downarrow f_*\mathcal{F} \\
 & & \text{Rings}
 \end{array}
 \qquad
 \begin{array}{ccc}
 X & \xrightarrow{f} & Y \\
 f^{\text{pre}}(U) & \xleftarrow{f^{\text{pre}}} & U \\
 & \searrow \mathcal{F} & \downarrow f_*\mathcal{F} \\
 & & \mathcal{F}(f^{\text{pre}}(U)) = (f_*\mathcal{F})(U) \in \text{Rings}
 \end{array}$$

I haven't actually checked that $f_*\mathcal{F}$ is a sheaf (as opposed to a pre-sheaf), but this isn't hard to do. Also, as f^{pre} is a functor $\text{OpenSets}(Y)^{\text{op}} \rightarrow \text{OpenSets}(X)^{\text{op}}$, the pushforward $f_*\mathcal{F}$ is just the composition of these two functors.

Question 55.8.3. Technically $f_*\mathcal{F}$ is supposed to be a functor $\text{OpenSets}(Y)^{\text{op}} \rightarrow \text{Rings}$, so it also needs to come with restriction arrows. What are they?

Now that we moved both sheaves onto Y , we mimic the construction with quasi-projective varieties by asking for a sheaf morphism from $\mathcal{O}_Y$ to $f_*\mathcal{O}_X$:

Definition 55.8.4. A **morphism of ringed spaces** $f : (X, \mathcal{O}_X) \rightarrow (Y, \mathcal{O}_Y)$ consists of a continuous map of topological spaces $f : X \rightarrow Y$, plus additionally a morphism of sheaves $f^* : \mathcal{O}_Y \rightarrow f_*\mathcal{O}_X$.

Exercise 55.8.5. (Mandatory) Check that this coincides with our work on baby ringed spaces if we choose f^* to be the pullback.

Morphisms of locally ringed spaces

Last step! Suppose now that $(X, \mathcal{O}_X)$ and $(Y, \mathcal{O}_Y)$ are locally ringed spaces. Thus we need to deal with some information about the stalks.

Given a morphism of ringed spaces $f : (X, \mathcal{O}_X) \rightarrow (Y, \mathcal{O}_Y)$, we can actually use the f_U^* above to induce maps on the stalks, as follows. For a point $p \in X$, construct the map

$$\begin{aligned}
 f_p^* : \mathcal{O}_{Y, f(p)} &\longrightarrow \mathcal{O}_{X, p} \\
 (s, U) &\longmapsto (f_U^*(s), f^{\text{pre}}(U)).
 \end{aligned}$$

where $s \in \mathcal{O}_Y(U)$.

Definition 55.8.6. Let R and S be local rings with maximal ideals $\mathfrak{m}_R$ and $\mathfrak{m}_S$. A **morphism of local rings** is a homomorphism of rings $\psi : R \rightarrow S$ such that $\psi^{\text{pre}}(\mathfrak{m}_S) = \mathfrak{m}_R$.

Definition 55.8.7. A **morphism of locally ringed spaces** is a morphism of ringed spaces $f : (X, \mathcal{O}_X) \rightarrow (Y, \mathcal{O}_Y)$ such that for every point p the induced map of stalks is a morphism of local rings.

Recalling that $\mathfrak{m}_{X, p}$ is the maximal ideal of $\mathcal{O}_X$ at p , the new condition is saying that

$$(f_p^*)^{\text{pre}}(\mathfrak{m}_{X, p}) = \mathfrak{m}_{Y, f(p)}.$$

Concretely, if g is a function vanishing on $f(p)$, then its pullback $f_U^*(g)$ vanishes on p .

This completes the definition of a morphism of locally ringed spaces. Isomorphisms of (locally) ringed spaces are defined in the usual way.

§55.9 Problems to think about

Problem 55A. Prove that if $\mathcal{F}$ is a sheaf, then $\mathcal{F} \cong \mathcal{F}^{\text{sh}}$.

Problem 55B*. Suppose X is a finite topological space equipped with the discrete topology (i.e. X is a finite set of points). Let $\mathcal{F}$ be a sheaf on X . Show that for any open set U , we have a ring isomorphism

$$\mathcal{F}(U) \cong \prod_{p \in U} \mathcal{F}_p.$$

Here the product is the product ring as defined in Example 12.2.10.

Problem 55C[†] (Skyscraper sheaf). Let Y be a *Hausdorff* topological space and let $p \in Y$ be a point. Fix a ring R . Construct a sheaf $\mathcal{F}$ on Y by

$$\mathcal{F}(U) = \begin{cases} R & p \in U \\ 0 & \text{otherwise} \end{cases}$$

with restriction maps in the obvious manner. We call $\mathcal{F}$ a **skyscraper sheaf** at p . Compute all the stalks of $\mathcal{F}$.

Problem 55D[†] (Pushforward sheaf). Suppose $f : X \rightarrow Y$ is a continuous map of spaces and $\mathcal{F}$ is a sheaf on X . Define a sheaf $f_*\mathcal{F}$ on Y from $\mathcal{F}$; we call this the pushforward of $\mathcal{F}$ onto Y .

Problem 55E. Interpret the skyscraper sheaf as the pushforward of a constant sheaf on a one-point space.

56 Schemes

Now that we understand sheaves well, we can readily define the scheme. It will be a locally ringed space, so we need to define

- The set of points,
- The topology on it, and
- The structure sheaf on it.

In the case of $\mathbb{A}^n$, we used $\mathbb{C}^n$ as the set of points and $\mathbb{C}[x_1, \dots, x_n]$ as the ring of functions but then remarked that the set of points of $\mathbb{C}^n$ corresponded to the maximal ideals of $\mathbb{C}[x_1, \dots, x_n]$.

In an *affine scheme*, we will take an *arbitrary* ring R , and generate the entire structure from just R itself. The final result is called $\text{Spec } R$, the **spectrum** of R . The affine varieties $\mathcal{V}(I)$ we met earlier will just be $\mathbb{C}[x_1, \dots, x_n]/I$, but now we will be able to take *any* ideal I , thus finally completing the table at the end of the “affine variety” chapter.

In particular, $\text{Spec } \mathbb{C}[x]/(x^2)$ will be the double point we sought for so long.

§56.1 The set of points

Prototypical example for this section: $\text{Spec } \mathbb{C}[x_1, \dots, x_n]/I$.

First surprise, for a ring R :

Spec R is defined as the set of prime ideals of R .

This might be a little surprising, since we might have guessed that $\text{Spec } R$ should just have the maximal ideals. What do the remaining ideals correspond to? The answer is that they will be so-called *generic points* points which are “somewhere” in the space, but nowhere in particular.

Remark 56.1.1. As usual R itself is not a prime ideal, but (0) is prime if R is an integral domain.

Example 56.1.2 (Examples of spectrums)

- $\text{Spec } \mathbb{C}[x]$ consists of a point $(x - a)$ for every $a \in \mathbb{C}$, which correspond to what we geometrically think of as $\mathbb{A}^1$. It additionally consists of a point (0) , which we think of as a “generic point”, nowhere in particular.
- $\text{Spec } \mathbb{C}[x, y]$ consists of points $(x - a, y - b)$ (which are the maximal ideals) as well as (0) again, a generic point that is thought of as “somewhere in $\mathbb{C}^2$, but nowhere in particular”. It also consists of generic points corresponding to irreducible polynomials $f(x, y)$, for example $(y - x^2)$, which is a “generic point on the parabola”.
- If k is a field, $\text{Spec } k$ is a single point, since the only maximal ideal of k is (0) .

Example 56.1.3 (Complex affine varieties)

Let $I \subseteq \mathbb{C}[x_1, \dots, x_n]$ be an ideal. Then

$$\operatorname{Spec} \mathbb{C}[x_1, \dots, x_n]/I$$

contains a point for every closed irreducible subvariety of $\mathcal{V}(I)$. So in addition to the “geometric points” we have “generic points” along each of the varieties.

Example 56.1.4 (More examples of spectrums)

- (a) $\operatorname{Spec} \mathbb{Z}$ consists of a point for every prime p , plus a generic point that is somewhere, but no where in particular.
- (b) $\operatorname{Spec} \mathbb{C}[x]/(x^2)$ has only (x) as a prime ideal. The ideal (0) is not prime since $0 = x \cdot x$. Thus as a *topological space*, $\operatorname{Spec} \mathbb{C}[x]/(x^2)$ is a single point.
- (c) $\operatorname{Spec} \mathbb{Z}_{60}$ consists of three points. What are they?

§56.2 The Zariski topology of the spectrum

Prototypical example for this section: Still $\operatorname{Spec} \mathbb{C}[x_1, \dots, x_n]/I$.

Now, we endow a topology on $\operatorname{Spec} R$. Since the points on $\operatorname{Spec} R$ are the prime ideals, we continue the analogy by thinking of the points f as functions on $\operatorname{Spec} R$. That is,

Definition 56.2.1. Let $f \in R$ and $\mathfrak{p} \in \operatorname{Spec} R$. Then the **value** of f at $\mathfrak{p}$ is defined to be $f \pmod{\mathfrak{p}}$. We denote it $f(\mathfrak{p})$.

Example 56.2.2 (Vanishing loci in $\mathbb{A}^n$)

Suppose $R = \mathbb{C}[x_1, \dots, x_n]$, and $\mathfrak{p} = (x_1 - a_1, x_2 - a_2, \dots, x_n - a_n)$ is a maximal ideal of R . Then for a polynomial $f \in \mathbb{C}$,

$$f \pmod{\mathfrak{p}} = f(a_1, \dots, a_n)$$

with the identification that $\mathbb{C}/\mathfrak{p} \cong \mathbb{C}$.

Indeed if you replace R with $\mathbb{C}[x_1, \dots, x_n]$ and $\operatorname{Spec} R$ with $\mathbb{A}^n$ in everything that follows, then everything will be clear.

Definition 56.2.3. Let $f \in R$. We define the **vanishing locus** of f to be

$$\mathcal{V}(f) = \{\mathfrak{p} \in \operatorname{Spec} R \mid f(\mathfrak{p}) = 0\} = \{\mathfrak{p} \in \operatorname{Spec} R \mid f \in \mathfrak{p}\}.$$

More generally, just as in the affine case, we define the vanishing locus for an ideal I as

$$\begin{aligned} \mathcal{V}(I) &= \{\mathfrak{p} \in \operatorname{Spec} R \mid f(\mathfrak{p}) = 0 \forall f \in I\} \\ &= \{\mathfrak{p} \in \operatorname{Spec} R \mid f \in \mathfrak{p} \forall f \in I\} \\ &= \{\mathfrak{p} \in \operatorname{Spec} R \mid I \subseteq \mathfrak{p}\}. \end{aligned}$$

Finally, we define the **Zariski topology** on $\operatorname{Spec} R$ by declaring that the sets of the form $\mathcal{V}(I)$ are closed.

We now define a useful topological notion:

Definition 56.2.4. A point $p \in X$ is a **closed point** if the set $\{p\}$ is closed.

Question 56.2.5. Show that $\mathfrak{p} \in \operatorname{Spec} R$ is a closed point if and only if $\mathfrak{p}$ is a maximal ideal.

Therefore the Zariski topology lets us refer back to the old “geometric” as just the closed points.

Example 56.2.6 (Generic points, continued)

Let $R = \mathbb{C}[x, y]$ and let $\mathfrak{p} = (y - x^2) \in \operatorname{Spec} R$; this is the “generic point” on a parabola. It is not closed, but we can compute its closure:

$$\overline{\{\mathfrak{p}\}} = \mathcal{V}(\mathfrak{p}) = \{\mathfrak{q} \in \operatorname{Spec} R \mid \mathfrak{q} \supseteq \mathfrak{p}\}.$$

This closure contains the point $\mathfrak{p}$ as well as several maximal ideals $\mathfrak{q}$, such as $(x - 2, y - 4)$ and $(x - 3, y - 9)$. In other words, the closure of the “generic point” of the parabola is literally the set of all points that are actually on the parabola (including generic points).

That means the way to picture $\mathfrak{p}$ is a point that is “somewhere on the parabola”, but nowhere in particular. It makes sense then that if we take the closure, we get the entire parabola, since $\mathfrak{p}$ “could have been” any of those points.

The previous example illustrates an important observation:

Exercise 56.2.7. Let $I \subsetneq R$ be a proper ideal. Construct a bijection between the maximal ideals of R containing I , and the maximal ideals of R/I . ([Va15] labels this exercise as “Essential Algebra Exercise (Mandatory if you haven’t done it before)”.)

Example 56.2.8 (The generic point of the y -axis isn’t on the x -axis)

Let $R = \mathbb{C}[x, y]$ again. Consider $\mathcal{V}(y)$, which is the x -axis of $\operatorname{Spec} R$. Then consider $\mathfrak{p} = (x)$, which is the generic point on the y -axis. Observe that

$$\mathfrak{p} \notin \mathcal{V}(y).$$

The geometric way of saying this is that a *generic point* on the y -axis does not lie on the x -axis.

Finally, as before, we can define distinguished open sets, which form a basis of the Zariski topology.

Definition 56.2.9. Let $f \in \operatorname{Spec} R$. Then $D(f)$ is the set of $\mathfrak{p}$ such that $f(\mathfrak{p}) \neq 0$, a **distinguished open set**. These open sets form a basis for the Zariski topology on $\operatorname{Spec} R$.

§56.3 The structure sheaf

Prototypical example for this section: Still $\mathbb{C}[x_1, \dots, x_n]/I$.

We have now endowed $\text{Spec } R$ with the Zariski topology, and so all that remains is to put a sheaf $\mathcal{O}_{\text{Spec } R}$ on it. To do this we want a notion of “regular functions” as before.

In order for this to make sense, we have to talk about rational functions in a meaningful way. If R was an integral domain, then we could just use the field of fractions; for example if $R = \mathbb{C}[x_1, \dots, x_n]$ then we could just look at rational quotients of polynomials.

Unfortunately, in the general situation R may not be an integral domain: for example, the ring $R = \mathbb{C}[x]/(x^2)$ corresponding to the double point. So we will need to define something a little different: we will construct the *localization* of R at a set S , which we think of as the “set of allowed denominators”.

Definition 56.3.1. Let $S \subseteq R$, where R is a ring, and assume S is closed under multiplication. Then the **localization of R at S** , denoted $S^{-1}R$, is defined as the set of fractions

$$\{r/s \mid r \in R, s \in S\}$$

where we declare two fractions $r_1/s_1 = r_2/s_2$ to be equal if

$$\exists s \in S : s(r_1 s_2 - r_2 s_1) = 0.$$

In particular, if $0 \in S$ then $S^{-1}R$ is the trivial ring. So we usually only take situations where $0 \notin S$.

Question 56.3.2. Assume R is an integral domain and $S = R \setminus \{0\}$. Show that $S^{-1}R$ is just the field of fractions.

Example 56.3.3 (Why the extra s ?)

The reason we need the condition $s(r_1 s_2 - r_2 s_1) = 0$ rather than the simpler $r_1 s_2 - r_2 s_1 = 0$ is that otherwise the equivalence relation may fail to be transitive. Here is a counterexample: take

$$R = \mathbb{Z}_{12} \quad S = \{2, 4, 8\}.$$

Then we have for example

$$\frac{1}{2} = \frac{2}{4} = \frac{6}{4} = \frac{3}{2}.$$

So we need to have $\frac{1}{2} = \frac{3}{2}$ which is only true with the first definition.

Of course, if R is an integral domain (and $0 \notin S$) then this is a moot point.

The most important special case is the localization at a prime ideal.

Definition 56.3.4. Let R be a ring and $\mathfrak{p}$ a prime ideal. Then $R_{\mathfrak{p}}$ is defined to be $S^{-1}R$ for $S = R \setminus \mathfrak{p}$. We call this **localization at $\mathfrak{p}$** . Addition is defined in the obvious way.

Question 56.3.5. Why is S multiplicative closed in the above definition?

Thus,

If R is functions on the space $\text{Spec } R$, we think of $R_{\mathfrak{p}}$ as rational quotients f/g where $g(\mathfrak{p}) \neq 0$.

In particular, if $R = \mathbb{C}[x_1, \dots, x_n]$ then this is precisely the definition of rational function from before!

Now, we can define the sheaf as “locally rational” functions. This is done by a sheafification. First, let $\mathcal{F}$ be the pre-sheaf of “globally rational” functions: i.e. we define $\mathcal{F}(U)$ to be

$$\mathcal{F}(U) = \left\{ \frac{f}{g} \mid f, g \in R \text{ and } g(\mathfrak{p}) \neq 0 \forall \mathfrak{p} \in U \right\} = \left(R \setminus \bigcup_{\mathfrak{p} \in U} \mathfrak{p} \right)^{-1} R.$$

This is the localization of R to the functions vanishing outside U . For every $\mathfrak{p} \in U$ we can view f/g as an element in $R_{\mathfrak{p}}$ (since $g(\mathfrak{p}) \neq 0$). As one might expect this is an isomorphism

Lemma 56.3.6 (Stalks of the “globally rational” pre-sheaf)

The stalk of $\mathcal{F}$ defined above $\mathcal{F}_{\mathfrak{p}}$ is isomorphic to $R_{\mathfrak{p}}$.

Proof. There is an obvious map $\mathcal{F}_{\mathfrak{p}} \rightarrow R_{\mathfrak{p}}$ on germs by

$$(U, f/g \in \mathcal{F}(U)) \mapsto f/g \in R_{\mathfrak{p}}.$$

(Note the f/g on the left lives in $\mathcal{F}(U) = \left(R \setminus \bigcup_{\mathfrak{p} \in U} \mathfrak{p} \right)^{-1} R$ but the one on the right lives in $R_{\mathfrak{p}}$). Now suppose $(U_1, f_1/g_1)$ and $(U_2, f_2/g_2)$ are germs with $f_1/g_1 = f_2/g_2 \in R_{\mathfrak{p}}$.

Exercise 56.3.7. Now, show both germs are equal to $(U_1 \cap U_2 \cap D(h), f_1/g_1)$ with $D(h)$ the distinguished open set.

It is also surjective, since given $f/g \in R_{\mathfrak{p}}$ we take $U = D(g)$ the distinguished open set for g . $\square$

Then, we set

$$\mathcal{O}_{\text{Spec } R} = \mathcal{F}^{\text{sh}}.$$

In fact, we can even write out the definition of the sheafification, by viewing the germ at each point as an element of $R_{\mathfrak{p}}$.

Definition 56.3.8. Let R be a ring. Then $\text{Spec } R$ is made into a ringed space by setting

$$\mathcal{O}_{\text{Spec } R}(U) = \{(f_{\mathfrak{p}} \in R_{\mathfrak{p}})_{\mathfrak{p} \in U} \text{ which are locally } f/g\}.$$

That is, it consists of sequence $(f_{\mathfrak{p}})_{\mathfrak{p} \in U}$, with each $f_{\mathfrak{p}} \in R_{\mathfrak{p}}$, such that for every point $\mathfrak{p}$ there is a neighborhood $U_{\mathfrak{p}}$ and an $f, g \in R$ such that $f_{\mathfrak{q}} = \frac{f}{g} \in R_{\mathfrak{q}}$ for all $\mathfrak{q} \in U_{\mathfrak{p}}$.

§56.4 Example: fat points

So, let me elaborate a little on the “double point” scheme

$$X_2 = \text{Spec } \mathbb{C}[x]/(x^2)$$

since it is such an important motivating example. How it does differ from the “one-point” scheme $X_1 = \text{Spec } \mathbb{C}[x]/(x)$?

The difference can only be seen on the level of the structure sheaves. Indeed,

- As a set of points, X_2 has only one point, call it $p_2 = (x)$. Similarly, X_1 has only one point, call it $p_1 = (0)$.
- The two Zariski topologies are the same (at any rate there is only one topology on a one-point space).
- But at the structure sheaf level, X_2 has a “bigger” sheaf: the ring of functions on the single point p_2 is instead a *two-dimensional* $\mathbb{C}$ -vector space $\mathbb{C}[x]/(x^2)$. This formalizes the notion that this point is “fat”: specifying a function from p_2 to $\mathbb{C}$ now gives you *two* degrees of freedom instead of just one.

Another way to think about is in terms of functions. Consider polynomials $f = a_0 + a_1x + a_2x^2 + \dots$ on $\mathbb{C}[x]$. Then we have a sequence of maps

$$\mathbb{C}[x] \longrightarrow \mathbb{C}[x]/(x^2) \longrightarrow \mathbb{C}[x]/(x)$$

$$f \longmapsto a_0 + a_1x \longmapsto a_0$$

So p_1 only remembers the value of f , i.e. it remembers $f(0)$. But the point p_2 remembers more than just the value of f : it also remembers the first derivative of f . In [Va15] one draws a picture of this by taking $0 \in \mathbb{C}$ and adding a little bit of “infinitesimal fuzz” around it.

One can play the analogy more. There’s a “triple point” $X_3 = \text{Spec } \mathbb{C}[x]/(x^3) = \{p_3\}$ whose ring of functions has three degrees of freedom: specifying a “function” on p_3 gives you three degrees of freedom. Analogously, it remembers both the first and second derivatives of any polynomial in $\mathbb{C}[x]$. In [Va15], this is “the point 0 with even more fuzz”. Going even further, $\mathbb{C}[x, y]/(x^2, y)$ is “the origin with fuzz in the x -direction”, $\mathbb{C}[x, y]/(x, y)^2$ is “the origin with fuzz in all directions”, and so on and so forth.

§56.5 Properties of affine schemes

Now that we’re done defining an affine scheme, we state some important results about them.

Definition 56.5.1. For $g \in R$, we define the **localization of R at g** , denoted R_g to be $\{1, g, g^2, g^3, \dots\}^{-1}R$. (Note that $\{1, g, g^2, \dots\}$ is multiplicatively closed.)

This is admittedly somewhat terrible notation, since $R_{\mathfrak{p}}$ is the localization with multiplicative set $R \setminus \mathfrak{p}$, while R_g is the localization with multiplicative set $\{1, g, g^2, \dots\}$; these two are quite different beasts!

Example 56.5.2 (Localization at an element)

Let $R = \mathbb{C}[x, y, z]$ and let $g = x$. Then

$$R_g = \left\{ \frac{P(x, y, z)}{x^n} \mid P \in \mathbb{C}[x, y, z], n \geq 0 \right\}.$$

Theorem 56.5.3 (On the affine structure sheaf)

Let R be a ring and $\operatorname{Spec} R$ the associated affine scheme.

- (a) Let $\mathfrak{p}$ be a prime ideal. Then $\mathcal{O}_{\operatorname{Spec} R, \mathfrak{p}} \cong R_{\mathfrak{p}}$.
- (b) Let $D(g)$ be a distinguished open set. Then $\mathcal{O}_{\operatorname{Spec} R}(D(g)) \cong R_g$.

This matches the results that we've seen when $\operatorname{Spec} R$ is an affine variety.

Proof. Part (a) follows by Lemma 56.3.6 and Lemma 55.5.5. For part (b), we need:

Question 56.5.4. If I and J are ideals of a ring R , then $\mathcal{V}(I) \subseteq \mathcal{V}(J)$ if and only if $\sqrt{J} \subseteq \sqrt{I}$. (Use the fact that $\sqrt{I} = \cap_{I \subseteq \mathfrak{p}} \mathfrak{p}$.)

Then we can repeat the proof of Theorem 51.6.1. □

Example 56.5.5 (Examples of structure sheaves)

- (a) Let $X = \operatorname{Spec} \mathbb{C}[x]$. Then $\mathcal{O}_X(U)$ is the set of regular functions on U in the sense that we've seen before.
- (b) Let $X = \operatorname{Spec} \mathbb{Z}$. Let $U = X \setminus \{(7), (11)\}$. Then $\mathcal{O}_X(U)$ can be identified with the set of rational numbers which have denominator of the form $7^x 11^y$.

We now state the most important theorem. In fact, this theorem is the justification that our definition of a scheme is the correct one.

Theorem 56.5.6 (Affine schemes and commutative rings are the same category)

Let R and S be rings. There is a natural bijection between maps of schemes $\operatorname{Spec} R \rightarrow \operatorname{Spec} S$ and ring homomorphisms $\psi : S \rightarrow R$.

Proof. First, we need to do the construction: given a map of schemes $f : \operatorname{Spec} R \rightarrow \operatorname{Spec} S$ we need to construct a homomorphism $S \rightarrow R$. This should be the easy part, because f has a lot of data. Indeed, recall that for every $U \subseteq \operatorname{Spec} S$ there is supposed to be a map

$$f_U^* : \mathcal{O}_{\operatorname{Spec} S}(U) \rightarrow \mathcal{O}_{\operatorname{Spec} R}(f^{\text{pre}}(U)).$$

If we take U to be the entire space $\operatorname{Spec} S$ we now have a ring homomorphism

$$S = \mathcal{O}_{\operatorname{Spec} S}(\operatorname{Spec} S) \rightarrow \mathcal{O}_{\operatorname{Spec} R}(\operatorname{Spec} R) = R.$$

which is one part of the construction.

The more involved direction is building from a homomorphism $\psi : S \rightarrow R$ a map $f : \operatorname{Spec} R \rightarrow \operatorname{Spec} S$. We define it by:

- On points, $f(\mathfrak{p}) = \psi^{\text{pre}}(\mathfrak{p})$ for each prime ideal $\mathfrak{p}$. One can check $\psi^{\text{pre}}(\mathfrak{p})$ is indeed prime and that f is continuous in the Zariski topology.

- We first construct a map on the stalks of the sheaf as follows: for every prime ideal $\mathfrak{q}$ in R , let $\psi^\#(\mathfrak{q}) = \mathfrak{p}$ be its image (so that $f(\mathfrak{q}) = \mathfrak{p}$) and consider the map

$$\psi_{\mathfrak{q}} : \mathcal{O}_{\text{Spec } S, \mathfrak{q}} = S_{\mathfrak{q}} \rightarrow R_{\psi^\#(\mathfrak{q})} = \mathcal{O}_{\text{Spec } R, \psi^\#(\mathfrak{q})} \quad \text{by} \quad f/g \mapsto \psi(f)/\psi(g).$$

Again, one can check this is well-defined and a map of local rings. This is called the localization of ψ at $\mathfrak{q}$.

- Finally, we use the above map to construct an $f_U^* : \mathcal{O}_{\text{Spec } S}(U) \rightarrow \mathcal{O}_{\text{Spec } R}(f^{\text{pre}}(U))$ for every open set $U \subseteq \text{Spec } S$. Since $\mathcal{O}_{\text{Spec } S}$ is a sheafification, we think of its sections as sequences of compatible germs $(g_{\mathfrak{q}})_{\mathfrak{q} \in U}$ and then map it via the map $\psi_{\mathfrak{q}}$ above.

One then has to check that everything is well-defined. This is left as an exercise to the diligent reader; for an actual proof see [Va15, Proposition 6.3.2]. $\square$

Remark 56.5.7. Categorically, this says that the “global section functor” $\text{Spec } R \mapsto \mathcal{O}_{\text{Spec } R}(R) = R$ is a fully faithful functor from the category AffSch to the category CRing . (It is also essentially surjective on objects.)

To be more philosophical,

The category of affine schemes and the category of commutative rings are exactly the same, down to the morphisms between two pairs of objects.

§56.6 Schemes

Definition 56.6.1. A **scheme** is a locally ringed space $(X, \mathcal{O}_X)$ with an open cover $\{U_\alpha\}$ of X such that each pair $(U_\alpha, \mathcal{O}_X|_{U_\alpha})$ is isomorphic to an affine scheme.

Hooray!

§56.7 Projective scheme

Prototypical example for this section: Projective varieties, in the same way.

The most important class of schemes which are not affine are *projective* schemes. The complete the obvious analogy:

$$\frac{\text{Affine variety}}{\text{Projective variety}} = \frac{\text{Affine scheme}}{\text{Projective scheme}}.$$

Let S be *any* (commutative) graded ring, like $\mathbb{C}[x_0, \dots, x_n]$.

Definition 56.7.1. We define $\text{Proj } S$, the **projective scheme over S** :

- As a set, $\text{Proj } S$ consists of *homogeneous prime ideals* $\mathfrak{p}$ which do not contain S^+ .
- If $I \subseteq S$ is homogeneous, then we let $\mathcal{V}_{\text{pr}}(I) = \{\mathfrak{p} \in \text{Proj } S \mid I \subseteq \mathfrak{p}\}$. Then the **Zariski topology** is imposed by declaring sets of the form $\mathcal{V}_{\text{pr}}(I)$ to be closed.

- We now define a pre-sheaf $\mathcal{F}$ on $\text{Proj } S$ by

$$\mathcal{F}(U) = \left\{ \frac{f}{g} \mid g(\mathfrak{p}) \neq 0 \ \forall \mathfrak{p} \in U \text{ and } \deg f = \deg g \right\}.$$

In other words, the rational functions are quotients f/g where f and g are *homogeneous of the same degree*. Then we let

$$\mathcal{O}_{\text{Proj } S} = \mathcal{F}^{\text{sh}}$$

be the sheafification.

Definition 56.7.2. The **distinguished open sets** $D(f)$ of the $\text{Proj } S$ are defined as $\{\mathfrak{p} \in \text{Proj } S : f(\mathfrak{p}) \neq 0\}$, as before; these form a basis for the Zariski topology of $\text{Proj } S$.

Now, we want analogous results as we did for affine structure sheaf. So, we define a slightly modified localization:

Definition 56.7.3. Let S be a graded ring.

- (i) For a prime ideal $\mathfrak{p}$, let

$$S_{(\mathfrak{p})} = \left\{ \frac{f}{g} \mid g(\mathfrak{p}) \neq 0 \text{ and } \deg f = \deg g \right\}$$

denote the elements of $S_{(\mathfrak{p})}$ with “degree zero”.

- (ii) For any homogeneous $g \in S$ of degree d , let

$$S_{(g)} = \left\{ \frac{f}{g^r} \mid \deg f = r \deg g \right\}$$

denote the elements of S_g with “degree zero”.

Theorem 56.7.4 (On the projective structure sheaf)

Let S be a graded ring and let $\text{Proj } S$ the associated projective scheme.

- (a) Let $\mathfrak{p} \in \text{Proj } S$. Then $\mathcal{O}_{\text{Proj } S, \mathfrak{p}} \cong S_{(\mathfrak{p})}$.
 (b) Suppose g is homogeneous with $\deg g > 0$. Then

$$D(g) \cong \text{Spec } S_{(g)}$$

as locally ringed spaces. In particular, $\mathcal{O}_{\text{Proj } S}(D(g)) \cong S_{(g)}$.

Question 56.7.5. Conclude that $\text{Proj } S$ is a scheme.

Of course, the archetypal example is that

$$\text{Proj } \mathbb{C}[x_0, x_1, \dots, x_n]/I$$

corresponds to the projective subvariety of $\mathbb{CP}^n$ cut out by I (when I is radical). In the general case of an arbitrary ideal I , we call such schemes **projective subscheme** of $\mathbb{CP}^n$. For example, the “double point” in $\mathbb{CP}^1$ is given by $\text{Proj}[x_0, x_1]/(x_0^2)$.

Remark 56.7.6. No comment yet on what the global sections of $\mathcal{O}_{\text{Proj } S}(\text{Proj } S)$ are. (The theorem above requires $\deg g > 0$, so we cannot just take $g = 1$.) One might hope that in general $\mathcal{O}_{\text{Proj } S}(\text{Proj } S) \cong S^0$ in analogy to our complex projective varieties, but one needs some additional assumptions on S for this to hold.

§56.8 Where to go from here

This chapter concludes the long setup for the definition of a scheme. Unfortunately, with respect to algebraic geometry this is as much as I have the patience or knowledge to write about. So, if you want to actually see how schemes are used in “real life”, you’ll have to turn elsewhere.

A good reference I happily recommend is [Ga03]. More intense is [Va15]. See Appendix A for further remarks.

§56.9 Problems to think about

Problem 56A. Describe the points of $\operatorname{Spec} \mathbb{R}[x, y]$.

Problem 56B[†] (Chinese remainder theorem). Consider $X = \operatorname{Spec} \mathbb{Z}_{60}$, which as a topological space has three points. By considering $\mathcal{O}_X(X)$ prove the Chinese theorem

$$\mathbb{Z}_{60} \cong \mathbb{Z}_4 \times \mathbb{Z}_3 \times \mathbb{Z}_5.$$

Problem 56C. Given an affine scheme $X = \operatorname{Spec} R$, show that there is a unique morphism of schemes $X \rightarrow \operatorname{Spec} \mathbb{Z}$, and describe where it sends points of X .



Problem 56D (From Andrew Critch). Let R be a Noetherian ring. Show that R is an integral domain if and only if it has no idempotents, and $R_{\mathfrak{p}}$ is an integral domain for every prime $\mathfrak{p}$.

XV

Set Theory I: ZFC, Ordinals, and Cardinals

57 Bonus: Cauchy's functional equation and Zorn's lemma	519
57.1 Let's construct a monster	519
57.2 Review of finite induction	520
57.3 Transfinite induction	520
57.4 Wrapping up functional equations	523
57.5 Zorn's lemma	524
58 Zermelo-Fraenkel with choice	526
58.1 The ultimate functional equation	526
58.2 Cantor's paradox	526
58.3 The language of set theory	527
58.4 The axioms of ZFC	528
58.5 Encoding	530
58.6 Choice and well-ordering	530
58.7 Sets vs classes	531
58.8 Problems to think about	532
59 Ordinals	533
59.1 Counting for preschoolers	533
59.2 Counting for set theorists	534
59.3 Definition of an ordinal	536
59.4 Ordinals are "tall"	537
59.5 Transfinite induction and recursion	538
59.6 Ordinal arithmetic	539
59.7 The hierarchy of sets	540
59.8 Problems to think about	542
60 Cardinals	543
60.1 Equinumerous sets and cardinals	543
60.2 Cardinalities	544
60.3 Aleph numbers	544
60.4 Cardinal arithmetic	545
60.5 Cardinal exponentiation	547
60.6 Cofinality	547
60.7 Inaccessible cardinals	549
60.8 Problems to think about	549

57


Bonus: Cauchy's functional equation and Zorn's lemma

This is an informal chapter on Zorn's lemma, which will give an overview of what's going to come in the last parts of the Napkin. It may be omitted, as all important terms used here will be defined in full later.

In the world of olympiad math, there's a famous functional equation that goes as follows:

$$f : \mathbb{R} \rightarrow \mathbb{R} \quad f(x + y) = f(x) + f(y).$$

Everyone knows what its solutions are! There's an obvious family of solutions $f(x) = cx$. Then there's also this family of... uh... noncontinuous solutions (mumble grumble) pathological (mumble mumble) Axiom of Choice (grumble).



Don't worry, I know what I'm doing!
Image from [Mu]

There's also this thing called Zorn's lemma. It sounds terrifying, because it's equivalent to the Axiom of Choice, which is also terrifying because why not.

In this post I will try to de-terrify these things, because they're really not terrifying and I'm not sure why no one bothered to explain this properly yet. I have yet to see an olympiad handout that explains how you would construct a pathological solution, even though it's really quite natural. So let me fix this problem now...

§57.1 Let's construct a monster

Let us just see if we can try and construct a "bad" f and see what happens.

By scaling, let's assume WLOG that $f(1) = 1$. Thus $f(n) = n$ for every integer n , and you can easily show from here that

$$f\left(\frac{m}{n}\right) = \frac{m}{n}.$$

So f is determined for all rationals. And then you get stuck.

None of this is useful for determining, say, $f(\sqrt{2})$. You could add and subtract rational numbers all day and, say, $\sqrt{2}$ isn't going to show up at all.

Well, we're trying to set things on fire anyways, so let's set

$$f(\sqrt{2}) = 2015$$

because why not? By the same induction, we get $f(n\sqrt{2}) = 2015n$, and then that

$$f(a + b\sqrt{2}) = a + 2015b.$$

Here a and b are rationals. Well, so far so good – as written, this is a perfectly good solution, other than the fact that we’ve only defined f on a tiny portion of the real numbers.

Well, we can do this all day:

$$f(a + b\sqrt{2} + c\sqrt{3} + d\pi) = a + 2015b + 1337c - 999d.$$

Perfectly consistent.

You can kind of see how we should keep going now. Just keep throwing in new real numbers which are “independent” to the previous few, assigning them to whatever junk we want. It feels like it *should* be workable. . .

In a moment I’ll explain what “independent” means (though you might be able to guess already), but at the moment there’s a bigger issue: no matter how many numbers we throw, it seems like we’ll never finish. Let’s address the second issue first.

§57.2 Review of finite induction

When you do induction, you get to count off 1, 2, 3, ... and so on. So for example, suppose we had a “problem” such as:

Prove that the intersection of n open intervals is either $\emptyset$ or an open interval.

You can do this by induction easily: it’s true for $n = 2$, and for the larger cases it’s similarly easy.

But you can’t conclude from this that *infinitely* many open intervals intersect at some open interval. Indeed, this is false: consider the intervals

$$(-1, 1), \quad \left(-\frac{1}{2}, \frac{1}{2}\right), \quad \left(-\frac{1}{3}, \frac{1}{3}\right), \quad \left(-\frac{1}{4}, \frac{1}{4}\right), \quad \dots$$

This *infinite* set of intervals intersects at a single point $\{0\}$!

The moral of the story is that induction doesn’t let us reach infinity. Too bad, because we’d have loved to use induction to help us construct a monster. That’s what we’re doing, after all – adding things in one by one.

§57.3 Transfinite induction

Well, it turns out we can, but we need a new notion of number.

For this we need a notion of an *ordinal number*. I defined these in their full glory a previous blog post, but I don’t need the full details of that. Here’s what I want to say: after all the natural numbers

$$0, 1, \dots,$$

I’ll put a *new number* called ω , representing how large the natural numbers are. After that there’s a number called

$$\omega + 1, \omega + 2, \dots$$

and eventually a number called 2ω .

The list goes on:

$$\begin{aligned}
 &0, 1, 2, 3, \dots, \omega \\
 &\omega + 1, \omega + 2, \dots, \omega + \omega \\
 &2\omega + 1, 2\omega + 2, \dots, 3\omega \\
 &\vdots \\
 &\omega^2 + 1, \omega^2 + 2, \dots \\
 &\vdots \\
 &\omega^3, \dots, \omega^4, \dots, \omega^\omega, \dots, \omega^{\omega^{\omega^{\dots}}}
 \end{aligned}$$

Pictorially, it kind of looks like this:




Image from [ca]

Anyways, in the same way that natural numbers dominate all finite sets, the ordinals dominate *all the sets*.

Theorem 57.3.1

Assume the Axiom of Choice. Then for every set S there's some ordinal α which is bigger than it.

But it turns out (and you can intuitively see) that as large as the ordinals grow, there is no *infinite descending chain*. Meaning: if I start at an ordinal (like $2\omega + 4$) and jump down, I can only take finitely many jumps before I hit 0. (To see this, try writing down a chain starting at $2\omega + 4$ yourself.) Hence, induction and recursion still work verbatim:

Theorem 57.3.2 (Transfinite induction)

Given a statement $P(-)$, suppose that

- $P(0)$ is true, and
- If $P(\alpha)$ is true for all $\alpha < \beta$, then $P(\beta)$ is true.

Then $P(\beta)$ is true.

Similarly, you're allowed to do recursion to define x_β if you know the value of x_α for all $\alpha < \beta$.

The difference from normal induction or recursion is that we'll often only do things like "define $x_{n+1} = \dots$ ". But this is not enough to define x_α for all α . To see this, try using our normal induction and see how far we can climb up the ladder.

Answer: you can't get ω ! It's not of the form $n + 1$ for any of our natural numbers n — our finite induction only lets us get up to the ordinals less than ω . Similarly, the simple $+1$ doesn't let us hit the ordinal 2ω , even if we already have $\omega + n$ for all n . Such ordinals are called **limit ordinals**. The ordinal that *are* of the form $\alpha + 1$ are called **successor ordinals**.

So a transfinite induction or recursion is very often broken up into three cases. In the induction phrasing, it looks like

- (Zero Case) First, resolve $P(0)$.
- (Successor Case) Show that from $P(\alpha)$ we can get $P(\alpha + 1)$.
- (Limit Case) Show that $P(\lambda)$ holds given $P(\alpha)$ for all $\alpha < \lambda$, where λ is a limit ordinal.

Similarly, transfinite recursion often is split into cases too.

- (Zero Case) First, define x_0 .
- (Successor Case) Define $x_{\alpha+1}$ from x_α .
- (Limit Case) Define x_λ from x_α for all $\alpha < \lambda$, where λ is a limit ordinal.

In both situations, finite induction only does the first two cases, but if we're able to do the third case we can climb far above the barrier ω .

§57.4 Wrapping up functional equations

Let's return to solving our problem.

Let S_n denote the set of “base” numbers we have at the n th step. In our example, we might have

$$S_1 = \{1\}, \quad S_2 = \{1, \sqrt{2}\}, \quad S_3 = \{1, \sqrt{2}, \sqrt{3}\}, \quad S_4 = \{1, \sqrt{2}, \sqrt{3}, \pi\}, \quad \dots$$

and we'd like to keep building up S_i until we can express all real numbers. For completeness, let me declare $S_0 = \emptyset$.

First, I need to be more precise about “independent”. Intuitively, this construction is working because

$$a + b\sqrt{2} + c\sqrt{3} + d\pi$$

is never going to equal zero for rational numbers a, b, c, d (other than all zeros). In general, a set X of numbers is “independent” if the combination

$$c_1x_1 + c_2x_2 + \dots + c_mx_m = 0$$

never occurs for rational numbers $\mathbb{Q}$ unless $c_1 = c_2 = \dots = c_m = 0$. Here $x_i \in X$ are distinct. Note that even if X is infinite, I can only take finite sums! (This notion has a name: we want X to be **linearly independent** over $\mathbb{Q}$; see the chapter on vector spaces for more on this!)

When do we stop? We'd like to stop when we have a set $S_{\text{something}}$ that's so big, every real number can be written in terms of the independent numbers. (This notion also has a name: it's called a $\mathbb{Q}$ -basis.) Let's call such a set **spanning**; we stop once we hit a spanning set.

The idea that we can induct still seems okay: suppose S_α isn't spanning. Then there's some number that is independent of S_α , say $\sqrt{2015}\pi$ or something. Then we just add it to get $S_{\alpha+1}$. And we keep going.

Unfortunately, as I said before it's not enough to be able to go from S_α to $S_{\alpha+1}$ (successor case); we need to handle the limit case as well. But it turns out there's a trick we can do. Suppose we've constructed *all* the sets $S_0, S_1, S_2, \dots$, one for each positive integer n , and none of them are spanning. The next thing I want to construct is S_ω ; somehow I have to “jump”. To do this, I now take the infinite union

$$S_\omega \stackrel{\text{def}}{=} S_0 \cup S_1 \cup S_2 \cup \dots$$

The elements of this set are also independent (why?).

Ta-da! With the simple trick of “union all the existing sets”, we've just jumped the hurdle to the first limit ordinal ω . Then we can construct $S_{\omega+1}, S_{\omega+2}, \dots$, once again – just keep throwing in elements. Then when we need to jump the next hurdle to $S_{2\omega}$, we just do the same trick of “union-ing” all the previous sets.

So we can formalize the process as follows:

1. Let $S_0 = \emptyset$.
2. For a successor stage $S_{\alpha+1}$, add any element to S_α to obtain $S_{\alpha+1}$.
3. For a limit stage S_λ , take the union $\bigcup_{\gamma < \lambda} S_\gamma$.

How do we know that we'll stop eventually? Well, the thing is that this process consumes a lot of real numbers. In particular, the ordinals get larger than the size of $\mathbb{R}$ (assuming

Choice). Hence if we don't stop we will quite literally reach a point where we have used up every single real number. Clearly that's impossible, because by then the elements can't possibly be independent!

So by transfinite recursion, we eventually hit some S_γ which is spanning: the elements are all independent, but every real number can be expressed using it. Done!

§57.5 Zorn's lemma

Now I can tell you what Zorn's lemma is: it lets us do the same thing in any poset.

We can think of the above example as follows: consider all sets of independent elements. These form a partially ordered set by inclusion, and what we did was quite literally climb up a chain

$$S_0 \subsetneq S_1 \subsetneq S_2 \subsetneq \dots$$

It's not quite climbing since we weren't just going one step at a time: we had to do "jumps" to get up to S_ω and resume climbing. But the main idea is to climb up a poset until we're at the very top; in the previous case, when we reached the spanning set.

The same thing works verbatim with any partially ordered set $\mathcal{P}$. Let's define some terminology. A **local maximum** of the entire poset $\mathcal{P}$ is an element which has no other elements strictly greater than it. (Most authors refer to this as "maximal element", but I think "local maximum" is a more accurate term.)

Now a **chain of length γ** is a set of elements p_α for every $\alpha < \gamma$ such that $p_0 < p_1 < p_2 < \dots$ (Observe that a chain has a last element if and only if γ is a successor ordinal, like $\omega + 3$.) An **upper bound** to a chain is an element $\tilde{p}$ which is greater than or equal to all elements of the chain; In particular, if γ is a successor ordinal, then just taking the last element of the chain works.

In this language, Zorn's lemma states that

Theorem 57.5.1 (Zorn's lemma)

Let $\mathcal{P}$ be a nonempty partially ordered set. If every chain has an upper bound, then $\mathcal{P}$ has a local maximum.

Chains with length equal to a successor ordinal always have upper bounds, but this is not true in the limit case. So the hypothesis of Zorn's lemma is exactly what lets us "jump" up to define p_ω and other limit ordinals. And the proof of Zorn's lemma is straightforward: keep climbing up the poset at successor stages, using Zorn's condition to jump up at limit stages, and thus building a really long chain. But we have to eventually stop, or we literally run out of elements of $\mathcal{P}$. And the only possible stopping point is a local maximum.

If we want to phrase our previous solution in terms of Zorn's lemma, we'd say:

Proof. Look at the poset whose elements are sets of independent real numbers. Every chain $S_0 \subsetneq S_1 \subsetneq \dots$ has an upper bound $\bigcup S_\alpha$ (which you have to check is actually an element of the poset). Thus by Zorn, there is a local maximum S . Then S must be spanning, because otherwise we could add an element to it. $\square$

So really, Zorn's lemma is encoding all of the work of climbing that I argued earlier. It's a neat little package that captures all the boilerplate, and tells you exactly what you need to check.




Image from [Go09]

One last thing you might ask: where is the Axiom of Choice used? Well, the idea is that for any chain there could be lots of $\tilde{p}$'s, and you need to pick one of them. Since you are making arbitrary choices infinitely many times, you need the Axiom of Choice. (Actually, you also need choice to talk about cardinalities as in theorem 1.) But really, it's nothing special.

58 Zermelo-Fraenkel with choice

Chapter 2 $\frac{1}{2}$ of [Le14] has a nice description of this.

§58.1 The ultimate functional equation

In abstract mathematics, we often define structures by what *properties* they should have; for example, a group is a set and a binary operation satisfying so-and-so axioms, while a metric space is a set and a distance function satisfying so-and-so axioms.

Nevertheless, these definitions rely on previous definitions. The colorful illustration of [Le14] on this:

- A *vector space* is an abelian group with...
- An *abelian group* has a binary operation such that...
- A *binary operation* on a set is...
- A *set* is...

and so on.

We have to stop at some point, because infinite lists of definitions are bad. The stopping turns out to be a set, “defined” by properties. The trick is that we never actually define what a set is, but nonetheless postulate that these sets satisfy certain properties: these are the ZFC axioms. Loosely, ZFC can be thought of as the *ultimate functional equation*.

Before talking about what these axioms are, I should talk about the caveats.

§58.2 Cantor’s paradox

Intuitively, a set is an unordered collection of elements. Two sets are equal if they share the same elements:

$$\{x \mid x \text{ is a featherless biped}\} = \{x \mid x \text{ is human}\}$$

(let’s put aside the issue of dinosaurs).

As another example, we have our empty set $\emptyset$ that contains no objects. We can have a set $\{1, 2, 3\}$, or maybe the set of natural numbers $\mathbb{N} = \{0, 1, 2, \dots\}$. (For the purposes of set theory, 0 is usually considered a natural number.) Sets can even contain other sets, like $\{\mathbb{Z}, \mathbb{Q}, \mathbb{N}\}$. Fine and dandy, right?

The trouble is that this definition actually isn’t good enough, and here’s why. If we just say “a set is any collection of objects”, then we can consider a really big set V , the set of all sets. So far no problem, right? V has the oddity that it has to contain itself $V \in V$, but oh well, no big deal.

Unfortunately, this existence of this V leads immediately to a paradox. The classical one is Bertrand’s Paradox. I will instead present a somewhat simpler one: not only does V contain itself, *every subset* $S \subseteq V$ is itself an element of V (i.e. $S \in V$). If we let $\mathcal{P}(V)$ denote the **power set** of V (i.e. all the subsets of V), then we have an inclusion

$$\mathcal{P}(V) \hookrightarrow V.$$

This is bad, since:

Lemma 58.2.1 (Cantor's diagonal argument)

For *any* set X , it's impossible to construct an injective map $\iota : \mathcal{P}(X) \hookrightarrow X$.

Proof. Assume for contradiction ι exists.

Exercise 58.2.2. Show that there exists a surjective map $j : X \rightarrow \mathcal{P}(X)$. (This is easier than it appears, just “invert ι ”).

We now claim that j can't exist. The idea is contained inside the picture:

		x_1	x_2	x_3	x_4	x_5	$\dots$
x_1	$\mapsto$	0	1	1	0	1	$\dots$
x_2	$\mapsto$	1	1	0	1	1	$\dots$
x_3	$\mapsto$	0	1	0	0	1	$\dots$
x_4	$\mapsto$	1	0	0	1	0	$\dots$
x_5	$\mapsto$	0	1	1	1	1	$\dots$
$\vdots$		$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$

here for each $j(x) \subseteq X$, I'm writing “1” to mean that the element is inside $j(x)$, and “0” otherwise. So $j(x_1) = \{x_2, x_3, x_5, \dots\}$. (Here the indices are ordinals rather than integers as X may be uncountable.)

Then we can read off the diagonal to get a new set. In our example, the diagonal specifies a set $A = \{x_2, x_4, x_5, \dots\}$. Then we invert it to get a set $B = \{x_1, x_3, \dots\}$. In symbols, we consider the set

$$B = \{x \mid x \notin j(x)\}$$

By construction, $B \subseteq X$ is not in the image of j , which is a contradiction since j was supposed to be surjective. $\square$

Now if you're not a set theorist, you could probably just brush this off, saying “oh well, I guess you can't look at certain sets”. But if you're a set theorist, this worries you, because you realize it means that you can't just define a set as “a collection of objects”, because then everything would blow up. Something more is necessary.

§58.3 The language of set theory

We need a way to refer to sets other than the informal description of “collection of objects”.

So here's what we're going to do. We'll start by defining a formal *language of set theory*, a way of writing logical statements. First of all we can throw in our usual logical operators:

- $\forall$ means “for all”
- $\exists$ means “exists”
- $=$ means “equal”
- $X \implies Y$ means “if X then Y ”
- $A \wedge B$ means “ A and B ”

- $A \vee B$ means “ A or B ”
- $\neg A$ means “not A ”.

Since we’re doing set theory, there’s only one more operator we add in: the inclusion $\in$. And that’s all we’re going to use (for now).

So how do we express something like “the set $\{1, 2\}$ ”? The trick is that we’re not going to actually “construct” any sets, but rather refer to them indirectly, like so:

$$\exists S : x \in S \iff ((x = 1) \vee (x = 2)).$$

This reads: “there exists an S such that x is in S if and only if either 1 is in S or 2 is in S ”. We don’t have to refer to sets as objects in and of themselves anymore – we now have a way to “create” our sets, by writing formulas for exactly what they contain. This is something a machine can parse.

Well, what are we going to do with things like 1 and 2, which are not sets? Answer:

Elements of sets are themselves sets.

We’re going to make **everything** into a set. Natural numbers will be sets. Ordered pairs will be sets. Functions will be sets. Later, I’ll tell you exactly how we manage to do something like encode 1 as a set. For now, all you need to know is that that sets don’t just hold objects; they hold other sets.

So now it makes sense to talk about whether something is a set or not: $\exists x$ means “ x is a set”, while $\nexists x$ means “ x is not a set”. In other words, we’ve rephrased the problem of deciding whether something is a set to whether it exists, which makes it easier to deal with in our formal language. That means that our axiom system had better find some way to let us show a lot of things exist, without letting us prove

$$\exists S \forall x : x \in S.$$

For if we prove this formula, then we have our “bad” set that caused us to go down the rabbit hole in the first place.

§58.4 The axioms of ZFC

I don’t especially want to get into details about these axioms; if you’re interested, read:

- <https://usamo.wordpress.com/2014/11/13/set-theory-an-intro-to-zfc-part-1/>
- <https://usamo.wordpress.com/2014/11/18/set-theory-part-2-constructing-the-ordinals/>

Here is a much terser description of the axioms, which also includes the corresponding sentence in the language of set theory. It is worth the time to get some practice parsing $\forall$, $\exists$, etc. and you can do so by comparing the formal sentences with the natural statement of the axiom.

First, the two easiest axioms:

- Extensionality is the sentence $\forall x \forall y ((\forall a (a \in x \iff a \in y)) \implies x = y)$, which says that if two sets x and y have the same elements, then $x = y$.
- EmptySet is the sentence $\exists a : \forall x \neg(x \in a)$; it says there exists a set with no elements. By Extensionality this set is unique, so we denote it $\emptyset$.

The next two axioms give us basic ways of building new sets.

- Given two elements x and y , there exists a set a containing only those two elements. In machine code, this is the sentence Pairing, written

$$\forall x \forall y \exists a \quad \forall z, z \in a \iff ((z = x) \vee (z = y)).$$

By Extensionality this set a is unique, so we write $a = \{x, y\}$.

- Given a set a , we can create the union of the elements of a . For example, if $a = \{\{1, 2\}, \{3, 4\}\}$, then $U = \{1, 2, 3, 4\}$ is a set. Formally, this is the sentence Union:

$$\forall a \exists U \quad \forall x [(x \in U) \iff (\exists y : x \in y \in a)].$$

Since U is unique by Extensionality, we denote it $\cup a$.

- We can construct the **power set** $\mathcal{P}(x)$. Formally, the sentence PowerSet says that

$$\forall x \exists P \forall y (y \in P \iff y \subseteq x)$$

where $y \subseteq x$ is short for $\forall z (z \in y \implies z \in x)$. As Extensionality gives us uniqueness of P , we denote it $\mathcal{P}(x)$.

- Foundation says there are no infinite descending chains

$$x_0 \ni x_1 \ni x_2 \ni \dots$$

This is important, because it lets us induct. In particular, **no set contains itself**.

- Infinity implies that $\omega = \{0, 1, \dots\}$ is a set.

These are all things you are already used to, so keep your intuition there. The next one is less intuitive:

- The **schema of restricted comprehension** says: if we are *given a set* X , and some formula $\phi(x)$ then we can *filter* through the elements of X to get a subset

$$Y = \{x \in X \mid \phi(x)\}.$$

Formally, given a formula ϕ :

$$\forall X \quad \exists Y \quad \forall y (y \in Y \iff y \in X \wedge \phi(y)).$$

Notice that we may *only* do this filtering over an already given set. So it is not valid to create $\{x \mid x \text{ is a set}\}$. We are thankful for this, because this lets us evade Cantor's paradox.

Abuse of Notation 58.4.1. Note that technically, there are infinitely many sentences, a $\text{Comprehension}_\phi$ for every possible formula ϕ . By abuse of notation, we let Comprehension abbreviate the infinitely many axioms $\text{Comprehension}_\phi$ for every ϕ .

There is one last schema called Replacement_ϕ . Suppose X is a set and $\phi(x, y)$ is some formula such that for every $x \in X$, there is a *unique* y in the universe such that $\phi(x, y)$ is true: for example “ $y = x \cup \{x\}$ ” works. (In effect, ϕ is defining a function f on X .) Then there exists a set Y consisting exactly of these images: (i.e. $f[X]$ is a set).

Abuse of Notation 58.4.2. By abuse of notation, we let Replacement abbreviate the infinitely many axioms Replacement_ϕ for every ϕ .

We postpone discussion of the Axiom of Choice momentarily.

§58.5 Encoding

Now that we have this rickety universe of sets, we can start re-building math. You'll get to see this more in the next chapter on ordinal numbers.

Definition 58.5.1. An **ordered pair** (x, y) is a set of the form

$$(x, y) \stackrel{\text{def}}{=} \{\{x\}, \{x, y\}\}.$$

Note that $(x, y) = (a, b)$ if and only if $x = a$ and $y = b$. Ordered k -tuples can be defined recursively: a three-tuple (a, b, c) means $(a, (b, c))$.

Definition 58.5.2. A **function** $f : X \rightarrow Y$ is defined as a collection of ordered pairs such that

- If $(x, y) \in f$, then $x \in X$ and $y \in Y$.
- For every $x \in X$, there is a unique $y \in Y$ such that $(x, y) \in f$. We denote this y by $f(x)$.

Definition 58.5.3. The **natural numbers** are defined inductively as

$$\begin{aligned} 0 &= \emptyset \\ 1 &= \{0\} \\ 2 &= \{0, 1\} \\ 3 &= \{0, 1, 2\} \\ &\vdots \end{aligned}$$

The set of all natural numbers is denoted ω .

Abuse of Notation 58.5.4. Yes, I'm sorry, in set theory 0 is considered a natural number. For this reason I'm using ω and not $\mathbb{N}$ since I explicitly have $0 \notin \mathbb{N}$ in all other parts of this book.

Et cetera, et cetera.

§58.6 Choice and well-ordering

The Axiom of Choice states that given a collection Y of nonempty sets, there is a function $g : Y \rightarrow \cup Y$ which “picks” an element of each member of Y . That means $g(y) \in y$ for every $y \in Y$. (The typical illustration is that Y contains infinitely many drawers, and each drawer (a y) has some sock in it.)

Formally, it is the sentence

$$\forall Y (\emptyset \notin Y \implies \exists g : Y \rightarrow \cup Y \text{ such that } \forall y \in Y (g(y) \in y).)$$

The tricky part is not that we can conceive of such a function g , but that in fact this function g is *actually* a set.

There is an equivalent formulation which is often useful.

Definition 58.6.1. A **well-ordering** $<$ of X is a strict, total order on X which has no infinite descending chains.

Well-orderings on a set are very nice, because we can pick minimal elements: this lets us do induction, for example.

Example 58.6.2 (Examples and non-examples of well-orderings)

- (a) The natural numbers $\omega = \{0, 1, 2, \dots\}$ are well-ordered by $<$.
- (b) The integers $\mathbb{Z} = \{\dots, -2, -1, 0, 1, 2, \dots\}$ are not well-ordered by $<$, because there are infinite descending chains (take $-1 > -2 > -3 > \dots$).
- (c) The positive real numbers are not well-ordered by $<$, again because of the descending chain $\frac{1}{1} > \frac{1}{2} > \frac{1}{3} > \dots$.
- (d) The positive integers are not well-ordered by the divisibility operation $|$. While there are no descending chains, there are elements which cannot be compared (for example $3 \nmid 5$, $5 \nmid 3$ and $3 \neq 5$).

Theorem 58.6.3 (Well-ordering theorem)

Assuming Choice, for every set we can place some well-ordering on it.

In fact, the well-ordering theorem is actually equivalent to the axiom of choice.

§58.7 Sets vs classes

Prototypical example for this section: The set of all sets is the standard example of a proper class.

We close the discussion of ZFC by mentioning “classes”.

Roughly, the “bad thing” that happened was that we considered a set S , the “set of all sets”, and it was *too big*. That is,

$$\{x \mid x \text{ is a set}\}$$

is not good. Similarly, we cannot construct a set

$$\{x \mid x \text{ is an ordered pair}\}.$$

The lesson of Cantor’s Paradox is that we cannot create any sets we want; we have to be more careful than that.

Nonetheless, if we are given a set we can still tell whether or not it is an ordered pair. So for convenience, we will define a **class** to be a “concept” like the “class of all ordered pairs”. Formally, a class is defined by some formula ϕ : it consists of the sets which satisfy the formula.

In particular:

Definition 58.7.1. The class of all sets is denoted V , defined by $V = \{x \mid x = x\}$. It is called the **von Neumann universe**.

A class is a **proper class** if it is not a set, so for example we have:

Theorem 58.7.2 (There is no set of all sets)

V is a proper class.

Proof. Assume not, and V is a set. Then $V \in V$, which violates Foundation. (In fact, V cannot be a set even without Foundation, as we saw earlier). $\square$

Abuse of Notation 58.7.3. Given a class C , we will write $x \in C$ to mean that x has the defining property of C . For example, $x \in V$ means “ x is a set”.

It does not mean x is an element of V – this doesn’t make sense as V is not a set.

§58.8 Problems to think about

Problem 58A. Let A and B be sets. Show that $A \cap B$ and $A \times B$ are sets.

Problem 58B. Show that the class of all groups is a proper class. (You can take the definition of a group as a pair $(G, \cdot)$ where $\cdot$ is a function $G \times G \rightarrow G$.)

Problem 58C. Show that the axiom of choice follows from the well-ordering theorem.

Problem 58D[†]. Prove that actually, Replacement $\implies$ Comprehension.

Problem 58E (From Taiwan IMO training camp). Consider infinitely many people each wearing a hat, which is either red, green, or blue. Each person can see the hat color of everyone except themselves. Simultaneously each person guesses the color of their hat. Show that they can form a strategy such that at most finitely many people guess their color incorrectly.

59 Ordinals

§59.1 Counting for preschoolers

In preschool, we were told to count as follows. We defined a set of symbols $1, 2, 3, 4, \dots$. Then the teacher would hold up three apples and say:

“One . . . two . . . three! There are three apples.”



Image from [Ho]

The implicit definition is that the *last* number said is the final answer. This raises some obvious problems if we try to count infinite sets, but even in the finite world, this method of counting fails for the simplest set of all: how many apples are in the following picture?



Image from [Kr]

Answer: 0. There is nothing to say, and our method of counting has failed for the simplest set of all: the empty set.

§59.2 Counting for set theorists

Prototypical example for this section: $\omega + 1 = \{0, 1, 2, \dots, \omega\}$ might work.

Rather than using the *last* number listed, I propose instead starting with a list of symbols $0, 1, 2, \dots$ and making the final answer the *first* number which was *not* said. Thus to count three apples, we would say

“Zero . . . one . . . two! There are three apples.”

We will call these numbers *ordinal numbers* (rigorous definition later). In particular, we’ll *define* each ordinal to be the set of things we say:

$$\begin{aligned} 0 &= \emptyset \\ 1 &= \{0\} \\ 2 &= \{0, 1\} \\ 3 &= \{0, 1, 2\} \\ &\vdots \end{aligned}$$

In this way we can write out the natural numbers. You can have some fun with this, by saying things like

$$4 \stackrel{\text{def}}{=} \{\{\}, \{\{\}\}, \{\{\}, \{\{\}\}\}, \{\{\}, \{\{\}\}, \{\{\}, \{\{\}\}\}\}$$

In this way, we soon write down all the natural numbers. The next ordinal, ω ,¹ is defined as

$$\omega = \{0, 1, 2, \dots\}$$

Then comes

$$\begin{aligned} \omega + 1 &= \{0, 1, 2, \dots, \omega\} \\ \omega + 2 &= \{0, 1, 2, \dots, \omega, \omega + 1\} \\ \omega + 3 &= \{0, 1, 2, \dots, \omega, \omega + 1, \omega + 2\} \\ &\vdots \end{aligned}$$

And in this way we define $\omega + n$, and eventually reach


$$\begin{aligned} \omega \cdot 2 &= \omega + \omega = \{0, 1, 2, \dots, \omega, \omega + 1, \omega + 2, \dots\} \\ \omega \cdot 2 + 1 &= \{0, 1, 2, \dots, \omega, \omega + 1, \omega + 2, \dots, \omega \cdot 2\}. \end{aligned}$$

¹ That ω is actually a set is not obvious. The proof follows from the actual statement of Infinity, which I’ve dutifully omitted. In fact, Infinity is equivalent to the existence of ω .

In this way we obtain

$$\begin{aligned}
 &0, 1, 2, 3, \dots, \omega \\
 &\omega + 1, \omega + 2, \dots, \omega + \omega \\
 &\omega \cdot 2 + 1, \omega \cdot 2 + 2, \dots, \omega \cdot 3, \\
 &\vdots \\
 &\omega^2 + 1, \omega^2 + 2, \dots \\
 &\vdots \\
 &\omega^3, \dots, \omega^4, \dots, \omega^\omega \\
 &\vdots \\
 &\omega^{\omega^{\omega^{\dots}}}
 \end{aligned}$$

The first several ordinals can be illustrated in a nice spiral.



Remark 59.2.1. (Digression) The number $\omega^{\omega^{\omega^{\dots}}}$ has a name, ε_0 ; it has the property that $\omega^{\varepsilon_0} = \varepsilon_0$. The reason for using “ ε ” (which is usually used to denote small quantities) is that, despite how huge it may appear, it is actually a countable set. More on that later.

§59.3 Definition of an ordinal

Our informal description of ordinals gives us a chain

$$0 \in 1 \in 2 \in \cdots \in \omega \in \omega + 1 \in \dots$$

To give the actual definition of an ordinal, I need to define two auxiliary terms first.

Definition 59.3.1. A set x is **transitive** if whenever $z \in y \in x$, we have $z \in x$ also.

Example 59.3.2 (7 is transitive)

The set 7 is transitive: for example, $2 \in 5 \in 7 \implies 2 \in 7$.

Question 59.3.3. Show that this is equivalent to: whenever $y \in x$, $y \subseteq x$.

Moreover, recall the definition of “well-ordering”: a strict linear order with no infinite descending chains.

Example 59.3.4 ($\in$ is a well-ordering on $\omega \cdot 3$)

In $\omega \cdot 3$, we have an ordering

$$0 \in 1 \in 2 \in \cdots \in \omega \in \omega + 1 \in \cdots \in \omega \cdot 2 \in \omega \cdot 2 + 1 \in \dots$$

which has no infinite descending chains. Indeed, a typical descending chain might look like

$$\omega \cdot 2 + 6 \ni \omega \cdot 2 \ni \omega + 2015 \ni \omega + 3 \ni \omega \ni 1000 \ni 256 \ni 42 \ni 7 \ni 0.$$

Even though there are infinitely many elements, there is no way to make an infinite descending chain.

Exercise 59.3.5. (Important) Convince yourself there are no infinite descending chains of ordinals at all, even without the Foundation axiom (which certainly implies this).

Definition 59.3.6. An **ordinal** is a transitive set which is well-ordered by $\in$. The class of all ordinals is denoted On .

Question 59.3.7. Satisfy yourself that this definition works.

We typically use Greek letters α , β , etc. for ordinal numbers.

Definition 59.3.8. We write

- $\alpha < \beta$ to mean $\alpha \in \beta$, and $\alpha > \beta$ to mean $\alpha \ni \beta$.
- $\alpha \leq \beta$ to mean $\alpha \in \beta$ or $\alpha = \beta$, and $\alpha \geq \beta$ to mean $\alpha \ni \beta$ or $\alpha = \beta$,

Theorem 59.3.9 (Ordinals are strictly ordered)

Given any two ordinal numbers α and β , either $\alpha < \beta$, $\alpha = \beta$ or $\alpha > \beta$.

Proof. Surprisingly annoying, thus omitted. $\square$

Theorem 59.3.10 (Ordinals represent all order types)

Suppose $<$ is a well-ordering on a set X . Then there exists a unique ordinal α such that there is a bijection $\alpha \rightarrow X$ which is order preserving.

Thus ordinals represent the possible *equivalence classes* of order types. Any time you have a well-ordered set, it is isomorphic to a unique ordinal.

We now formalize the “+1” operation we were doing:

Definition 59.3.11. Given an ordinal α , we let $\alpha + 1 = \alpha \cup \{\alpha\}$. An ordinal of the form $\alpha + 1$ is called a **successor ordinal**.

Definition 59.3.12. If λ is an ordinal which is neither zero nor a successor ordinal, then we say λ is a **limit ordinal**.

Example 59.3.13 (Successor and limit ordinals)

7 , $\omega + 3$, $\omega \cdot 2 + 2015$ are successor ordinals, but ω and $\omega \cdot 2$ are limit ordinals.

§59.4 Ordinals are “tall”

First, we note that:

Theorem 59.4.1 (There is no set of all ordinals)

On is a proper class.

Proof. Assume for contradiction not. Then On is well-ordered by $\in$ and transitive, so On is an ordinal, i.e. $\text{On} \in \text{On}$, which violates Foundation. $\square$

Exercise 59.4.2 (Unimportant). Give a proof without Foundation by considering $\text{On} + 1$.

From this we deduce:

Theorem 59.4.3 (Sets of ordinals are bounded)

Let $A \subseteq \text{On}$. Then there is some ordinal α such that $A \subseteq \alpha$ (i.e. A must be bounded).

Proof. Otherwise, look at $\bigcup A$. It is a set. But if A is unbounded it must equal On , which is a contradiction. $\square$

In light of this, every set of ordinals has a **supremum**, which is the least upper bound. We denote this by $\sup A$.

Question 59.4.4. Show that

- (a) $\sup(\alpha + 1) = \alpha$ for any ordinal α .
- (b) $\sup \lambda = \lambda$ for any limit ordinal λ .

The pictorial “tall” will be explained in a few sections.

§59.5 Transfinite induction and recursion

The fact that $\in$ has no infinite descending chains means that induction and recursion still work verbatim.

Theorem 59.5.1 (Transfinite induction)

Given a statement $P(-)$, suppose that

- $P(0)$ is true, and
- If $P(\alpha)$ is true for all $\alpha < \beta$, then $P(\beta)$ is true.

Then $P(\alpha)$ is true for every ordinal α .

Theorem 59.5.2 (Transfinite recursion)

To define a sequence x_α for every ordinal α , it suffices to

- define x_0 , then
- for any β , define x_β for any $\alpha < \beta$.

The difference between this and normal induction lies in the *limit ordinals*. In real life, we might only do things like “define $x_{n+1} = \dots$ ”. But this is not enough to define x_α for all α , because we can’t hit ω this way. Similarly, the simple $+1$ doesn’t let us hit the ordinal 2ω , even if we already have $\omega + n$ for all n . In other words, simply incrementing by 1 cannot get us past limit stages, but using transfinite induction to jump upwards lets us sidestep this issue.

So a transfinite induction or recursion is very often broken up into three cases. In the induction phrasing, it looks like

- (Zero Case) First, resolve $P(0)$.
- (Successor Case) Show that from $P(\alpha)$ we can get $P(\alpha + 1)$.
- (Limit Case) For λ a limit ordinal, show that $P(\lambda)$ holds given $P(\alpha)$ for all $\alpha < \lambda$, where λ is a limit ordinal.

Similarly, transfinite recursion often is split into cases too.

- (Zero Case) First, define x_0 .
- (Successor Case) Define $x_{\alpha+1}$ from x_α .
- (Limit Case) Define x_λ from x_α for all $\alpha < \lambda$, where λ is a limit ordinal.

In both situations, finite induction only does the first two cases, but if we’re able to do the third case we can climb above the barrier ω .

§59.6 Ordinal arithmetic

Prototypical example for this section: $1 + \omega = \omega \neq \omega + 1$.

To give an example of transfinite recursion, let's define addition of ordinals. Recall that we defined $\alpha + 1 = \alpha \cup \{\alpha\}$. By transfinite recursion, let

$$\begin{aligned}\alpha + 0 &= \alpha \\ \alpha + (\beta + 1) &= (\alpha + \beta) + 1 \\ \alpha + \lambda &= \bigcup_{\beta < \lambda} (\alpha + \beta).\end{aligned}$$

Here $\lambda \neq 0$.

We can also do this explicitly: The picture is to just line up α after β . That is, we can consider the set

$$X = (\{0\} \times \alpha) \cup (\{1\} \times \beta)$$

(i.e. we tag each element of α with a 0, and each element of β with a 1). We then impose a well-ordering on X by a lexicographic ordering $<_{\text{lex}}$ (sort by first component, then by second). This well-ordering is isomorphic to a unique ordinal,

Example 59.6.1 ($2 + 3 = 5$)

Under the explicit construction for $\alpha = 2$ and $\beta = 3$, we get the set

$$X = \{(0, 0) < (0, 1) < (1, 0) < (1, 1) < (1, 2)\}$$

which is isomorphic to 5.

Example 59.6.2 (Ordinal arithmetic is not commutative)

Note that $1 + \omega = \omega$! Indeed, under the transfinite definition, we have

$$1 + \omega = \bigcup_n (1 + n) = 2 \cup 3 \cup 4 \cup \dots = \omega.$$

With the explicit construction, we have

$$X = \{(0, 0) < (1, 0) < (1, 1) < (1, 2) < \dots\}$$

which is isomorphic to ω .

Exercise 59.6.3. Show that $n + \omega = \omega$ for any $n \in \omega$.

Remark 59.6.4. Ordinal addition is not commutative. However, from the explicit construction we can see that it is at least associative.

Similarly, we can define multiplication in two ways. By transfinite induction:

$$\begin{aligned}\alpha \cdot 0 &= 0 \\ \alpha \cdot (\beta + 1) &= (\alpha \cdot \beta) + \alpha \\ \alpha \cdot \lambda &= \bigcup_{\beta < \lambda} \alpha \cdot \beta.\end{aligned}$$

We can also do an explicit construction: $\alpha \cdot \beta$ is the order type of

$$<_{\text{lex}} \text{ applied to } \beta \times \alpha.$$

Example 59.6.5 (Ordinal multiplication is not commutative)

We have $\omega \cdot 2 = \omega + \omega$, but $2 \cdot \omega = \omega$.

Exercise 59.6.6. Prove this.

Exercise 59.6.7. Verify that ordinal multiplication (like addition) is associative but not commutative. (Look at $\gamma \times \beta \times \alpha$.)

Exponentiation can also be so defined, though the explicit construction is less natural.

$$\begin{aligned}\alpha^0 &= 1 \\ \alpha^{\beta+1} &= \alpha^\beta \cdot \alpha \\ \alpha^\lambda &= \bigcup_{\beta < \lambda} \alpha^\beta.\end{aligned}$$

Exercise 59.6.8. Verify that $2^\omega = \omega$.

§59.7 The hierarchy of sets

We now define the **von Neumann Hierarchy** by transfinite recursion.

Definition 59.7.1. By transfinite recursion, we set

$$\begin{aligned}V_0 &= \emptyset \\ V_{\alpha+1} &= \mathcal{P}(V_\alpha) \\ V_\lambda &= \bigcup_{\alpha < \lambda} V_\alpha\end{aligned}$$

By transfinite induction, we see V_α is transitive and that $V_\alpha \subseteq V_\beta$ for all $\alpha < \beta$.

Example 59.7.2 (V_α for $\alpha \leq 3$)

The first few levels of the hierarchy are:

$$\begin{aligned}V_0 &= \emptyset \\ V_1 &= \{0\} \\ V_2 &= \{0, 1\} \\ V_3 &= \{0, 1, 2, \{1\}\}.\end{aligned}$$

Notice that for each n , V_n consists of only finite sets, and each n appears in V_{n+1} for the first time. Observe that

$$V_\omega = \bigcup_{n \in \omega} V_n$$

consists only of finite sets; thus ω appears for the first time in $V_{\omega+1}$.

Question 59.7.3. How many sets are in V_5 ?

Definition 59.7.4. The **rank** of a set y , denoted $\text{rank}(y)$, is the smallest ordinal α such that $y \in V_{\alpha+1}$.

Example 59.7.5

$\text{rank}(2) = 2$, and actually $\text{rank}(\alpha) = \alpha$ for any ordinal α (problem later). This is the reason for the extra “+1”.

Question 59.7.6. Show that $\text{rank}(y)$ is the smallest ordinal α such that $y \subseteq V_\alpha$.

It’s not yet clear that the rank of a set actually exists, so we prove:

Theorem 59.7.7 (The von Neumann hierarchy is complete)


The class V is equal to $\bigcup_{\alpha \in \text{On}} V_\alpha$. In other words, every set appears in some V_α .

Proof. Assume for contradiction this is false. The key is that because $\in$ satisfies Foundation, we can take a $\in$ -minimal counterexample x . Thus $\text{rank}(y)$ is defined for every $y \in x$, and we can consider (by Replacement) the set

$$\{\text{rank}(y) \mid y \in x\}.$$

Since it is a set of ordinals, it is bounded. So there is some large ordinal α such that $y \in V_\alpha$ for all $y \in x$, i.e. $x \subseteq V_\alpha$, so $x \in V_{\alpha+1}$. $\square$

This leads us to a picture of the universe V :



We can imagine the universe V as a triangle, built in several stages or layers, $V_0 \subsetneq V_1 \subsetneq V_2 \subsetneq \dots$. This universe doesn't have a top: but each of the V_i do. However, the universe has a very clear bottom. Each stage is substantially wider than the previous one.

In the center of this universe are the ordinals: for every successor V_α , exactly one new ordinal appears, namely α . Thus we can picture the class of ordinals as a thin line that stretches the entire height of the universe. A set has rank α if it appears at the same stage that α does.

All of number theory, the study of the integers, lives inside V_ω . Real analysis, the study of real numbers, lives inside $V_{\omega+1}$, since a real number can be encoded as a subset of $\mathbb{N}$ (by binary expansion). Functional analysis lives one step past that, $V_{\omega+2}$. For all intents and purposes, most mathematics does not go beyond $V_{\omega+\omega}$. This pales in comparison to the true magnitude of the whole universe.

§59.8 Problems to think about

Problem 59A. Prove that $\text{rank}(\alpha) = \alpha$ for any α by transfinite induction.

Problem 59B (Online Math Open). Count the number of transitive sets in V_5 .

Problem 59C (Goodstein). Let a_2 be any positive integer. We define the infinite sequence $a_2, a_3, \dots$ recursively as follows. If $a_n = 0$, then $a_{n+1} = 0$. Otherwise, we write a_n in base n , then write all exponents in base n , and so on until all numbers in the expression are at most n . Then we replace all instances of n by $n+1$ (including the exponents!), subtract 1, and set the result to a_{n+1} . For example, if $a_2 = 11$ we have

$$\begin{aligned} a_2 &= 2^3 + 2 + 1 = 2^{2+1} + 2 + 1 \\ a_3 &= 3^{3+1} + 3 + 1 - 1 = 3^{3+1} + 3 \\ a_4 &= 4^{4+1} + 4 - 1 = 4^{4+1} + 3 \\ a_5 &= 5^{5+1} + 3 - 1 = 5^{5+1} + 2 \end{aligned}$$

and so on. Prove that $a_N = 0$ for some integer $N > 2$.

60 Cardinals

An ordinal measures a total ordering. However, it does not do a fantastic job at measuring size. For example, there is a bijection between the elements of ω and $\omega + 1$:

$$\begin{aligned}\omega + 1 &= \{ \omega \ 0 \ 1 \ 2 \ \dots \} \\ \omega &= \{ \ 0 \ 1 \ 2 \ 3 \ \dots \}.\end{aligned}$$

In fact, as you likely already know, there is even a bijection between ω and ω^2 :

+	0	1	2	3	4	...
0	0	1	3	6	10	...
ω	2	4	7	11	...	
$\omega \cdot 2$	5	8	12	...		
$\omega \cdot 3$	9	13	...			
$\omega \cdot 4$	14	...				

So ordinals do not do a good job of keeping track of size. For this, we turn to the notion of a cardinal number.

§60.1 Equinumerous sets and cardinals

Definition 60.1.1. Two sets A and B are **equinumerous**, written $A \approx B$, if there is a bijection between them.

Definition 60.1.2. A **cardinal** is an ordinal κ such that for no $\alpha < \kappa$ do we have $\alpha \approx \kappa$.

Example 60.1.3 (Examples of cardinals)

Every finite number is a cardinal. Moreover, ω is a cardinal. However, $\omega + 1$, ω^2 , ω^{2015} are not, because they are countable.

Example 60.1.4 (ω^ω is countable)

Even ω^ω is not a cardinal, since it is a countable union

$$\omega^\omega = \bigcup_n \omega^n$$

and each ω^n is countable.

Question 60.1.5. Why must an infinite cardinal be a limit ordinal?

Remark 60.1.6. There is something fishy about the definition of a cardinal: it relies on an *external* function f . That is, to verify κ is a cardinal I can't just look at κ itself; I need to examine the entire universe V to make sure there does not exist a bijection $f : \kappa \rightarrow \alpha$ for $\alpha < \kappa$. For now this is no issue, but later in model theory this will lead to some highly counterintuitive behavior.

§60.2 Cardinalities

Now that we have defined a cardinal, we can discuss the size of a set by linking it to a cardinal.

Definition 60.2.1. The **cardinality** of a set X is the *least* ordinal κ such that $X \approx \kappa$. We denote it by $|X|$.

Question 60.2.2. Why must $|X|$ be a cardinal?

Remark 60.2.3. One needs the well-ordering theorem (equivalently, choice) in order to establish that such an ordinal κ actually exists.

Since cardinals are ordinals, it makes sense to ask whether $\kappa_1 \leq \kappa_2$, and so on. Our usual intuition works well here.

Proposition 60.2.4 (Restatement of cardinality properties)

Let X and Y be sets.

- (i) $X \approx Y$ if and only $|X| = |Y|$, if and only if there's a bijection from X to Y .
- (ii) $|X| \leq |Y|$ if and only if there is an injective map $X \hookrightarrow Y$.

Diligent readers are invited to try and prove this.

§60.3 Aleph numbers

Prototypical example for this section: $\aleph_0$ is ω , and $\aleph_1$ is the first uncountable.

First, let us check that cardinals can get arbitrarily large:

Proposition 60.3.1

We have $|X| < |\mathcal{P}(X)|$ for every set X .

Proof. There is an injective map $X \hookrightarrow \mathcal{P}(X)$ but there is no injective map $\mathcal{P}(X) \hookrightarrow X$ by Lemma 58.2.1. $\square$

Thus we can define:

Definition 60.3.2. For a cardinal κ , we define κ^+ to be the least cardinal above κ , called the **successor cardinal**.

This κ^+ exists and has $\kappa^+ \leq |\mathcal{P}(\kappa)|$.

Next, we claim that:

Exercise 60.3.3. Show that if A is a set of cardinals, then $\cup A$ is a cardinal.

Thus by transfinite induction we obtain that:

Definition 60.3.4. For any $\alpha \in \text{On}$, we define the **aleph numbers** as

$$\begin{aligned}\aleph_0 &= \omega \\ \aleph_{\alpha+1} &= (\aleph_\alpha)^+ \\ \aleph_\lambda &= \bigcup_{\alpha < \lambda} \aleph_\alpha.\end{aligned}$$

Thus we have the sequence of cardinals

$$0 < 1 < 2 < \cdots < \aleph_0 < \aleph_1 < \cdots < \aleph_\omega < \aleph_{\omega+1} < \cdots$$

By definition, $\aleph_0$ is the cardinality of the natural numbers, $\aleph_1$ is the first uncountable ordinal, $\dots$

We claim the aleph numbers constitute all the cardinals:

Lemma 60.3.5 (Aleph numbers constitute all infinite cardinals)

If κ is a cardinal then either κ is finite (i.e. $\kappa \in \omega$) or $\kappa = \aleph_\alpha$ for some $\alpha \in \text{On}$.

Proof. Assume κ is infinite, and take α minimal with $\aleph_\alpha \geq \kappa$. Suppose for contradiction that we have $\aleph_\alpha > \kappa$. We may assume $\alpha > 0$, since the case $\alpha = 0$ is trivial.

If $\alpha = \bar{\alpha} + 1$ is a successor, then

$$\aleph_{\bar{\alpha}} < \kappa < \aleph_\alpha = (\aleph_{\bar{\alpha}})^+$$

which contradicts the definition of the successor cardinal.

If $\alpha = \lambda$ is a limit ordinal, then $\aleph_\lambda$ is the supremum $\bigcup_{\gamma < \lambda} \aleph_\gamma$. So there must be some $\gamma < \lambda$ with $\aleph_\gamma > \kappa$, which contradicts the minimality of α . $\square$

Definition 60.3.6. An infinite cardinal which is not a successor cardinal is called a **limit cardinal**. It is exactly those cardinals of the form $\aleph_\lambda$, for λ a limit ordinal, plus $\aleph_0$.

§60.4 Cardinal arithmetic

Prototypical example for this section: $\aleph_0 \cdot \aleph_0 = \aleph_0 + \aleph_0 = \aleph_0$

Recall the way we set up ordinal arithmetic. Note that in particular, $\omega + \omega > \omega$ and $\omega^2 > \omega$. Since cardinals count size, this property is undesirable, and we want to have

$$\begin{aligned}\aleph_0 + \aleph_0 &= \aleph_0 \\ \aleph_0 \cdot \aleph_0 &= \aleph_0\end{aligned}$$

because $\omega + \omega$ and $\omega \cdot \omega$ are countable. In the case of cardinals, we simply “ignore order”.

The definition of cardinal arithmetic is as expected:

Definition 60.4.1 (Cardinal arithmetic). Given cardinals κ and μ , define

$$\kappa + \mu \stackrel{\text{def}}{=} |(\{0\} \times \kappa) \cup (\{1\} \times \mu)|$$

and

$$\kappa \cdot \mu \stackrel{\text{def}}{=} |\mu \times \kappa|.$$

Question 60.4.2. Check this agrees with what you learned in pre-school for finite cardinals.

Abuse of Notation 60.4.3. This is a slight abuse of notation since we are using the same symbols as for ordinal arithmetic, even though the results are different ($\omega \cdot \omega = \omega^2$ but $\aleph_0 \cdot \aleph_0 = \aleph_0$). In general, I'll make it abundantly clear whether I am talking about cardinal arithmetic or ordinal arithmetic.

To help combat this confusion, we use separate symbols for ordinals and cardinals. Specifically, ω will always refer to $\{0, 1, \dots\}$ viewed as an ordinal; $\aleph_0$ will always refer to the same set viewed as a cardinal. More generally,

Definition 60.4.4. Let $\omega_\alpha = \aleph_\alpha$ viewed as an ordinal.

However, as we've seen already we have that $\aleph_0 \cdot \aleph_0 = \aleph_0$. In fact, this holds even more generally:

Theorem 60.4.5 (Infinite cardinals squared)

Let κ be an infinite cardinal. Then $\kappa \cdot \kappa = \kappa$.

Proof. Obviously $\kappa \cdot \kappa \geq \kappa$, so we want to show $\kappa \cdot \kappa \leq \kappa$.

The idea is to try to repeat the same proof that we had for $\aleph_0 \cdot \aleph_0 = \aleph_0$, so we re-iterate it here. We took the “square” of elements of $\aleph_0$, and then *re-ordered* it according to the diagonal:

	0	1	2	3	4	...
0	0	1	3	6	10	...
1	2	4	7	11	...	
2	5	8	12	...		
3	9	13	...			
4	14	...				

We'd like to copy this idea for a general κ ; however, since addition is less well-behaved for infinite ordinals it will be more convenient to use $\max(\alpha, \beta)$ rather than $\alpha + \beta$. Specifically, we put the ordering $<_{\max}$ on $\kappa \times \kappa$ as follows: for (α_1, β_1) and (α_2, β_2) in $\kappa \times \kappa$ we declare $(\alpha_1, \beta_1) <_{\max} (\alpha_2, \beta_2)$ if

- $\max\{\alpha_1, \beta_1\} < \max\{\alpha_2, \beta_2\}$ or
- $\max\{\alpha_1, \beta_1\} = \max\{\alpha_2, \beta_2\}$ and (α_1, β_1) is lexicographically earlier than (α_2, β_2) .

This alternate ordering (which deliberately avoids referring to the addition) looks like:

	0	1	2	3	4	...
0	0	1	4	9	16	...
1	2	3	5	10	17	...
2	6	7	8	11	18	...
3	12	13	14	15	19	...
4	20	21	22	23	24	...
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$

Now we proceed by transfinite induction on κ . The base case is $\kappa = \aleph_0$, done above. Now, $<_{\max}$ is a well-ordering of $\kappa \times \kappa$, so we know it is in order-preserving bijection with some ordinal γ . Our goal is to show that $|\gamma| \leq \kappa$. To do so, it suffices to prove that for any $\bar{\gamma} \in \gamma$, we have $|\bar{\gamma}| < \kappa$.

Suppose $\bar{\gamma}$ corresponds to the point $(\alpha, \beta) \in \kappa$ under this bijection. If α and β are both finite then certainly $\bar{\gamma}$ is finite too. Otherwise, let $\bar{\kappa} = \max\{\alpha, \beta\} < \kappa$; then the number of points below $\bar{\gamma}$ is at most

$$|\alpha| \cdot |\beta| \leq \bar{\kappa} \cdot \bar{\kappa} = \bar{\kappa}$$

by the inductive hypothesis. So $|\bar{\gamma}| \leq \bar{\kappa} < \kappa$ as desired. $\square$

From this it follows that cardinal addition and multiplication is really boring:

Theorem 60.4.6 (Infinite cardinal arithmetic is trivial)

Given cardinals κ and μ , one of which is infinite, we have

$$\kappa \cdot \mu = \kappa + \mu = \max(\kappa, \mu).$$

Proof. The point is that both of these are less than the square of the maximum. Writing out the details:

$$\begin{aligned} \max(\kappa, \mu) &\leq \kappa + \mu \\ &\leq \kappa \cdot \mu \\ &\leq \max(\kappa, \mu) \cdot \max(\kappa, \mu) \\ &= \max(\kappa, \mu). \end{aligned}$$

$\square$

§60.5 Cardinal exponentiation

Prototypical example for this section: $2^\kappa = |\mathcal{P}(\kappa)|$.

Definition 60.5.1. Suppose κ and λ are cardinals. Then

$$\kappa^\lambda \stackrel{\text{def}}{=} |\mathcal{F}(\lambda, \kappa)|.$$

Here $\mathcal{F}(A, B)$ is the set of functions from A to B .

Abuse of Notation 60.5.2. As before, we are using the same notation for both cardinal and ordinal arithmetic. Sorry!

In particular, $2^\kappa = |\mathcal{P}(\kappa)| > \kappa$, and so from now on we can use the notation 2^κ freely. (Note that this is totally different from ordinal arithmetic; there we had $2^\omega = \bigcup_{n \in \omega} 2^n = \omega$. In cardinal arithmetic $2^{\aleph_0} > \aleph_0$.)

I have unfortunately not told you what $2^{\aleph_0}$ equals. A natural conjecture is that $2^{\aleph_0} = \aleph_1$; this is called the **Continuum Hypothesis**. It turns out that this is *undecidable* – it is not possible to prove or disprove this from the ZFC axioms.

§60.6 Cofinality

Prototypical example for this section: $\aleph_0, \aleph_1, \dots$ are all regular, but $\aleph_\omega$ has cofinality ω .

Definition 60.6.1. Let λ be a limit ordinal, and α another ordinal. A map $f : \alpha \rightarrow \lambda$ of ordinals is called **cofinal** if for every $\bar{\lambda} < \lambda$, there is some $\bar{\alpha} \in \alpha$ such that $f(\bar{\alpha}) \geq \bar{\lambda}$. In other words, the map reaches arbitrarily high into λ .

Example 60.6.2 (Example of a cofinal map)

- (a) The map $\omega \rightarrow \omega^\omega$ by $n \mapsto \omega^n$ is cofinal.
- (b) For any ordinal α , the identity map $\alpha \rightarrow \alpha$ is cofinal.

Definition 60.6.3. Let λ be a limit ordinal. The **cofinality** of λ , denoted $\text{cof}(\lambda)$, is the smallest ordinal α such that there is a cofinal map $\alpha \rightarrow \lambda$.

Question 60.6.4. Why must α be an infinite cardinal?

Usually, we are interested in taking the cofinality of a cardinal κ .

Pictorially, you can imagine standing at the bottom of the universe and looking up the chain of ordinals to κ . You have a machine gun and are firing bullets upwards, and you want to get arbitrarily high but less than κ . The cofinality is then the number of bullets you need to do this.

We now observe that “most” of the time, the cofinality of a cardinal is itself. Such a cardinal is called **regular**.

Example 60.6.5 ($\aleph_0$ is regular)

$\text{cof}(\aleph_0) = \aleph_0$, because no finite subset of ω can reach arbitrarily high.

Example 60.6.6 ($\aleph_1$ is regular)

$\text{cof}(\aleph_1) = \aleph_1$. Indeed, assume for contradiction that some countable set of ordinals $A = \{\alpha_0, \alpha_1, \dots\} \subseteq \omega_1$ reaches arbitrarily high inside ω_1 . Then $\Lambda = \cup A$ is a *countable* ordinal, because it is a countable union of countable ordinals. In other words $\Lambda \in \omega_1$. But Λ is an upper bound for A , contradiction.

On the other hand, there *are* cardinals which are not regular; since these are the “rare” cases we call them **singular**.

Example 60.6.7 ($\aleph_\omega$ is not regular)

Notice that $\aleph_0 < \aleph_1 < \aleph_2 < \dots$ reaches arbitrarily high in $\aleph_\omega$, despite only having $\aleph_0$ terms. It follows that $\text{cof}(\aleph_\omega) = \aleph_0$.

We now confirm a suspicion you may have:

Theorem 60.6.8 (Successor cardinals are regular)

If $\kappa = \bar{\kappa}^+$ is a successor cardinal, then it is regular.

Proof. We copy the proof that $\aleph_1$ was regular.

Assume for contradiction that for some $\mu \leq \bar{\kappa}$, there are μ sets reaching arbitrarily high in κ as a cardinal. Observe that each of these sets must have cardinality at most $\bar{\kappa}$. We take the union of all μ sets, which gives an ordinal Λ serving as an upper bound.

The number of elements in the union is at most

$$\#\text{sets} \cdot \#\text{elms} \leq \mu \cdot \bar{\kappa} = \bar{\kappa}$$

and hence $|\Lambda| \leq \bar{\kappa} < \kappa$. □

So, what about limit cardinals? It seems to be that most of them are singular: if $\aleph_\lambda \neq \aleph_0$ is a limit ordinal, then the sequence $\{\aleph_\alpha\}_{\alpha \in \lambda}$ (of length λ) is certainly cofinal.

Consider the monstrous cardinal

$$\kappa = \aleph_{\aleph_{\aleph}}$$

This might look frighteningly huge, as $\kappa = \aleph_\kappa$, but its cofinality is ω as it is the limit of the sequence

$$N_0, N_{N_0}, N_{N_{N_0}}, \dots$$

More generally, one can in fact prove that

$$\mathrm{cof}(\aleph_\lambda) = \mathrm{cof}(\lambda).$$

But it is actually conceivable that λ is so large that $|\lambda| = |\aleph_\lambda|$.

A regular limit cardinal other than $\aleph_0$ has a special name: it is **weakly inaccessible**. Such cardinals are so large that it is impossible to prove or disprove their existence in ZFC. It is the first of many so-called “large cardinals”.

An infinite cardinal κ is a strong limit cardinal if

$$\forall \bar{\kappa} < \kappa \quad 2^{\bar{\kappa}} < \kappa$$

for any cardinal $\bar{\kappa}$. For example, $\aleph_0$ is a strong limit cardinal.

Question 60.7.2. Why must strong limit cardinals actually be limit cardinals? (This is offensively easy.)

A regular strong limit cardinal other than $\aleph_0$ is called **strongly inaccessible**.

§60.8 Problems to think about

Problem 60A. Compute $|V_\omega|$.

Problem 60B. Prove that for any limit ordinal α , $\text{cof}(\alpha)$ is a *regular* cardinal.

Problem 60C[★] (Strongly inaccessible cardinals). Show that for any strongly inaccessible κ , we have $|V_\kappa| = \kappa$.

Problem 60D (König's theorem). Show that

$$\kappa^{\text{cof}(\kappa)} > \kappa$$

for every infinite cardinal κ .

XVI

Set Theory II: Model Theory and Forcing

61 Inner model theory	551
61.1 Models	551
61.2 Sentences and satisfaction	552
61.3 The Levy hierarchy	554
61.4 Substructures, and Tarski-Vaught	555
61.5 Obtaining the axioms of ZFC	556
61.6 Mostowski collapse	557
61.7 Adding an inaccessible	557
61.8 FAQ's on countable models	559
61.9 Picturing inner models	559
61.10 Problems to think about	561
62 Forcing	562
62.1 Setting up posets	563
62.2 More properties of posets	564
62.3 Names, and the generic extension	565
62.4 Fundamental theorem of forcing	568
62.5 (Optional) Defining the relation	568
62.6 The remaining axioms	570
62.7 Problems to think about	570
63 Breaking the continuum hypothesis	571
63.1 Adding in reals	571
63.2 The countable chain condition	572
63.3 Preserving cardinals	573
63.4 Infinite combinatorics	574
63.5 Problems to think about	575

61 Inner model theory

Model theory is *really* meta, so you will have to pay attention here.

Roughly, a “model of ZFC” is a set with a binary relation that satisfies the ZFC axioms, just as a group is a set with a binary operation that satisfies the group axioms. Unfortunately, unlike with groups, it is very hard for me to give interesting examples of models, for the simple reason that we are literally trying to model the entire universe.

§61.1 Models

Prototypical example for this section: $(\omega, \in)$ obeys PowerSet, V_κ is a model for κ inaccessible (later).

Definition 61.1.1. A **model** $\mathcal{M}$ consists of a set M and a binary relation $E \subseteq M \times M$. (The E relation is the “ $\in$ ” for the model.)

Remark 61.1.2. I’m only considering *set-sized* models where M is a set. Experts may be aware that I can actually play with M being a class, but that would require too much care for now.

If you have a model, you can ask certain things about it, such as “does it satisfy EmptySet?”. Let me give you an example of what I mean and then make it rigorous.

Example 61.1.3 (A stupid model)

Let’s take $\mathcal{M} = (M, E) = (\omega, \in)$. This is not a very good model of ZFC, but let’s see if we can make sense of some of the first few axioms.

(a) $\mathcal{M}$ satisfies Extensionality, which is the sentence

$$\forall x \forall y \forall a : (a \in x \iff a \in y) \implies x = y.$$

This just follows from the fact that E is actually $\in$.

(b) $\mathcal{M}$ satisfies EmptySet, which is the sentence

$$\exists a : \forall x \neg(x \in a).$$

Namely, take $a = \emptyset \in \omega$.

(c) $\mathcal{M}$ does not satisfy Pairing, since $\{1, 3\}$ is not in ω , even though $1, 3 \in \omega$.

(d) Miraculously, $\mathcal{M}$ satisfies Union, since for any $n \in \omega$, $\cup n$ is $n - 1$ (unless $n = 0$). The Union axiom statements that

$$\forall a \exists U \quad \forall x [(x \in U) \iff (\exists y : x \in y \in a)].$$

An important thing to notice is that the “ $\forall a$ ” ranges only over the sets in the model of the universe, $\mathcal{M}$.

Example 61.1.4 (Important: this stupid model satisfies PowerSet)

Most incredibly of all: $\mathcal{M} = (\omega, \in)$ satisfies PowerSet. This is a really important example.

You might think this is ridiculous. Look at $2 = \{0, 1\}$. The power set of this is $\{0, 1, 2, \{1\}\}$ which is not in the model, right?

Well, let's look more closely at PowerSet. It states that:

$$\forall x \exists a \forall y (y \in a \iff y \subseteq x).$$

What happens if we set $x = 2 = \{0, 1\}$? Well, actually, we claim that $a = 3 = \{0, 1, 2\}$ works. The key point is “for all y ” – this *only ranges over the objects in $\mathcal{M}$* . In $\mathcal{M}$, the only subsets of 2 are $0 = \emptyset$, $1 = \{0\}$ and $2 = \{0, 1\}$. The “set” $\{1\}$ in the “real world” (in V) is not a set in the model $\mathcal{M}$.

In particular, you might say that in this strange new world, we have $2^n = n + 1$, since $n = \{0, 1, \dots, n - 1\}$ really does have only $n + 1$ subsets.

Example 61.1.5 (Sentences with parameters)

The sentences we ask of our model are allowed to have “parameters” as well. For example, if $\mathcal{M} = (\omega, \in)$ as before then $\mathcal{M}$ satisfies the sentence

$$\forall x \in 3 (x \in 5).$$

§61.2 Sentences and satisfaction

With this intuitive notion, we can define what it means for a model to satisfy a sentence.

Definition 61.2.1. Note that any sentence ϕ can be written in one of five forms:

- $x \in y$
- $x = y$
- $\neg\psi$ (“not ψ ”) for some shorter sentence ψ
- $\psi_1 \vee \psi_2$ (“ ψ_1 or ψ_2 ”) for some shorter sentences ψ_1, ψ_2
- $\exists x\psi$ (“exists x ”) for some shorter sentence ψ .

Question 61.2.2. What happened to $\wedge$ (and) and $\forall$ (for all)? (Hint: use $\neg$.)

Often (almost always, actually) we will proceed by so-called “induction on formula complexity”, meaning that we define or prove something by induction using this. Note that we require all formulas to be finite.

Now suppose we have a sentence ϕ , like $a = b$ or $\exists a \forall x \neg(x \in a)$, plus a model $\mathcal{M} = (M, E)$. We want to ask whether $\mathcal{M}$ satisfies ϕ .

To give meaning to this, we have to designate certain variables as **parameters**. For example, if I asked you

“Does $a = b$?”

the first question you would ask is what a and b are. So a, b would be parameters: I have to give them values for this sentence to make sense.

On the other hand, if I asked you

“Does $\exists a \forall x \neg(x \in a)$?”

then you would just say “yes”. In this case, x and a are *not* parameters. In general, parameters are those variables whose meaning is not given by some $\forall$ or $\exists$.

In what follows, we will let $\phi(x_1, \dots, x_n)$ denote a formula ϕ , whose parameters are $x_1, \dots, x_n$. Note that possibly $n = 0$, for example all ZFC axioms have no parameters.

Question 61.2.3. Try to guess the definition of satisfaction before reading it below. (It’s not very hard to guess!)

Definition 61.2.4. Let $\mathcal{M} = (M, E)$ be a model. Let $\phi(x_1, \dots, x_n)$ be a sentence, and let $b_1, \dots, b_n \in M$. We will define a relation

$$\mathcal{M} \models \phi[b_1, \dots, b_n]$$

and say $\mathcal{M}$ **satisfies** the sentence ϕ with parameters $b_1, \dots, b_n$.

The relationship is defined by induction on formula complexity as follows:

- If ϕ is “ $x_1 = x_2$ ” then $\mathcal{M} \models \phi[b_1, b_2] \iff b_1 = b_2$.
- If ϕ is “ $x_1 \in x_2$ ” then $\mathcal{M} \models \phi[b_1, b_2] \iff b_1 E b_2$.
(This is what we mean by “ E interprets $\in$ ”.)
- If ϕ is “ $\neg\psi$ ” then $\mathcal{M} \models \phi[b_1, \dots, b_n] \iff \mathcal{M} \not\models \psi[b_1, \dots, b_n]$.
- If ϕ is “ $\psi_1 \vee \psi_2$ ” then $\mathcal{M} \models \phi[b_1, \dots, b_n]$ means $\mathcal{M} \models \psi_i[b_1, \dots, b_n]$ for some $i = 1, 2$.
- Most important case: suppose ϕ is $\exists x \psi(x, x_1, \dots, x_n)$. Then $\mathcal{M} \models \phi[b_1, \dots, b_n]$ if and only if

$$\exists b \in M \text{ such that } \mathcal{M} \models \psi[b, b_1, \dots, b_n].$$

Note that ψ has one extra parameter.

Notice where the information of the model actually gets used. We only ever use E in interpreting $x_1 \in x_2$; unsurprising. But we only ever use the set M when we are running over $\exists$ (and hence $\forall$). That’s well-worth keeping in mind:

The behavior of a model essentially comes from $\exists$ and $\forall$, which search through the entire model M .

And finally,

Definition 61.2.5. A **model of ZFC** is a model $\mathcal{M} = (M, E)$ satisfying all ZFC axioms.

We are especially interested in models of the form $(M, \in)$, where M is a *transitive* set. (We want our universe to be transitive, otherwise we would have elements of sets which are not themselves in the universe, which is very strange.) Such a model is called a **transitive model**.

Abuse of Notation 61.2.6. If M is a transitive set, the model $(M, \in)$ will be abbreviated to just M .

Definition 61.2.7. An **inner model** of ZFC is a transitive model satisfying ZFC.

§61.3 The Levy hierarchy

Prototypical example for this section: `isSubset`(x, y) is absolute. The axiom `EmptySet` is Σ_1 , `isPowerSetOf`(X, x) is Π_1 .

A key point to remember is that the behavior of a model is largely determined by $\exists$ and $\forall$. It turns out we can say even more than this.

Consider a formula such as

$$\text{isEmpty}(x) : \neg \exists a (a \in x)$$

which checks whether a given set x has an element in it. Technically, this has an “ $\exists$ ” in it. But somehow this $\exists$ does not really search over the entire model, because it is *bounded* to search in x . That is, we might informally rewrite this as

$$\neg(\exists a \in x)$$

which doesn’t fit into the strict form, but points out that we are only looking over $a \in x$. We call such a quantifier a **bounded quantifier**.

We like sentences with bounded quantifiers because they designate properties which are **absolute** over transitive models. It doesn’t matter how strange your surrounding model M is. As long as M is transitive,

$$M \models \text{isEmpty}(\emptyset)$$

will always hold. Similarly, the sentence

$$\text{isSubset}(x, y) : x \subseteq y \text{ i.e. } \forall a \in x (a \in y)$$

is absolute. Sentences with this property are called Σ_0 or Π_0 .

The situation is different with a sentence like

$$\text{isPowerSetOf}(y, x) : \forall z (z \subseteq x \iff z \in y)$$

which in English means “ y is the power set of x ”, or just $y = \mathcal{P}(x)$. The $\forall z$ is *not* bounded here. This weirdness is what allows things like

$$\omega \models “\{0, 1, 2\} \text{ is the power set of } \{0, 1\}”$$

and hence

$$\omega \models \text{PowerSet}$$

which was our stupid example earlier. The sentence `isPowerSetOf` consists of an unbounded $\forall$ followed by an absolute sentence, so we say it is Π_1 .

More generally, the **Levy hierarchy** keeps track of how bounded our quantifiers are. Specifically,

- Formulas which have only bounded quantifiers are $\Delta_0 = \Sigma_0 = \Pi_0$.
- Formulas of the form $\exists x_1 \dots \exists x_k \psi$ where ψ is Π_n are considered Σ_{n+1} .
- Formulas of the form $\forall x_1 \dots \forall x_k \psi$ where ψ is Σ_n are considered Π_{n+1} .

(A formula which is both Σ_n and Π_n is called Δ_n , but we won’t use this except for $n = 0$.)

Example 61.3.1 (Examples of Δ_0 sentences)

- (a) The sentences $\text{isEmpty}(x)$, $x \subseteq y$, as discussed above.
- (b) The formula “ x is transitive” can be expanded as a Δ_0 sentence.
- (c) The formula “ x is an ordinal” can be expanded as a Δ_0 sentence.

Exercise 61.3.2. Write out the expansions for “ x is transitive” and “ x is ordinal” in a Δ_0 form.

Example 61.3.3 (More complex formulas)

- (a) The axiom EmptySet is Σ_1 ; it is $\exists a(\text{isEmpty}(a))$, and $\text{isEmpty}(a)$ is Δ_0 .
- (b) The formula “ $y = \mathcal{P}(x)$ ” is Π_1 , as discussed above.
- (c) The formula “ x is countable” is Σ_1 . One way to phrase it is “ $\exists f$ an injective map $x \hookrightarrow \omega$ ”, which necessarily has an unbounded “ $\exists f$ ”.
- (d) The axiom PowerSet is Π_3 :

$$\forall y \exists P \forall x (x \subseteq y \iff x \in P).$$

§61.4 Substructures, and Tarski-Vaught

Let $\mathcal{M}_1 = (M_1, E_1)$ and $\mathcal{M}_2 = (M_2, E_2)$ be models.

Definition 61.4.1. We say that $\mathcal{M}_1 \subseteq \mathcal{M}_2$ if $M_1 \subseteq M_2$ and E_1 agrees with E_2 ; we say $\mathcal{M}_1$ is a **substructure** of $\mathcal{M}_2$.

That’s boring. The good part is:

Definition 61.4.2. We say $\mathcal{M}_1 \prec \mathcal{M}_2$, or $\mathcal{M}_1$ is an **elementary substructure** of $\mathcal{M}_2$, if $\mathcal{M}_1 \subseteq \mathcal{M}_2$ and for *every* sentence $\phi(x_1, \dots, x_n)$ and parameters $b_1, \dots, b_n \in M_1$, we have

$$\mathcal{M}_1 \models \phi[b_1, \dots, b_n] \iff \mathcal{M}_2 \models \phi[b_1, \dots, b_n].$$

In other words, $\mathcal{M}_1$ and $\mathcal{M}_2$ agree on every sentence possible. Note that the b_i have to come from $\mathcal{M}_1$; if the b_i came from $\mathcal{M}_2$ then asking something of $\mathcal{M}_1$ wouldn’t make sense.

Let’s ask now: how would $\mathcal{M}_1 \prec \mathcal{M}_2$ fail to be true? If we look at the possibly sentences, none of the atomic formulas, nor the “ $\wedge$ ” and “ $\neg$ ”, are going to cause issues.

The intuition you should be getting by now is that things go wrong once we hit $\forall$ and $\exists$. They won’t go wrong for bounded quantifiers. But unbounded quantifiers search the entire model, and that’s where things go wrong.

To give a “concrete example”: imagine $\mathcal{M}_1$ is MIT, and $\mathcal{M}_2$ is the state of Massachusetts. If $\mathcal{M}_1$ thinks there exist hackers at MIT, certainly there exist hackers in Massachusetts. Where things go wrong is something like:

$$\mathcal{M}_2 \models “\exists x : x \text{ is a course numbered } > 50”.$$

This is true for $\mathcal{M}_2$ because we can take the witness $x = \text{Math 55}$, say. But it’s false for $\mathcal{M}_1$, because at MIT all courses are numbered 18.701 or something similar.

The issue is that the *witness* for statements in $\mathcal{M}_2$ do not necessarily propagate down to witnesses for $\mathcal{M}_1$.

The Tarski-Vaught test says this is the only impediment: if every witness in $\mathcal{M}_2$ can be replaced by one in $\mathcal{M}_1$ then $\mathcal{M}_1 \prec \mathcal{M}_2$.

Lemma 61.4.3 (Tarski-Vaught)

Let $\mathcal{M}_1 \subseteq \mathcal{M}_2$. Then $\mathcal{M}_1 \prec \mathcal{M}_2$ if and only if: For every sentence $\phi(x, x_1, \dots, x_n)$ and parameters $b_1, \dots, b_n \in M_1$: if there is a witness $\tilde{b} \in M_2$ to $\mathcal{M}_2 \models \phi(\tilde{b}, b_1, \dots, b_n)$ then there is a witness $b \in M_1$ to $\mathcal{M}_1 \models \phi(b, b_1, \dots, b_n)$.

Proof. Easy after the above discussion. To formalize it, use induction on formula complexity. $\square$

§61.5 Obtaining the axioms of ZFC

We now want to write down conditions for M to satisfy ZFC axioms. The idea is that almost all the ZFC axioms are just Σ_1 claims about certain desired sets, and so verifying an axiom reduces to checking some appropriate “closure” condition: that the witness to the axiom is actually in the model.

For example, the EmptySet axiom is “ $\exists a(\text{isEmpty}(a))$ ”, and so we’re happy as long as $\emptyset \in M$, which is of course true for any nonempty transitive set M .

Lemma 61.5.1 (Transitive sets inheriting ZFC)

Let M be a nonempty transitive set. Then

- (i) M satisfies Extensionality, Foundation, EmptySet.
- (ii) $M \models \text{Pairing}$ if $x, y \in M \implies \{x, y\} \in M$.
- (iii) $M \models \text{Union}$ if $x \in M \implies \cup x \in M$.
- (iv) $M \models \text{PowerSet}$ if $x \in M \implies \mathcal{P}(x) \cap M \in M$.
- (v) $M \models \text{Replacement}$ if for every $x \in M$ and every function $F : x \rightarrow M$ which is M -definable with parameters, we have $F''x \in M$ as well.
- (vi) $M \models \text{Infinity}$ as long as $\omega \in M$.

Here, a set $X \subseteq M$ is **M -definable with parameters** if it can be realized as

$$X = \{x \in M \mid \phi[x, b_1, \dots, b_n]\}$$

for some (fixed) choice of parameters $b_1, \dots, b_n \in M$. We allow $n = 0$, in which case we say X is **M -definable without parameters**. Note that X need not itself be in M ! As a trivial example, $X = M$ is M -definable without parameters (just take $\phi[x]$ to always be true), and certainly we do not have $X \in M$.

Exercise 61.5.2. Verify (i)-(iv) above.

Remark 61.5.3. Converses to the statements of Lemma 61.5.1 are true for all claims other than (vii).

§61.6 Mostowski collapse

Up until now I have been only talking about transitive models, because they were easier to think about. Here's a second, better reason we might only care about transitive models.

Lemma 61.6.1 (Mostowski collapse lemma)

Let $X = (X, \in)$ be a model, where X is a set (possibly not transitive). Then there exists an isomorphism $\pi : X \rightarrow M$ for a transitive model $M = (M, \in)$.

This is also called the *transitive collapse*. In fact, both π and M are unique.

Proof. The idea behind the proof is very simple. Since $\in$ is well-founded and extensional, we can look at the $\in$ -minimal element $x_\emptyset$ of X with respect to $\in$. Clearly, we want to send that to $0 = \emptyset$.

Then we take the next-smallest set under $\in$, and send it to $1 = \{\emptyset\}$. We “keep doing this”; it's not hard to see this does exactly what we want.

To formalize, define π by transfinite recursion:

$$\pi(x) \stackrel{\text{def}}{=} \{\pi(y) \mid y \in x\}.$$

This π , by construction, does the trick. □

Remark 61.6.2 (Digression for experts). Earlier versions of Napkin claimed this was true for general models $\mathcal{X} = (X, E)$ with $\mathcal{X} \models \text{Foundation} + \text{Extensionality}$. This is false; it does not even imply E is well-founded, because there may be infinite descending chains of subsets of X which do not live in X itself. Another issue is that E may not be set-like.

The picture of this is quite “collapsing” the elements of M down to the bottom of V , hence the name.

§61.7 Adding an inaccessible

Prototypical example for this section: V_κ

At this point you might be asking, well, where's my model of ZFC?

I unfortunately have to admit now: ZFC can never prove that there is a model of ZFC (unless ZFC is inconsistent, but that would be even worse). This is a result called Gödel's incompleteness theorem.

Nonetheless, with some very modest assumptions added, we can actually show that a model *does* exist: for example, assuming that there exists a strongly inaccessible cardinal κ would do the trick, V_κ will be such a model (Problem 61D*). Intuitively you can see why: κ is so big that any set of rank lower than it can't escape it even if we take their power sets, or any other method that ZFC lets us do.

More pessimistically, this shows that it's impossible to prove in ZFC that such a κ exists. Nonetheless, we now proceed under ZFC^+ for convenience, which adds the existence of such a κ as a final axiom. So we now have a model V_κ to play with. Joy!

Great. Now we do something *really* crazy.

Theorem 61.7.1 (Countable transitive model)

Assume ZFC^+ . Then there exists a transitive model X of ZFC such that X is a countable set.

Proof. Fasten your seat belts.

First, since ZFC^+ we can take $M = (V_\kappa, \in)$ as our model of ZFC . Start with the set $X_0 = \emptyset$. Then for every integer n , we do the following to get X_{n+1} .

- Start with X_{n+1} containing every element of X_n .
- Consider a formula $\phi(x, x_1, \dots, x_n)$ and $b_1, \dots, b_n$ in X_n . Suppose that M thinks there is an $b \in M$ for which

$$M \models \phi[b, b_1, \dots, b_n].$$

We then add in the element b to X_{n+1} .

- We do this for *every possible formula in the language of set theory*. We also have to put in *every possible set of parameters* from the previous set X_n .

At every step X_n is countable. Reason: there are countably many possible finite sets of parameters in X_n , and countably many possible formulas, so in total we only ever add in countably many things at each step. This exhibits an infinite nested sequence of countable sets

$$X_0 \subseteq X_1 \subseteq X_2 \subseteq \dots$$

None of these is a substructure of M , because each X_n relies on witnesses in X_{n+1} . So we instead *take the union*:

$$X = \bigcup_n X_n.$$

This satisfies the Tarski-Vaught test, and is countable.

There is one minor caveat: X might not be transitive. We don't care, because we just take its Mostowski collapse. $\square$

Please take a moment to admire how insane this is. It hinges irrevocably on the fact that there are countably many sentences we can write down.

Remark 61.7.2. This proof relies heavily on the Axiom of Choice when we add in the element b to X_{n+1} . Without Choice, there is no way of making these decisions all at once.

Usually, the right way to formalize the Axiom of Choice usage is, for every formula $\phi(x, x_1, \dots, x_n)$, to pre-commit (at the very beginning) to a function $f_\phi(x_1, \dots, x_n)$, such that given any $b_1, \dots, b_n$ $f_\phi(b_1, \dots, b_n)$ will spit out the suitable value of b (if one exists). Personally, I think this is hiding the spirit of the proof, but it does make it clear how exactly Choice is being used.

These f_ϕ 's have a name: **Skolem functions**.

The trick we used in the proof works in more general settings:

Theorem 61.7.3 (Downward Löwenheim-Skolem theorem)

Let $\mathcal{M} = (M, E)$ be a model, and $A \subseteq M$. Then there exists a set B (called the **Skolem hull** of A) with $A \subseteq B \subseteq M$, such that $(B, E) \prec \mathcal{M}$, and

$$|B| < \max\{\omega, |A|\}.$$

In our case, what we did was simply take A to be the empty set.

Question 61.7.4. Prove this. (Exactly the same proof as before.)

§61.8 FAQ's on countable models

The most common one is “how is this possible?”, with runner-up “what just happened”.

Let me do my best to answer the first question. It seems like there are two things running up against each other:

- (1) M is a transitive model of ZFC, but its universe is uncountable.
- (2) ZFC tells us there are uncountable sets!

(This has confused so many people it has a name, Skolem's paradox.)

The reason this works I actually pointed out earlier: *countability is not absolute, it is a Σ_1 notion.*

Recall that a set x is countable if *there exists* an injective map $x \hookrightarrow \omega$. The first statement just says that *in the universe V* , there is a injective map $F : M \hookrightarrow \omega$. In particular, for any $x \in M$ (hence $x \subseteq M$, since M is transitive), x is countable *in V* . This is the content of the first statement.

But for M to be a model of ZFC, M only has to think statements in ZFC are true. More to the point, the fact that ZFC tells us there are uncountable sets means

$$M \models \exists x \text{ uncountable.}$$

In other words,

$$M \models \exists x \forall f \text{ If } f : x \rightarrow \omega \text{ then } f \text{ isn't injective.}$$


The key point is the $\forall f$ searches only functions in our tiny model M . It is true that in the “real world” V , there are injective functions $f : x \rightarrow \omega$. But M has no idea they exist! It is a brain in a vat: M is oblivious to any information outside it.

So in fact, every ordinal which appears in M is countable in the real world. It is just not countable in M . Since $M \models \text{ZFC}$, M is going to think there is some smallest uncountable cardinal, say $\aleph_1^M$. It will be the smallest (infinite) ordinal in M with the property that there is no bijection *in the model M* between $\aleph_1^M$ and ω . However, we necessarily know that such a bijection is going to exist in the real world V .

Put another way, cardinalities in M can look vastly different from those in the real world, because cardinality is measured by bijections, which I guess is inevitable, but leads to chaos.

§61.9 Picturing inner models

Here is a picture of a countable transitive model M .



Note that M and V must agree on finite sets, since every finite set has a formula that can express it. However, past V_{ω} the model and the true universe start to diverge.

The entire model M is countable, so it only occupies a small portion of the universe, below the first uncountable cardinal $\aleph_1^V$ (where the superscript means “of the true universe V ”). The ordinals in M are precisely the ordinals of V which happen to live inside the model, because the sentence “ α is an ordinal” is absolute. On the other hand, M has only a portion of these ordinals, since it is only a lowly set, and a countable set at that. To denote the ordinals of M , we write On^M , where the superscript means “the ordinals as computed in M ”. Similarly, On^V will now denote the “set of true ordinals”.

Nonetheless, the model M has its own version of the first uncountable cardinal $\aleph_1^M$. In the true universe, $\aleph_1^M$ is countable (below $\aleph_1^V$), but the necessary bijection witnessing this might not be inside M . That’s why M can think $\aleph_1^M$ is uncountable, even if it is a countable cardinal in the original universe.

So our model M is a brain in a vat. It happens to believe all the axioms of ZFC, and so every statement that is true in M could conceivably be true in V as well. But M can’t see the universe around it; it has no idea that what it believes is the uncountable $\aleph_1^M$ is really just an ordinary countable cardinal.

§61.10 Problems to think about

Problem 61A*. Show that for any transitive model M , the set of ordinals in M is itself some ordinal.

Problem 61B[†]. Assume $\mathcal{M}_1 \subseteq \mathcal{M}_2$. Show that

(a) If ϕ is Δ_0 , then $\mathcal{M}_1 \models \phi[b_1, \dots, b_n] \iff \mathcal{M}_2 \models \phi[b_1, \dots, b_n]$.

(b) If ϕ is Σ_1 , then $\mathcal{M}_1 \models \phi[b_1, \dots, b_n] \implies \mathcal{M}_2 \models \phi[b_1, \dots, b_n]$.

(c) If ϕ is Π_1 , then $\mathcal{M}_2 \models \phi[b_1, \dots, b_n] \implies \mathcal{M}_1 \models \phi[b_1, \dots, b_n]$.

(This should be easy if you've understood the chapter.)



Problem 61C[†] (Reflection). Let κ be an inaccessible cardinal such that $|V_\alpha| < \kappa$ for all $\alpha < \kappa$. Prove that for any $\delta < \kappa$ there exists $\delta < \alpha < \kappa$ such that $V_\alpha \prec V_\kappa$; in other words, the set of α such that $V_\alpha \prec V_\kappa$ is *unbounded* in κ . This means that properties of V_κ reflect down to properties of V_α .




Problem 61D* (Inaccessible cardinal produce models). Let κ be an inaccessible cardinal. Prove that V_κ is a model of ZFC.

62 Forcing

We are now going to introduce Paul Cohen's technique of **forcing**, which we then use to break the Continuum Hypothesis.

Here is how it works. Given a transitive model M and a poset $\mathbb{P}$ inside it, we can consider a "generic" subset $G \subseteq \mathbb{P}$, where G is not in M . Then, we are going to construct a bigger universe $M[G]$ which contains both M and G . (This notation is deliberately the same as $\mathbb{Z}[\sqrt{2}]$, for example – in the algebra case, we are taking $\mathbb{Z}$ and adding in a new element $\sqrt{2}$, plus everything that can be generated from it.) By choosing $\mathbb{P}$ well, we can cause $M[G]$ to have desirable properties.

Picture:



The model M is drawn in green, and its extension $M[G]$ is drawn in red.

The models M and $M[G]$ will share the same ordinals, which is represented here as M being no taller than $M[G]$. But one issue with this is that forcing may introduce some new bijections between cardinals of M that were not there originally; this leads to the phenomenon called **cardinal collapse**: quite literally, cardinals in M will no longer be

cardinals in $M[G]$, and instead just an ordinal. This is because in the process of adjoining G , we may accidentally pick up some bijections which were not in the earlier universe. In the diagram drawn, this is the function f mapping ω to $\aleph_1^M$. Essentially, the difficulty is that “ κ is a cardinal” is a Π_1 statement.

In the case of the Continuum Hypothesis, we’ll introduce a $\mathbb{P}$ such that any generic subset G will “encode” $\aleph_2^M$ real numbers. We’ll then show cardinal collapse does not occur, meaning $\aleph_2^{M[G]} = \aleph_2^M$. Thus $M[G]$ will have $\aleph_2^{M[G]}$ real numbers, as desired.

§62.1 Setting up posets

Prototypical example for this section: Infinite Binary Tree

Let M be a transitive model of ZFC. Let $\mathbb{P} = (\mathbb{P}, \leq) \in M$ be a poset with a maximal element $1_{\mathbb{P}}$ which lives inside a model M . The elements of $\mathbb{P}$ are called **conditions**; because they will force things to be true in $M[G]$.


Definition 62.1.1. A subset $D \subseteq \mathbb{P}$ is **dense** if for all $p \in \mathbb{P}$, there exists a $q \in D$ such that $q \leq p$.

Examples of dense subsets include the entire $\mathbb{P}$ as well as any downwards “slice”.

Definition 62.1.2. For $p, q \in \mathbb{P}$ we write $p \parallel q$, saying “ p is **compatible** with q ”, if there is exists $r \in \mathbb{P}$ with $r \leq p$ and $r \leq q$. Otherwise, we say p and q are **incompatible** and write $p \perp q$.

Example 62.1.3 (Infinite binary tree)

Let $\mathbb{P} = 2^{<\omega}$ be the **infinite binary tree** shown below, extended to infinity in the obvious way:



- (a) The maximal element $1_{\mathbb{P}}$ is the empty string $\emptyset$.
- (b) $D = \{\text{all strings ending in } 001\}$ is an example of a dense set.
- (c) No two elements of $\mathbb{P}$ are compatible unless they are comparable.

Now, I can specify what it means to be “generic”.

Definition 62.1.4. A nonempty set $G \subseteq \mathbb{P}$ is a **filter** if

- (a) The set G is upwards-closed: $\forall p \in G (\forall q \geq p) (q \in G)$.
- (b) Any pair of elements in G is compatible.

We say G is **M -generic** if for all D which are *in the model* M , if D is dense then $G \cap D \neq \emptyset$.

Question 62.1.5. Show that if G is a filter then $1_{\mathbb{P}} \in G$.


Example 62.1.6 (Generic filters on the infinite binary tree)

Let $\mathbb{P} = 2^{<\omega}$. The generic filters on $\mathbb{P}$ are sets of the form

$$\{0, b_1, b_1b_2, b_1b_2b_3, \dots\}.$$

So every generic filter on $\mathbb{P}$ correspond to a binary number $b = 0.b_1b_2b_3\dots$.

It is harder to describe which reals correspond to generic filters, but they should really “look random”. For example, the set of strings ending in 011 is dense, so one should expect “011” to appear inside b , and more generally that b should contain every binary string. So one would expect the binary expansion of $\pi - 3$ might correspond to a generic, but not something like $0.010101\dots$. That’s why we call them “generic”.



Exercise 62.1.7. Verify that these are every generic filter $2^{<\omega}$ has the form above. Show that conversely, a binary number gives a filter, but it need not be generic.

Notice that if $p \geq q$, then the sentence $q \in G$ tells us more information than the sentence $p \in G$. In that sense q is a *stronger* condition. In another sense $1_{\mathbb{P}}$ is the weakest possible condition, because it tells us nothing about G ; we always have $1_{\mathbb{P}} \in G$ since G is upwards closed.

§62.2 More properties of posets

We had better make sure that generic filters exist. In fact this is kind of tricky, but for countable models it works:

Lemma 62.2.1 (Rasiowa-Sikorski lemma)

Suppose M is a *countable* transitive model of ZFC and $\mathbb{P}$ is a partial order. Then there exists an M -generic filter G .

Proof. Essentially, hit them one by one. Problem 62B. □

Fortunately, for breaking CH we would want M to be countable anyways.

The other thing we want to do to make sure we're on the right track is guarantee that a generic set G is not actually in M . (Analogy: $\mathbb{Z}[3]$ is a really stupid extension.) The condition that guarantees this is:

Definition 62.2.2. A partial order $\mathbb{P}$ is **splitting** if for all $p \in \mathbb{P}$, there exists $q, r \leq p$ such that $q \perp r$.

Example 62.2.3 (Infinite binary tree is (very) splitting)

The infinite binary tree is about as splitting as you can get. Given $p \in 2^{<\omega}$, just consider the two elements right under it.

Lemma 62.2.4 (Splitting posets omit generic sets)

Suppose $\mathbb{P}$ is splitting. Then if $F \subseteq \mathbb{P}$ is a filter such that $F \in M$, then $\mathbb{P} \setminus F$ is dense. In particular, if $G \subseteq \mathbb{P}$ is generic, then $G \notin M$.

Proof. Consider $p \notin \mathbb{P} \setminus F \iff p \in F$. Then there exists $q, r \leq p$ which are not compatible. Since F is a filter it cannot contain both; we must have one of them outside F , say q . Hence every element of $p \in \mathbb{P} \setminus (\mathbb{P} \setminus F)$ has an element $q \leq p$ in $\mathbb{P} \setminus F$. That's enough to prove $\mathbb{P} \setminus F$ is dense.

Question 62.2.5. Deduce the last assertion of the lemma about generic G . □

§62.3 Names, and the generic extension

We now define the *names* associated to a poset $\mathbb{P}$.

Definition 62.3.1. Suppose M is a transitive model of ZFC, $\mathbb{P} = (\mathbb{P}, \leq) \in M$ is a partial order. We define the hierarchy of **$\mathbb{P}$ -names** recursively by

$$\begin{aligned} \text{Name}_0 &= \emptyset \\ \text{Name}_{\alpha+1} &= \mathcal{P}(\text{Name}_\alpha \times \mathbb{P}) \\ \text{Name}_\lambda &= \bigcup_{\alpha < \lambda} \text{Name}_\alpha. \end{aligned}$$

Finally, $\text{Name} = \bigcup_\alpha \text{Name}_\alpha$ denote the class of all $\mathbb{P}$ -names.

(These Name_α 's are the analog of the V_α 's: each Name_α is just the set of all names with rank $\leq \alpha$.)

Definition 62.3.2. For a filter G , we define the **interpretation** of τ by G , denote τ^G , using the transfinite recursion

$$\tau^G = \{ \sigma^G \mid \langle \sigma, p \rangle \in \tau \text{ and } p \in G \}.$$

We then define the model

$$M[G] = \{ \tau^G \mid \tau \in \text{Name}^M \}.$$

In words, $M[G]$ is the interpretation of all the possible $\mathbb{P}$ -names (as computed by M).

You should think of a $\mathbb{P}$ -name as a “fuzzy set”. Here’s the idea. Ordinary sets are collections of ordinary sets, so fuzzy sets should be collections of fuzzy sets. These fuzzy sets can be thought of like the Ghosts of Christmases yet to come: they represent things that might be, rather than things that are certain. In other words, they represent the possible futures of $M[G]$ for various choices of G .

Every fuzzy set has an element $p \in \mathbb{P}$ pinned to it. When it comes time to pass judgment, we pick a generic G and filter through the universe of $\mathbb{P}$ -names. The fuzzy sets with an element of G attached to it materialize into the real world, while the fuzzy sets with elements outside of G fade from existence. The result is $M[G]$.

Example 62.3.3 (First few levels of the name hierarchy)

Let us compute

$$\begin{aligned} \text{Name}_0 &= \emptyset \\ \text{Name}_1 &= \mathcal{P}(\emptyset \times \mathbb{P}) \\ &= \{\emptyset\} \\ \text{Name}_2 &= \mathcal{P}(\{\emptyset\} \times \mathbb{P}) \\ &= \mathcal{P}(\{\langle \emptyset, p \rangle \mid p \in \mathbb{P}\}). \end{aligned}$$

Compare the corresponding von Neuman universe.

$$V_0 = \emptyset, V_1 = \{\emptyset\}, V_2 = \{\emptyset, \{\emptyset\}\}.$$

Example 62.3.4 (Example of an interpretation)

As we said earlier, $\text{Name}_1 = \{\emptyset\}$. Now suppose

$$\tau = \{\langle \emptyset, p_1 \rangle, \langle \emptyset, p_2 \rangle, \dots, \langle \emptyset, p_n \rangle\} \in \text{Name}_2.$$

Then

$$\tau^G = \{\emptyset \mid \langle \emptyset, p \rangle \in \tau \text{ and } p \in G\} = \begin{cases} \{\emptyset\} & \text{if some } p_i \in G \\ \emptyset & \text{otherwise.} \end{cases}$$

In particular, remembering that G is nonempty we see that

$$\{\tau^G \mid \tau \in \text{Name}_2\} = V_2^M.$$

In fact, this holds for any natural number n , not just 2.

So, $M[G]$ and M agree on finite sets.

Now, we want to make sure $M[G]$ contains the elements of M . To do this, we take advantage of the fact that $1_{\mathbb{P}}$ must be in G , and define for every $x \in M$ the set

$$\check{x} = \{\langle \check{y}, 1_{\mathbb{P}} \rangle \mid y \in x\}$$

by transfinite recursion. Basically, $\check{x}$ is just a copy of x where we add check marks and tag every element with $1_{\mathbb{P}}$.

Example 62.3.5

Compute $\check{0} = 0$ and $\check{1} = \{\langle \check{0}, 1_{\mathbb{P}} \rangle\}$. Thus

$$(\check{0})^G = 0 \quad \text{and} \quad (\check{1})^G = 1.$$

Question 62.3.6. Show that in general, $(\check{x})^G = x$. (Rank induction.)

However, we'd also like to cause G to be in $M[G]$. In fact, we can write down the name exactly:

$$\dot{G} = \{\langle \check{p}, p \rangle \mid p \in \mathbb{P}\}.$$

Question 62.3.7. Show that $(\dot{G})^G = G$.

Question 62.3.8. Verify that $M[G]$ is transitive: that is, if $\sigma^G \in \tau^G \in M[G]$, show that $\sigma^G \in M[G]$. (This is offensively easy.)

In summary,

$M[G]$ is a transitive model extending M (it contains G).

Moreover, it is reasonably well-behaved even if G is just a filter. Let's see what we can get off the bat.

Lemma 62.3.9 (Properties obtained from filters)

Let M be a transitive model of ZFC. If G is a filter, then $M[G]$ is transitive and satisfies Extensionality, Foundation, EmptySet, Infinity, Pairing, and Union.

This leaves PowerSet, Replacement, and Choice.

Proof. Hence, we get Extensionality and Foundation for free. Then Infinity and EmptySet follows from $M \subseteq M[G]$.

For Pairing, suppose $\sigma_1^G, \sigma_2^G \in M[G]$. Then

$$\sigma = \{\langle \sigma_1, 1_{\mathbb{P}} \rangle, \langle \sigma_2, 1_{\mathbb{P}} \rangle\}$$

works satisfies $\sigma^G = \{\sigma_1, \sigma_2\}$. (Note that we used $M \models \text{Pairing}$.) Union is left as a problem, which you are encouraged to try now. $\square$

Up to here, we don't need to know anything about when a sentence is true in $M[G]$; all we had to do was contrive some names like $\check{x}$ or $\{\langle \sigma_1, 1_{\mathbb{P}} \rangle, \langle \sigma_2, 1_{\mathbb{P}} \rangle\}$ to get the facts we wanted. But for the remaining axioms, we *are* going to need this extra power are true in $M[G]$. For this, we have to introduce the fundamental theorem of forcing.

§62.4 Fundamental theorem of forcing

The model M unfortunately has no idea what G might be, only that it is some generic filter.¹ Nonetheless, we are going to define a relation $\Vdash$, called the *forcing* relation. Roughly, we are going to write

$$p \Vdash \varphi(\sigma_1, \dots, \sigma_n)$$

where $p \in \mathbb{P}$, $\sigma_1, \dots, \sigma_n \in M[G]$, if and only if:

For *any* generic G , if $p \in G$, then $M[G] \models \varphi(\sigma_1^G, \dots, \sigma_n^G)$.

Note that $\Vdash$ is defined without reference to G : it is something that M can see. We say p **forces** the sentences $\varphi(\sigma_1, \dots, \sigma_n)$. And miraculously, we can define this relation in such a way that the converse is true: *a sentence holds if and only if some p forces it*.

Theorem 62.4.1 (Fundamental theorem of forcing)

Suppose M is a transitive model of ZF. Let $\mathbb{P} \in M$ be a poset, and $G \subseteq \mathcal{P}$ is an M -generic filter. Then,

- (1) Consider $\sigma_1, \dots, \sigma_n \in \text{Name}^M$, Then

$$M[G] \models \varphi[\sigma_1^G, \dots, \sigma_n^G]$$

if and only if there exists a condition $p \in G$ such that p *forces* the sentence $\varphi(\sigma_1, \dots, \sigma_n)$. We denote this by $p \Vdash \varphi(\sigma_1, \dots, \sigma_n)$.

- (2) This forcing relation is (uniformly) definable in M .

I'll tell you how the definition works in the next section.

§62.5 (Optional) Defining the relation

Here's how we're going to go. We'll define the most generous condition possible such that the forcing works in one direction ($p \Vdash \varphi(\sigma_1, \dots, \sigma_n)$ means $M[G] \models \varphi[\sigma_1^G, \dots, \sigma_n^G]$). We will then cross our fingers that the converse also works.

We proceed by induction on the formula complexity. It turns out in this case that the atomic formula (base cases) are hardest and themselves require induction on ranks.

For some motivation, let's consider how we should define $p \Vdash \tau_1 \in \tau_2$ given that we've already defined $p \Vdash \tau_1 = \tau_2$. We need to ensure this holds iff

$$\forall M\text{-generic } G \text{ with } p \in G : M[G] \models \tau_1^G \in \tau_2^G.$$

So it suffices to ensure that any generic $G \ni p$ hits a condition q which forces τ_1^G to *equal* a member τ^G of τ_2^G . In other words, we want to choose the definition of $p \Vdash \tau_1 \in \tau_2$ to hold if and only if

$$\{q \in \mathbb{P} \mid \exists \langle \tau, r \rangle \in \tau_2 (q \leq r \wedge q \Vdash (\tau = \tau_1))\}$$

¹You might say this is a good thing, here's why. We're trying to show that $\neg\text{CH}$ is consistent with ZFC, and we've started with a model M of the real universe V . But for all we know CH might be true in V (what if $V = L$?), in which case it would also be true of M .

Nonetheless, we boldly construct $M[G]$ an extension of the model M . In order for it to behave differently from M , it has to be out of reach of M . Conversely, if M could compute everything about $M[G]$, then $M[G]$ would have to conform to M 's beliefs.

That's why we worked so hard to make sure $G \in M[G]$ but $G \notin M$.

is dense below in p . In other words, if the set is dense, then the generic must hit q , so it must hit r , meaning that $\langle \tau_r \rangle \in \tau_2$ will get interpreted such that $\tau^G \in \tau_2^G$, and moreover the $q \in G$ will force $\tau_1 = \tau$.

Now let's write down the definition... In what follows, the $\Vdash$ omits the M and $\mathbb{P}$.

Definition 62.5.1. Let M be a countable transitive model of ZFC. Let $\mathbb{P} \in M$ be a partial order. For $p \in \mathbb{P}$ and $\varphi(\sigma_1, \dots, \sigma_n)$ a formula in LST, we write $\tau \Vdash \varphi(\sigma_1, \dots, \sigma_n)$ to mean the following, defined by induction on formula complexity plus rank.

(1) $p \Vdash \tau_1 = \tau_2$ means

(i) For all $\langle \sigma_1, q_1 \rangle \in \tau_1$ the set

$$D_{\sigma_1, q_1} \stackrel{\text{def}}{=} \{r \mid r \leq q_1 \rightarrow \exists \langle \sigma_2, q_2 \rangle \in \tau_2 (r \leq q_2 \wedge r \Vdash (\sigma_1 = \sigma_2))\}.$$

is dense in p . (This encodes " $\tau_1 \subseteq \tau_2$ ".)

(ii) For all $\langle \sigma_2, q_2 \rangle \in \tau_2$, the set D_{σ_2, q_2} defined similarly is dense below p .

(2) $p \Vdash \tau_1 \in \tau_2$ means

$$\{q \in \mathbb{P} \mid \exists \langle \tau, r \rangle \in \tau_2 (q \leq r \wedge q \Vdash (\tau = \tau_1))\}$$

is dense below p .

(3) $p \Vdash \varphi \wedge \psi$ means $p \Vdash \varphi$ and $p \Vdash \psi$.

(4) $p \Vdash \neg \varphi$ means $\forall q \leq p, q \nVdash \varphi$.

(5) $p \Vdash \exists x \varphi(x, \sigma_1, \dots, \sigma_n)$ means that the set

$$\{q \mid \exists \tau (q \Vdash \varphi(\tau, \sigma_1, \dots, \sigma_n))\}$$

is dense below p .

This is definable in M ! All we've referred to is $\mathbb{P}$ and names, which are in M . (Note that being dense is definable.) Actually, in parts (3) through (5) of the definition above, we use induction on formula complexity. But in the atomic cases (1) and (2) we are doing induction on the ranks of the names.

So, the construction above gives us one direction (I've omitted tons of details, but...).

Now, how do we get the converse: that a sentence is true if and only if something forces it? Well, by induction, we can actually show:

Lemma 62.5.2 (Consistency and Persistence)

We have

(1) (Consistency) If $p \Vdash \varphi$ and $q \leq p$ then $q \Vdash \varphi$.

(2) (Persistence) If $\{q \mid q \Vdash \varphi\}$ is dense below p then $p \Vdash \varphi$.

You can prove both of these by induction on formula complexity. From this we get:

Corollary 62.5.3 (Completeness)

The set $\{p \mid p \Vdash \varphi \text{ or } p \Vdash \neg\varphi\}$ is dense.

Proof. We claim that whenever $p \nVdash \varphi$ then for some $\bar{p} \leq p$ we have $\bar{p} \Vdash \neg\varphi$; this will establish the corollary.

By the contrapositive of the previous lemma, $\{q \mid q \Vdash \varphi\}$ is not dense below p , meaning for some $\bar{p} \leq p$, every $q \leq \bar{p}$ gives $q \nVdash \varphi$. By the definition of $p \Vdash \neg\varphi$, we have $\bar{p} \Vdash \neg\varphi$. $\square$

And this gives the converse: the M -generic G has to hit some condition that passes judgment, one way or the other. This completes the proof of the fundamental theorem.

§62.6 The remaining axioms**Theorem 62.6.1** (The generic extension satisfies ZFC)

Suppose M is a transitive model of ZFC. Let $\mathbb{P} \in M$ be a poset, and $G \subseteq \mathbb{P}$ is an M -generic filter. Then

$$M[G] \models \text{ZFC}.$$

Proof. We'll just do Comprehension, as the other remaining axioms are similar.

Suppose $\sigma^G, \sigma_1^G, \dots, \sigma_n^G \in M[G]$ are a set and parameters, and $\varphi(x, x_1, \dots, x_n)$ is an LST formula. We want to show that the set

$$A = \{x \in \sigma^G \mid M[G] \models \varphi[x, \sigma_1^G, \dots, \sigma_n^G]\}$$

is in $M[G]$; i.e. it is the interpretation of some name.

Note that every element of σ^G is of the form ρ^G for some $\rho \in \text{dom}(\sigma)$ (a bit of abuse here, σ is a bunch of pairs of names and p 's, and the domain is just the set of names). So by the fundamental theorem of forcing, we may write

$$A = \{\rho^G \mid \rho \in \text{dom}(\sigma) \text{ and } \exists p \in G (p \Vdash \rho \in \sigma \wedge \varphi(\rho, \sigma_1, \dots, \sigma_n))\}.$$

To show $A \in M[G]$ we have to write down a τ such that the name τ^G coincides with A . We claim that

$$\tau = \{\langle \rho, p \rangle \in \text{dom}(\sigma) \times \mathbb{P} \mid p \Vdash \rho \in \sigma \wedge \varphi(\rho, \sigma_1, \dots, \sigma_n)\}$$

is the correct choice. It's actually clear that $\tau^G = A$ by construction; the "content" is showing that τ is in actually a name of M , which follows from $M \models \text{Comprehension}$.

So really, the point of the fundamental theorem of forcing is just to let us write down this τ ; it lets us show that τ is in Name^M without actually referencing G . $\square$

§62.7 Problems to think about

Problem 62A. For a filter G and M a transitive model of ZFC, show that $M[G] \models \text{Union}$.

Problem 62B (Rasiowa-Sikorski lemma). Show that in a countable transitive model M of ZFC, one can find an M -generic filter on any partial order.

63 Breaking the continuum hypothesis


We now use the technique of forcing to break the Continuum Hypothesis by choosing a good poset $\mathbb{P}$. As I mentioned earlier, one can also build a model where the Continuum Hypothesis is true; this is called the *constructible universe*, i.e. the model $V = L$. However, I think it's more fun when things break...

§63.1 Adding in reals

Starting with a *countable* transitive model M .

We want to choose $\mathbb{P} \in M$ such that $(\aleph_2)^M$ many real numbers appear, and then worry about cardinal collapse later.

Recall the earlier situation where we set $\mathbb{P}$ to be the infinite complete binary tree; its nodes can be thought of as partial functions $n \rightarrow 2$ where $n < \omega$. Then G itself is a path down this tree; i.e. it can be encoded as a total function $G : \omega \rightarrow 2$, and corresponds to a real number.



We want to do something similar, but with ω_2 many real numbers instead of just one. In light of this, consider in M the poset

$$\mathbb{P} = \text{Add}(\omega_2, \omega) \stackrel{\text{def}}{=} (\{p : \omega_2 \times \omega \rightarrow 2, \text{dom}(p) < \omega\}, \supseteq).$$

These elements (conditions) are “partial functions”: we take some finite subset of $\omega \times \omega_2$ and map it into $2 = \{0, 1\}$. Moreover, we say $p \leq q$ if $\text{dom}(p) \supseteq \text{dom}(q)$ and the two functions agree over $\text{dom}(q)$.

Question 63.1.1. What is $1_{\mathbb{P}}$ here?

Exercise 63.1.2. Show that a generic G can be encoded as a function $\omega_2 \times \omega \rightarrow 2$.

Lemma 63.1.3 (G encodes distinct real numbers)

For $\alpha \in \omega_2$ define

$$G_\alpha = \{n \mid G(\alpha, n) = 0\} \in \mathcal{P}(\mathbb{N}).$$

Then $G_\alpha \neq G_\beta$ for any $\alpha \neq \beta$.

Proof. We claim that the set

$$D = \{q \mid \exists n \in \omega : q(\alpha, n) \neq q(\beta, n) \text{ are both defined}\}$$

is dense.

Question 63.1.4. Check this. (Use the fact that the domains are all finite.)

Since G is an M -generic it hits this dense set D . Hence $G_\alpha \neq G_\beta$. $\square$

Since $G \in M[G]$ and $M[G] \models \text{ZFC}$, it follows that each G_α is in $M[G]$. So there are at least $\aleph_2^M$ real numbers in M . We are done once we can show there is no cardinal collapse.

§63.2 The countable chain condition

It remains to show that with $\mathbb{P} = \text{Add}(\omega, \omega_2)$, we have that

$$\aleph_2^{M[G]} = \aleph_2^M.$$

In that case, since $M[G]$ will have $\aleph_2^M = \aleph_2^{M[G]}$ many reals, we will be done.

To do this, we'll rely on a combinatorial property of $\mathbb{P}$:

Definition 63.2.1. We say that $A \subseteq \mathcal{P}$ is a **strong antichain** if for any distinct p and q in A , we have $p \perp q$.

Example 63.2.2 (Example of an antichain)

In the infinite binary tree, the set $A = \{00, 01, 10, 11\}$ is a strong antichain (in fact maximal by inclusion).

This is stronger than the notion of “antichain” than you might be used to!¹ We don't merely require that every two elements are incomparable, but that they are in fact *incompatible*.

Question 63.2.3. Draw a finite poset and an antichain of it which is not strong.

Definition 63.2.4. A poset $\mathbb{P}$ has the **κ -chain condition** (where κ is a cardinal) if all strong antichains in $\mathbb{P}$ have size less than κ . The special case $\kappa = \aleph_1$ is called the **countable chain condition**, because it implies that every strong antichain is countable.

We are going to show that if the poset has the κ -chain condition then it preserves all cardinals greater than κ . In particular, the countable chain condition will show that $\mathbb{P}$ preserves all the cardinals. Then, we'll show that $\text{Add}(\omega, \omega_2)$ does indeed have this property. This will complete the proof.

We isolate a useful lemma:

¹In the context of forcing, some authors use “antichain” to refer to “strong antichain”. I think this is lame.

Lemma 63.2.5 (Possible values argument)

Suppose M is a transitive model of ZFC and $\mathbb{P}$ is a partial order such that $\mathbb{P}$ has the κ -chain condition in M . Let $X, Y \in M$ and let $f : X \rightarrow Y$ be some function in $M[G]$, but $f \notin M$.

Then there exists a function $F \in M$, with $F : X \rightarrow \mathcal{P}(Y)$ and such that for any $x \in X$,

$$f(x) \in F(x) \quad \text{and} \quad |F(x)|^M < \kappa.$$

What this is saying is that if f is some new function that's generated, M is still able to pin down the values of f to at most κ many values.

Proof. The idea behind the proof is easy: any possible value of f gives us some condition in the poset $\mathbb{P}$ which forces it. Since distinct values must have incompatible conditions, the κ -chain condition guarantees there are at most κ such values.

Here are the details. Let $\dot{f}$, $\check{X}$, $\check{Y}$ be names for f , X , Y . Start with a condition p such that p forces the sentence

“ $\dot{f}$ is a function from $\check{X}$ to $\check{Y}$ ”.

We'll work just below here.

For each $x \in X$, we can consider (using the Axiom of Choice) a maximal strong antichain $A(x)$ of incompatible conditions $q \leq p$ which forces $f(x)$ to equal some value $y \in Y$. Then, we let $F(x)$ collect all the resulting y -values. These are all possible values, and there are less than κ of them. $\square$

§63.3 Preserving cardinals

As we saw earlier, cardinal collapse can still occur. For the Continuum Hypothesis we want to avoid this possible, so we can add in $\aleph_2^M$ many real numbers and have $\aleph_2^{M[G]} = \aleph_2^M$. It turns out that to verify this, one can check a weaker result.

Definition 63.3.1. For M a transitive model of ZFC and $\mathbb{P} \in M$ a poset, we say $\mathbb{P}$ **preserves cardinals** if $\forall G \subseteq \mathbb{P}$ an M -generic, the model M and $M[G]$ agree on the sentence “ κ is a cardinal” for every κ . Similarly we say $\mathbb{P}$ **preserves regular cardinals** if M and $M[G]$ agree on the sentence “ κ is a regular cardinal” for every κ .

Intuition: In a model M , it's possible that two ordinals which are in bijection in V are no longer in bijection in M . Similarly, it might be the case that some cardinal $\kappa \in M$ is regular, but stops being regular in V because some function $f : \bar{\kappa} \rightarrow \kappa$ is cofinal but happened to only exist in V . In still other words, “ κ is a regular cardinal” turns out to be a Π_1 statement too.

Fortunately, each implies the other. We quote the following without proof.

Proposition 63.3.2 (Preserving cardinals $\iff$ preserving regular cardinals)

Let M be a transitive model of ZFC. Let $\mathbb{P} \in M$ be a poset. Then for any λ , $\mathbb{P}$ preserves cardinalities less than or equal to λ if and only if $\mathbb{P}$ preserves regular cardinals less than or equal to λ . Moreover the same holds if we replace “less than or equal to” by “greater than or equal to”.

Thus, to show that $\mathbb{P}$ preserves cardinality and cofinalities it suffices to show that $\mathbb{P}$ preserves regularity. The following theorem lets us do this:

Theorem 63.3.3 (Chain conditions preserve regular cardinals)

Let M be a transitive model of ZFC, and let $\mathbb{P} \in M$ be a poset. Suppose M satisfies the sentence “ $\mathbb{P}$ has the κ chain condition and κ is regular”. Then $\mathbb{P}$ preserves regularity greater than or equal to κ .

Proof. Use the Possible Values Argument. Problem 63A. $\square$

In particular, if $\mathbb{P}$ has the countable chain condition then $\mathbb{P}$ preserves *all* the cardinals (and cofinalities). Therefore, it remains to show that $\text{Add}(\omega, \omega_2)$ satisfies the countable chain condition.

§63.4 Infinite combinatorics

We now prove that $\text{Add}(\omega, \omega_2)$ satisfies the countable chain condition. This is purely combinatorial, and so we work briefly.

Definition 63.4.1. Suppose C is an uncountable collection of finite sets. C is a Δ -system if there exists a **root** R with the condition that for any distinct X and Y in C , we have $X \cap Y = R$.

Lemma 63.4.2 (Δ -System lemma)

Suppose C is an uncountable collection of finite sets. Then $\exists \bar{C} \subseteq C$ such that $\bar{C}$ is an uncountable Δ -system.

Proof. There exists an integer n such that C has uncountably many guys of length n . So we can throw away all the other sets, and just assume that all sets in C have size n .

We now proceed by induction on n . The base case $n = 1$ is trivial, since we can just take $R = \emptyset$. For the inductive step we consider two cases.

First, assume there exists an $a \in C$ contained in uncountably many $F \in C$. Throw away all the other guys. Then we can just delete a , and apply the inductive hypothesis.

Now assume that for every a , only countably many members of C have a in them. We claim we can even get a $\bar{C}$ with $R = \emptyset$. First, pick $F_0 \in C$. It's straightforward to construct an F_1 such that $F_1 \cap F_0 = \emptyset$. And we can just construct $F_2, F_3, \dots$ $\square$

Lemma 63.4.3

For all κ , $\text{Add}(\omega, \kappa)$ satisfies the countable chain condition.

Proof. Assume not. Let

$$\{p_\alpha : \alpha < \omega_1\}$$

be a strong antichain. Let

$$C = \{\text{dom}(p_\alpha) : \alpha < \omega_1\}.$$

Let $\bar{C} \subseteq C$ be such that $\bar{C}$ is uncountable, and $\bar{C}$ is a Δ -system which root R . Then let

$$B = \{p_\alpha : \text{dom}(p_\alpha) \in R\}.$$

Each $p_\alpha \in B$ is a function $p_\alpha : R \rightarrow \{0, 1\}$, so there are two that are the same. $\square$

Thus, we have proven that the Continuum Hypothesis cannot be proven in ZFC.

§63.5 Problems to think about

Problem 63A. Let M be a transitive model of ZFC, and let $\mathbb{P} \in M$ be a poset. Suppose M satisfies the sentence “ $\mathbb{P}$ has the κ chain condition and κ is regular”. Show that $\mathbb{P}$ preserves regularity greater than or equal to κ .

DRAFT (Evan Chen)
Updated February 15, 2017

XVII

Backmatter

A	Pedagogical comments and references	577
A.1	Basic algebra and topology	577
A.2	Second-year topics	578
A.3	Advanced topics	579
A.4	Further topics	580
B	Hints to selected problems	581
C	Sketches of selected solutions	588
D	Glossary of notations	602
D.1	General	602
D.2	Functions and sets	602
D.3	Abstract and linear algebra	603
D.4	Quantum computation	604
D.5	Topology and (complex) analysis	604
D.6	Category theory	606
D.7	Differential geometry	606
D.8	Algebraic number theory	607
D.9	Representation theory	608
D.10	Algebraic geometry	608
D.11	Set theory	609
E	Terminology on sets and functions	611
E.1	Sets	611
E.2	Functions	612
E.2.1	Injective / surjective / bijective functions	612
E.2.2	Images and pre-images	613
E.3	Equivalence relations	614

A Pedagogical comments and references

Here are some higher-level comments on the way specific topics were presented, as well as pointers to further reading.

§A.1 Basic algebra and topology

Linear algebra and multivariable calculus

Following the comments in Section 6.9, I think that most presentations of linear algebra and multivariable calculus are *immoral*, because they miss the two key ideas, namely:

- In linear algebra, we study *linear maps* between spaces.
- In calculus, we *approximate functions at points by linear functions*.

In particular, I believe linear algebra should *always* be taught before multivariable calculus. The fact that this is the opposite of what happens is a testament to how atrociously these subjects are usually taught.

In particular, I do not recommend most linear algebra or multivariable calculus books.

For linear algebra, I've heard that [Ax97] follows this approach, hence the appropriate name "Linear Algebra Done Right". I followed with heavy modifications the proceedings of Math 55a, see [Ga14].

For multivariable calculus and differential geometry, I found the notes [Sj05] to be unusually well written. I referred to it frequently while I was enrolled in Math 55b [Ga15].

General topology

My personal view on spaces is that every space I ever work with is either metrizable or is the Zariski topology.

I adopted the approach of [Pu02], using metric topology first. I find that metric spaces are far more intuitive, and are a much better way to get a picture of what open / closed / compact etc. sets look like. This is the approach history took; general topology grew out of metric topology.

In my opinion, it's a very bad idea to start any general topology class by defining what a general topological space is, because it doesn't communicate a good picture of open and closed sets to draw pictures of.

Groups and commutative algebra

I teach groups before commutative rings but might convert later. Rings have better examples, don't have the confusion of multiplicative notation for additive groups, and modding out by ideals is more intuitive.

There's a specific thing I have a qualm with in group theory: the way that the concept of a normal subgroup is introduced. Only [Go11] does something similar to what I do. Most other people simply *define* a normal subgroup N as one with gNg^{-1} and then proceed to define modding out, without taking the time to explain where this definition comes from. I remember distinctly this concept as the first time in learning math where

I didn't understand what was going on. Only in hindsight do I see where this definition came from; I tried hard to make sure my own presentation didn't have this issue.

I deliberately don't include a chapter on just commutative algebra; other than the chapter on rings and ideals. The reason is that I always found it easier to learn commutative algebra theorems on the fly, in the context of something like algebraic number theory or algebraic geometry. For example, I finally understand why radicals and the Nullstellensatz were important when I saw how they were used in algebraic geometry. Before then, I never understood why I cared about them.

§A.2 Second-year topics

Complex analysis

I picked the approach of presenting the Cauchy-Goursat theorem as given (rather than proving a weaker version by Stoke's theorem, or whatever), and then deriving the key result that holomorphic functions are analytic from it. I think this most closely mirrors the "real-life" use of complex analysis, i.e. the computation of contour integrals.

The main reference for this chapter was [Ya12], which I recommend.

Category theory

I enthusiastically recommend [Le14], from which my chapters are based, and which contains much more than I had time to cover.

You might try reading chapters 2-4 in reverse order though: I found that limits were much more intuitive than adjoints. But your mileage may vary.

The category theory will make more sense as you learn more examples of structures: it will help to have read, say, the chapters on groups, rings, and modules.

Quantum algorithms

The exposition given here is based off a full semester at MIT taught by Seth Lloyd, in 18.435J [Ll15]. It is written in a far more mathematical perspective.

I only deal with finite-dimensional Hilbert spaces, because that is all that is needed for Shor's algorithm, which is the point of this chapter. This is not an exposition intended for someone who wishes to seriously study quantum mechanics (though it might be a reasonable first read): the main purpose is to give students a little appreciation for what this "Shor's algorithm" that everyone keeps talking about is.

Representation theory

I staunchly support teaching the representation of algebras first, and then specializing to the case of groups by looking at $k[G]$. The primary influence for the chapters here is [Et11], and you might think of what I have here as just some selections from the first four chapters of this source.

Set theory

Set theory is far off the beaten path. The notes I have written are based off the class I took at Harvard College, Math 145a [Ko14].

My general impression is that the way I present set theory (trying to remain intuitive and informal in a logical minefield) is not standard. Possible other reference: [Mi14].

§A.3 Advanced topics

Algebraic topology

I cover the fundamental group π_1 first, because I think the subject is more intuitive this way. A possible reference in this topic is [Mu00]. Only later do I do the much more involved homology groups. The famous standard reference for algebraic topology is [Ha02], which is what almost everyone uses these days. But I also found [Ma13a] to be very helpful, particularly in the part about cohomology rings.

I don't actually do very much algebraic topology. In particular, I think the main reason to learn algebraic topology is to see the construction of the homology and cohomology groups from the chain complex, and watch the long exact sequence in action. The concept of a (co)chain complex comes up often in other contexts as well, like the cohomology of sheaves or Galois cohomology. Algebraic topology is by far the most natural one.

I use category theory extensively, because it makes life easier.

Algebraic number theory

I learned from [Og10], using [Le02] for the part about the Chebotarev density theorem.

When possible I try to keep the algebraic number theory chapter close at heart to an “olympiad spirit”. Factoring in rings like $\mathbb{Z}[i]$ and $\mathbb{Z}[\sqrt{-5}]$ is very much an olympiad-flavored topic at heart: one is led naturally to the idea of factoring in general rings of integers, around which the presentation is built. As a reward for the entire buildup, the exposition finishes with the application of the Chebotarev density theorem to IMO 2003, Problem 6.

Algebraic geometry

My preferred introduction to algebraic geometry is [Ga03]. I also referenced [Va15] frequently in writing these notes. Both sets of lecture notes are essentially self-contained.

I would like to confess now that I know relatively little algebraic geometry, and in my personal opinion the parts on algebraic geometry are the weakest part of the Napkin. This is reflected in my work here: in the entire set of notes I only barely finish defining a scheme, the first central definition of the subject.

Nonetheless, I will foolishly still make some remarks about my own studies. I think there are three main approaches to beginning the study of schemes:

- Only looking at affine and projective varieties, as part of an “introductory” class, typically an undergraduate course.
- Studying affine and projective varieties closely and using them as the motivating example of a *scheme*, and then developing algebraic geometry from there.
- Jumping straight into the definition of a scheme, as in the well-respected and challenging [Va15].

I have gone with the second approach, I think that if you don't know what a scheme is, then you haven't learned algebraic geometry. But on the other hand I think the definition of a scheme is difficult to digest without having a good handle first on varieties.

These opinions are based on my personal experience of having tried to learn the subject through all three approaches over a period of a year. Your mileage may vary.

§A.4 Further topics

Real analysis

I don't do very much analysis over $\mathbb{R}$ itself. If this interests you, I highly recommend [Pu02], who approaches the subject starting with metric topology, much in the same way the Napkin does. Pugh also takes the time to properly construct $\mathbb{R}$ from $\mathbb{Q}$, something which I don't do here.

Analytic number theory

I never had time to write up notes in Napkin for these. If you're interested though, I recommend [Hi13]. They are highly accessible and delightful to read. The only real prerequisites are a good handle on Cauchy's residue formula.

DRAFT (Evan Chen)
Updated February 15, 2017

B Hints to selected problems

- 1A.** Orders.
- 1B.** Copy the proof of Fermat's little theorem, using Lemma 1.2.5.)
- 1C.** For the former, decide where the isomorphism should send r and s , and the rest will follow through. For the latter, look at orders.
- 1D***. Generated groups.
- 1E[†]**. Use $n = |G|$.
- 1F.** Draw inspiration from D_6 .
- 2A.** $\pm x$ for good choices of $\pm$.
- 2B.** Note that $p\mathbb{Z}$ is closed for each p . If there were finitely many primes, then $\bigcup p\mathbb{Z} = \mathbb{Z} \setminus \{-1, 1\}$ would have to be closed; i.e. $\{-1, 1\}$ would be open, but all open sets here are infinite.
- 2C.** Use contradiction. Only take δ of the form $1/k$ ($k \in \mathbb{Z}$).
- 2D.** Project gaps onto the y -axis. Use the fact that uncountably many positive reals cannot have finite sum.
- 2E.** The balls at 0 should be of the form $n! \cdot \mathbb{Z}$.
- 3A.** Write it out: $\phi(ab) = \phi(a)\phi(b)$.
- 3B.** $\gcd(1000, 999) = 1$.
- 3D.** Look at the group of 2×2 matrices mod p with determinant ± 1 .
- 3E.** Find an example of order 8.
- 3F.** Get yourself a list of English homophones, I guess. Don't try too hard. Letter v is the worst; maybe *felt* = *veldt*?
- 4D.** Appeal to $\mathbb{Q}$.
- 4E.** Color a circle magenta if it contains p but not q and color a circle cyan if it contains q but not p . Color p itself magenta and q itself cyan as well. Finally, color a circle neon yellow if it contains both p and q .
- 5A.** $[0, 1]$ is compact.
- 5D[†]**. Mimic the proof of Theorem 5.2.2. The totally bounded condition lets you do Pigeonhole.
- 6A.** Interpret as $V \oplus V \rightarrow W$ for suitable V, W .
- 6B.** Plug in $y = -1, 0, 1$. Use dimensions of $\mathbb{R}[x]$.

6D*. Use the fact that the infinite chain of subspaces

$$\ker T \subseteq \ker T^2 \subseteq \ker T^3 \subseteq \dots$$

and the similar chain for $\operatorname{im} T$ must eventually stabilize (for dimension reasons.).

7A. The point is that

$$(v_1 + v_2) \wedge v_2 \cdots \wedge v_n = v_1 \wedge v_2 \cdots \wedge v_n + v_2 \wedge v_2 \cdots \wedge v_n$$

and the latter term is zero.

7B[†]. One solution is to just take a basis. Otherwise, interpret $T \otimes S \mapsto \operatorname{Tr}(T \circ S)$ as a linear map $(V^\vee \otimes W) \otimes (W^\vee \otimes V) \rightarrow k$, and verify that it is commutative.

7C[†]. Again one can just take a basis.

7D. Consider 1000×1000 matrix M with entries 0 on diagonal and ± 1 off-diagonal. Mod 2.

7E. Take bases, and do a fairly long calculation.

8A. This is actually immediate by taking any basis in which X is upper-triangular.

8B. Look at the trace of T .

8C. There is a family of solutions other than just $a = b = c = d$.

One can solve the problem using Cayley-Hamilton. A more “bare-hands” approach is to show the matrix is invertible (unless $a = b = c = d$) and then diagonalize the matrix as $M = \begin{bmatrix} s & -q \\ -r & p \end{bmatrix} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} \begin{bmatrix} p & q \\ r & s \end{bmatrix} = \begin{bmatrix} ps\lambda_1 - qr\lambda_2 & qs(\lambda_1 - \lambda_2) \\ pr(\lambda_2 - \lambda_1) & ps\lambda_2 - qr\lambda_1 \end{bmatrix}$.

9B. Dot products in $\mathbb{F}_2$.

9C*. Define it on simple tensors then extend linearly.

9D. $k = n^n$. Endow tensor products with an inner form. Note that “zero entry somewhere on its diagonal” is equivalent to the product of those entries being zero.

10A. Just apply Burnside’s lemma directly to get the answer of 198 (the relevant group is D_{14}).

10B. There are multiple ways to see this. One is to just do the algebraic manipulation. Another is to use Cayley’s theorem to embed G into a symmetric group.

10C. Double-count pairs (g, x) with $g \cdot x = x$.

10E[†]. Let G act on the left cosets $\{gH \mid g \in G\}$ by left multiplication: $g' \cdot gH = g'gH$. Consider the orbit $\mathcal{O}$ of the coset H . By the orbit-stabilizer theorem, $|\mathcal{O}|$ divides $|G|$. But $|\mathcal{O}| \leq p$ also. So either $\mathcal{O} = \{H\}$ or $\mathcal{O}$ contains all cosets. The first case is impossible.

11B. Count Sylow 8 and 7 groups and let them intersect.

11C. Construct a non-abelian group such that all elements have order three.

- 11D.** First, if G abelian it's trivial. Otherwise, let $Z(G)$ be the center of the group, which is always a normal subgroup of G . Do a mod p argument via conjugation (or use the class equation).
- 12B***. Just keep on adding in elements to get an ascending chain.
- 12E.** I think the result is true if you add the assumption A is Noetherian, so look for trouble by picking A not Noetherian.
- 13A[†].** In the structure theorem, $k/(s_i) \in \{0, k\}$.
- 13B[†].** By theorem $V \cong \bigoplus_i k[x]/(s_i)$ for some polynomials s_i . Write each block in the form described.
- 13C[†].** Copy the previous proof, except using the other form of the structure theorem. Since $k[x]$ is algebraically closed each p_i is a linear factor.
- 14A***. Look at the Taylor series of f , and use Cauchy's differentiation formula to show that each of the larger coefficients must be zero.
- 14B***. Proceed by contradiction, meaning there exists a sequence $z_1, z_2, \dots \rightarrow z$ where $0 = f(z_1) = f(z_2) = \dots$ all distinct. Prove that $f = 0$ on some small neighborhood and apply the identity theorem.
- 14C***. Take the interior of the agreeing points; show that this set is closed, which implies the conclusion.
- 14E.** Liouville. Look at $\frac{1}{f(z)-w}$.
- 15C.** It's $\lim_{a \rightarrow \infty} \int_{-a}^a \frac{\cos x}{x^2+1} dx$. For each a , construct a semicircle.
- 17A.** Rewrite $|\Psi_{-}\rangle = -\frac{1}{\sqrt{2}}(|\rightarrow\rangle_A \otimes |\leftarrow\rangle_B - |\leftarrow\rangle_A \otimes |\rightarrow\rangle_B)$.
- 17B.** $-1, 1, 1, 1$. When we multiply them all together, we get that $\text{id}^A \otimes \text{id}^B \otimes \text{id}^C$ has measurement -1 , which is the paradox. What this means is that the values of the measurements are created when we make the observation, and not prepared in advance.
- 18A.** One way is to create CCNOT using a few Fredkin gates.
- 18B.** Plug in $|\psi\rangle = |0\rangle, |\psi\rangle = |1\rangle, |\psi\rangle = |\rightarrow\rangle$ and derive a contradiction.
- 18C.** First show that the box sends $|x_1\rangle \otimes \dots \otimes |x_m\rangle \otimes |\leftarrow\rangle$ to $(-1)^{f(x_1, \dots, x_m)}(|x_1\rangle \otimes \dots \otimes |x_m\rangle \otimes |\leftarrow\rangle)$.
- 18D[†].** This is direct computation.
- 20C[†].** Prove and use the fact that a quotients of compact spaces remain compact.
- 24A.** The category $\mathcal{A} \times \mathbf{2}$ has "redundant arrows".
- 26B.** Simply induct, with the work having been done on the $k = 2$ case.
- 27B.** This is just a summation. You will need the fact that mixed partials are symmetric.
- 28A[†].** Direct application of Stokes' theorem to $\alpha = f dx + g dy$.
- 28B.** This is just an exercises in sigma notation.

- 28C.** This is a straightforward (but annoying) computation.
- 28D.** We would want $\alpha_p(v) = \|v\|$.
- 28E.** Show that $d^2 = 0$ implies $\int_{\partial c} \alpha = 0$ for exact α . Draw an annulus.
- 30A.** Take the $n - 1$ st homology groups.
- 30B.** Build F as follows: draw the ray from x through $f(x)$ and intersect it with the boundary S^{n-1} .
- 31A.** Induction on m , using hemispheres.
- 31B.** One strategy is induction on p , with base case $p = 1$. Another strategy is to let U be the desired space and let V be the union of p non intersecting balls.
- 31C***. Use Theorem 31.2.5. Note that $\mathbb{R}^n \setminus \{0\}$ is homotopy equivalent to S^{n-1} .
- 31D***. Find a new short exact sequence to apply Theorem 31.2.1 to.
- 31E***. It's possible to use two cylinders with U and V . This time the matrix is $\begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$ or some variant though; in particular, it's injective, so $\tilde{H}_2(X) = 0$.
- 32B.** Use Theorem 31.2.5.
- 33A[†].** $\mathbb{CP}^n$ has no cells in adjacent dimensions, so all d_k maps must be zero.
- 33B.** The space $S^n - \{x_0\}$ is contractible.
- 33D.** You won't need to refer to any elements. Start with
- $$H_2(X) \cong H_2(X^3) \cong H_2(X^2) / \ker [H_2(X^2) \rightarrow H_2(X^3)],$$
- say. Take note of the marked injective and surjective arrows.
- 33E[†].** There is one cell of each dimension. Show that the degree of d_k is $\deg(\text{id}) + \deg(-\text{id})$, hence d_k is zero or $\cdot 2$ depending on whether k is even or odd.
- 35A[†].** Write $H^k(M; \mathbb{Z})$ in terms of $H_k(M)$ using the UCT, and analyze the ranks.
- 35B.** Use the previous result on Betti numbers.
- 35C.** Use the $\mathbb{Z}_2$ cohomologies, and find the cup product.
- 35D.** Assume that $r : S^m \times S^n \rightarrow S^m \wedge S^n$ is such a map. Show that the induced map $H^\bullet(S^m \wedge S^n; \mathbb{Z}) \rightarrow H^\bullet(S^m \times S^n; \mathbb{Z})$ between their cohomology rings is monic (since there exists an inverse map i).
- 36A***. The norm is multiplicative and equal to product of Galois conjugates.
- 36B.** Taking the standard norm on $\mathbb{Q}(\sqrt{2})$ will destroy it.
- 36C.** View as roots of unity. Note $\frac{1}{2}$ isn't an algebraic integer.
- 36D.** Norm in $\mathbb{Q}(\sqrt[3]{2})$.
- 36E***. $\left| \frac{1}{n}(\varepsilon_1 + \cdots + \varepsilon_n) \right| \leq 1$.

- 36F.** Pick a $\mathbb{Q}$ -basis of $\alpha_1, \dots, \alpha_n$ of K and WLOG the α_i are in $\mathcal{O}_K$ by scaling. Look at $N(c_1\alpha_1 + \dots + c_n\alpha_n)$. Bound denominators.
- 36G[†].** Obviously $\mathbb{Z}[\zeta_p] \subseteq \mathcal{O}_K$, so our goal is to show the reverse inclusion. Show that for any $\alpha \in \mathcal{O}_K$, the trace of $\alpha(1-\zeta_p)$ is divisible by p . Given $x = a_0 + a_1\zeta_p + \dots + a_{p-2}\zeta_p^{p-2} \in \mathcal{O}_K$ (where $a_i \in \mathbb{Q}$), consider $(1-\zeta_p)x$.
- 36H.** Let $\alpha = \alpha_1, \alpha_2, \dots, \alpha_n$ be its conjugates. Look at the polynomial $(x-\alpha_1^e) \dots (x-\alpha_n^e)$ across $e \in \mathbb{N}$. Pigeonhole principle on all possible polynomials.
- 37C.** Copy the proof of the usual Fermat's little theorem.
- 37D[†].** Clear denominators!
- 37E.** (a) is straightforward. For (b) work mod p . For (c) use norms.
- 38A.** Repeat the previous procedure.
- 38B.** You should get a group of order three.
- 38C.** Mimic the proof of part (a) of Minkowski's theorem.
- 38D.** Linear algebra.
- 38E.** Factor in $\mathbb{Q}(i)$.
- 38F.** Factor p , show that the class group of $\mathbb{Q}(\sqrt{-5})$ has order two.
- 39A^{*}.** Direct linear algebra computation.
- 39B^{*}.** Let M be the "embedding" matrix. Look at $M^\top M$, where $M^\top$ is the transpose matrix.
- 39C^{*}.** Vandermonde matrices.
- 39D.** $M_K \geq 1$ must hold. Bash.
- 41A^{*}.** Look at the image of ζ_p .
- 41B.** Repeated quadratic extensions have degree 2, so one can only get powers of two.
- 41C.** Left-hand side has minimal polynomial of degree 7, but the right-hand side lives in a degree six extension.
- 41D.** Hint: $\sigma(x^2) = \sigma(x)^2 \geq 0$ plus Cauchy's Functional Equation.
- 41E.** By induction, suffices to show $\mathbb{Q}(\alpha, \beta) = \mathbb{Q}(\gamma)$ for some γ in terms of α and β . For all but finitely many rational λ , the choice $\gamma = \alpha + \lambda\beta$ will work.
- 43A[†].** Show that no rational prime p can remain inert if $\text{Gal}(K/\mathbb{Q})$ is not cyclic.
- 44A.** Modify the end of the proof of quadratic reciprocity.
- 44C[†].** Chebotarev Density on $\mathbb{Q}(\zeta_m)$.
- 44E.** By primitive roots, it's the same as the action of $\times 3$ on $\mathbb{Z}_{p-1}$. Let ζ be a $(p-1)$ st root of unity. Take $d = \prod_{i < j} (\zeta^i - \zeta^j)$, think about $\mathbb{Q}(d)$, and figure out how to act on it by $x \mapsto x^3$.

- 45A[†]**. Pick m so that $\mathfrak{f}(L/\mathbb{Q}) \mid m\infty$.
- 45B[†]**. Apply the Takagi existence theorem with $\mathfrak{m} = 1$.
- 46B[†]**. For any $a \in A$, the map $v \mapsto a \cdot v$ is intertwining.
- 46C^{*}**. For part (b), pick a basis and do $T \mapsto (T(e_1), \dots, T(e_n))$.
- 46D^{*}**. Right multiplication.
- 46E**. Apply Problem 6D^{*}.
- 47A**. They are all one-dimensional, n of them. What are the homomorphisms $\mathbb{Z}_n \rightarrow \mathbb{C}^\times$?
- 47B**. The span of $(1, 0)$ is a subrepresentation.
- 47C**. This is actually easy.
- 47D**. There are only two one-dimensional ones (corresponding to the only two homomorphisms $D_{10} \rightarrow \mathbb{C}^\times$). So the remaining ones are two-dimensional.
- 48A[†]**. Obvious. Let $W = \bigoplus V_i^{m_i}$ (possible since $\mathbb{C}[G]$ semisimple) thus $\chi_W = \sum_i m_i \chi_{V_i}$.
- 48B**. Use the previous problem, with $\chi_W = \chi_{\text{refl}_0}^2$.
- 48C**. Characters. Note that $|\chi_W| = 1$ everywhere.
- 48D**. There are five conjugacy classes, $1, -1$ and $\pm i, \pm j, \pm k$. Given four of the representations, orthogonality can give you the fifth one.
- 48E^{*}**. Write as
- $$\sum_{i=1}^r \chi_{V_i \otimes V_i^\vee}(gh^{-1}) = \chi_{\bigoplus_i V_i \otimes V_i^\vee}(gh^{-1}) = \chi_{\mathbb{C}[G]}(gh^{-1}).$$
- Now look at the usual basis for $\mathbb{C}[G]$.
- 50A**. Squares.
- 50B**. Use the weak Nullstellensatz on $n+1$ dimensions. Given f vanishing on everything, consider $x_{n+1}f - 1$.
- 53B**. You will need to know about complex numbers in Euclidean geometry to solve this problem.
- 54B[†]**. Use the standard affine charts.
- 54C**. Examine the global regular functions.
- 55B^{*}**. We have $\mathcal{F}_p \cong \mathcal{F}(\{p\})$. If X is discrete, then all sequences of germs are compatible.
- 56A**. Galois conjugates.
- 56B[†]**. Appeal to Problem 55B^{*}.
- 56C**. Use the proof that $\text{AffSch} \simeq \text{CRing}$.
- 56D**. Show that if $\text{Spec } R$ is connected and its stalks are irreducible, then $\text{Spec } R$ is itself irreducible. Consider nilradical $N = \sqrt{(0)}$.
- 60A**. $\sup_{k \in \omega} |V_k|$.

- 60B.** Rearrange the cofinal maps to be nondecreasing.
- 61C[†].** This is very similar to the proof of Löwenheim-Skolem. For a sentence ϕ , let f_ϕ send α to the least $\beta < \kappa$ such that for all $\vec{b} \in M_\alpha$, if there exists $a \in M$ such that $V_\kappa \models \phi[a, \vec{b}]$ then $\exists a \in M_\beta$ such that $V_\kappa \models \phi[a, \vec{b}]$. (To prove this β exists, use the fact that κ is cofinal.) Then, take the supremum over the countably many sentences for each α .
- 62B.** Let $D_1, D_2, \dots$ be the dense sets (there are countably many of them).
- 63A.** Assume not, and take $\lambda > \kappa$ regular in M ; if $f : \bar{\lambda} \rightarrow \lambda$, use the Possible Values Argument on f to generate a function in M that breaks cofinality of λ .

DRAFT (Evan Chen)
Updated February 15, 2017



Sketches of selected solutions

- 1A.** The point is that $\heartsuit$ is a group, $G \subsetneq \heartsuit$ a subgroup and $G \cong \heartsuit$. This can only occur if $|\heartsuit| = \infty$; otherwise, a proper subgroup would have strictly smaller size than the original.
- 1B.** Let $\{g_1, g_2, \dots, g_n\}$ denote the elements of G . For any $g \in G$, this is the same as the set $\{gg_1, \dots, gg_n\}$. Taking the entire product and exploiting commutativity gives $g^n \cdot g_1 g_2 \dots g_n = g_1 g_2 \dots g_n$, hence $g^n = 1$.
- 1C.** One can check manually that $D_6 \cong S_3$, using the map $r \mapsto (1\ 2\ 3)$ and $s \mapsto (1\ 2)$. On the other hand D_{24} contains an element of order 12 while S_4 does not.
- 1D***. Let G be a group of order p , and $1 \neq g \in G$. Look at the group H generated by g and use Lagrange's theorem.
- 1E†.** The idea is that each element $g \in G$ can be thought of as a permutation $G \rightarrow G$ by $x \mapsto gx$.
- 1F.** We have $www = bb$, $bww = wb$, $wwb = bw$, $bwb = ww$. Interpret these as elements of D_6 .
- 2A.** Let $f(x) = x$ for $x \in \mathbb{Q}$ and $f(x) = -x$ for irrational x .
- 2C.** Assume for contradiction that there is a bad $\varepsilon_0 > 0$ meaning that for any δ , there is a $x \in M$ which is within δ of $p \in M$ but $f(x)$ is at least ε away from $f(p) \in N$. For $\delta = 1/k$ let x_k be the said counterexample. Then x_k converges to p (by triangle inequality) so $f(x_k)$ is supposed to converge to $f(p)$, which is impossible by construction since the $f(x_k)$ are at least ε_0 away from $f(p)$.
- 2D.** Assume for contradiction it is completely discontinuous; by scaling set $f(0) = 0$, $f(1) = 1$ and focus just on $f : [0, 1] \rightarrow [0, 1]$. Since it's discontinuous everywhere, for every $x \in [0, 1]$ there's an $\varepsilon_x > 0$ such that the continuity condition fails. Since the function is strictly increasing, that can only happen if the function misses all y -values in the interval $(fx - \varepsilon_x, fx)$ or $(fx, fx + \varepsilon_x)$ (or both). Projecting these missing intervals to the y -axis you find uncountably many intervals (one for each $x \in [0, 1]$) all of which are disjoint. In particular, summing the ε_x you get that a sum of uncountably many positive reals is 1. But in general it isn't possible for an uncountable family $\mathcal{F}$ of positive reals to have finite sum. Indeed, just classify the reals into buckets $\frac{1}{k} \leq x < \frac{1}{k-1}$. If the sum is actually finite then each bucket is finite, so the collection $\mathcal{F}$ must be countable, contradiction.
- 2E.** Let $d(x, y) = 2017^{-n}$, where n is the largest integer such that $n!$ divides $|x - y|$.
- 3A.** Abelian groups: $abab = a^2b^2 \iff ab = ba$.
- 3B.** $G/\ker G$ is isomorphic to a subgroup of H . The order of the former divides 1000; the order of the latter divides 999. This can only occur if $G/\ker G = \{1\}$ so $\ker G = G$.

3D. Look at the group G of 2×2 matrices mod p with determinant ± 1 (whose entries are the integers mod p). Let $g = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix}$ and then use $g^{|G|} = 1_G$.

3E. Quaternion group.

3F. <http://andrew.gibiansky.com/blog/linguistics/homophony-groups/>.

4D. You can pick a rational number in each interval and there are only countably many rational numbers. Done!

4E. Color a circle magenta if it contains p but not q and color a circle cyan if it contains q but not p . Color p itself magenta and q itself cyan as well. Finally, color a circle neon yellow if it contains both p and q . By repeating the argument in (1) there are no circles enclosing neither p nor q . Hence every point is either magenta, cyan, or neon yellow. Now note that given any magenta circle, its interior is completely magenta. Actually, the magenta circles can be totally ordered by inclusion (since they can't intersect). So we consider two cases:

- If there is a maximal magenta circle (i.e. a magenta circle not contained in any other magenta circle) then the set of all magenta points is just a closed disk.
- If there is no maximal magenta circle, then the set of magenta points can also be expressed as the union over all magenta circles of their interiors. This is a union of open sets, so it is itself open.

We conclude the set of magenta points is either a closed disk or an open set. Similarly for the set of cyan points. Moreover, the set of such points is convex.

To finish the problem:

- Suppose there are no neon yellow points. If the magenta points form a closed disk, then the cyan points are $\mathbb{R}^2$ minus a disk which is not convex. Contradiction. So the magenta points must be open. Similarly the cyan points must be open. But $\mathbb{R}^2$ is connected, so it can't be written as the union of two open sets.
- Now suppose there are neon yellow points. We claim there is a neon yellow circle minimal by inclusion. If not, then repeat the argument of (1) to get a contradiction, since any neon yellow circle must have diameter the distance from p to q . So we can find a neon yellow circle $\mathcal{C}$ whose interior is all magenta and cyan. Now repeat the argument of the previous part, replacing $\mathbb{R}^2$ by the interior of $\mathcal{C}$.

5A. Compactness is preserved under homeomorphism, so this is obvious. To be clear, suppose $f : [0, 1] \rightarrow (0, 1)$ is continuous and surjective. This would imply $(0, 1)$ is compact, which is false.

5C. Suppose $p_i = (x_i, y_i)$ is a sequence in $X \times Y$ ($i = 1, 2, \dots$). Looking on the X side, some subsequence converges: for the sake of illustration say it's $x_1, x_4, x_9, x_{16}, \dots \rightarrow x$. Then look at the corresponding sequence $y_1, y_4, y_9, y_{16}, \dots$. Using compactness of Y , it has a convergent subsequence, say $y_1, y_{16}, y_{81}, y_{256}, \dots \rightarrow y$. Then $p_1, p_{16}, p_{81}, \dots$ will converge to (x, y) .

It's interesting to watch students with little proof experience try this problem. They'll invariably conclude that (x_n) has a convergent subsequence and that (y_n)

does too. But these sequences could be totally unrelated. For this proof to work, you do need to apply compactness of X first, and then compactness of Y on the resulting *filtered* sequence like we did here.


6A. Let V be the space of real polynomials with degree at most $d/2$ (which has dimension $1 + \lfloor d/2 \rfloor$), and W be the space of real polynomials modulo P (which has dimension d). Then $\dim(V \oplus V) > \dim W$. So the linear map $V \oplus V \rightarrow W$ by $(A, B) \mapsto A + Q \cdot B$ has a kernel of positive dimension (by rank-nullity, for example).

6D*. Consider

$$\{0\} \subsetneq \ker S \subseteq \ker S^2 \subseteq \ker S^3 \subseteq \dots \text{ and } V \supsetneq \operatorname{im} S \supseteq \operatorname{im} S^2 \supseteq \operatorname{im} S^3 \supseteq \dots$$

For dimension reasons, these subspaces must eventually stabilize: for some large integer N , $\ker T^N = \ker T^{N+1} = \dots$ and $\operatorname{im} T^N = \operatorname{im} T^{N+1} = \operatorname{im} T^{N+2} = \dots$. When this happens, $\ker T^N \cap \operatorname{im} T^N = \{0\}$, since T^N is an automorphism of $\operatorname{im} T^N$. On the other hand, by Rank-Nullity we also have $\dim \ker T^N + \dim \operatorname{im} T^N = \dim V$. Thus for dimension reasons, $V = \ker T^N \oplus \operatorname{im} T^N$.

7B[†]. This amounts to drawing the diagram



It is easy to check that the center rectangle commutes, by check it on simple tensors $\xi_W \otimes v \otimes \xi_V \otimes w$. So the outer hexagon commutes and we're done. This is really the same as the proof with bases; what it amounts to is checking the assertion is true for matrices that have a 1 somewhere and 0 elsewhere, then extending by linearity.

7D. We proceed by contradiction. Let v be a vector of length 1000 whose entries are weight of cows. Assume the existence of a matrix M such that $Mv = 0$, with entries 0 on diagonal and ± 1 off-diagonal. But $\det M \pmod{2}$ is equal to the number of derangements of $\{1, \dots, 1000\}$, which is odd. Thus $\det M$ is odd and in particular not zero, so M is invertible. Thus $Mv = 0 \implies v = 0$, contradiction.

7E. Pick a basis $e_1, \dots, e_n$ of V . Let T have matrix (x_{ij}) , and let $m = \dim V$. Let δ_{ij} be the Kronecker delta. Also, let $\operatorname{Fix}(\sigma)$ denote the fixed points of a permutation σ and let $\operatorname{NoFix}(\sigma)$ denote the non-fixed points.

Expanding then gives

$$\begin{aligned}
& \det(a \cdot \text{id} - T) \\
&= \sum_{\sigma \in S_m} \left(\text{sign}(\sigma) \cdot \prod_{i=1}^m (a \cdot \delta_{i\sigma(i)} - x_{i\sigma(i)}) \right) \\
&= \sum_{s=0}^m \sum_{1 \leq i_1 < \dots < i_s \leq m} \sum_{\substack{\sigma \in S_m \\ \sigma \text{ fixes } i_k}} \left(\text{sign}(\sigma) \cdot \prod_{i=1}^m (a \cdot \delta_{i\sigma(i)} - x_{i\sigma(i)}) \right) \\
&= \sum_{s=0}^m \sum_{1 \leq i_1 < \dots < i_s \leq m} \sum_{\substack{\sigma \in S_m \\ \sigma \text{ fixes } (i_k)}} \left(\text{sign}(\sigma) \cdot \prod_{i \notin (i_k)} -x_{i\sigma(i)} \prod_{i \in (i_k)}^n (a \cdot -x_{ii}) \right) \\
&= \sum_{\sigma \in S_m} \left(\text{sign}(\sigma) \cdot \prod_{i \in \text{NoFix}(\sigma)} -x_{i\sigma(i)} \prod_{i \in \text{Fix}(\sigma)} (a - x_{ii}) \right) \\
&= \sum_{\sigma \in S_m} \left(\text{sign}(\sigma) \cdot \left(\prod_{i \in \text{NoFix}(\sigma)} -x_{i\sigma(i)} \right) \left(\sum_{t=0}^{|\text{Fix}(\sigma)|} a^{|\text{Fix}(\sigma)|-t} \cdot \sum_{i_1 < \dots < i_t \in \text{Fix}(\sigma)} \prod_{k=1}^t -x_{i_k i_k} \right) \right) \\
&= \sum_{\sigma \in S_m} \left(\text{sign}(\sigma) \left(\sum_{t=0}^{|\text{Fix}(\sigma)|} a^{m-t-|\text{Fix}(\sigma)|} \sum_{\substack{X \subseteq \{1, \dots, m\} \\ \text{NoFix}(\sigma) \subseteq X \\ X \text{ has exactly } t \text{ fixed}}} \prod_{i \in X} -x_{i\sigma(i)} \right) \right) \\
&= \sum_{n=0}^m a^{m-n} \left(\sum_{\sigma \in S_m} \text{sign}(\sigma) \sum_{\substack{X \subseteq \{1, \dots, m\} \\ \text{NoFix}(\sigma) \subseteq X \\ |X|=n}} \prod_{i \in X} -x_{i\sigma(i)} \right) \\
&= \sum_{n=0}^m a^{m-n} (-1)^n \left(\sum_{\substack{X \subseteq \{1, \dots, m\} \\ |X|=n}} \sum_{\substack{\sigma \in S_m \\ \text{NoFix}(\sigma) \subseteq X}} \text{sign}(\sigma) \prod_{i \in X} x_{i\sigma(i)} \right).
\end{aligned}$$

Hence it's the same to show that

$$\sum_{\substack{X \subseteq \{1, \dots, m\} \\ |X|=n}} \sum_{\substack{\sigma \in S_m \\ \text{NoFix}(\sigma) \subseteq X}} \text{sign}(\sigma) \prod_{i \in X} x_{i\sigma(i)} = \text{Tr}_{\Lambda^n(V)} (\Lambda^n(T))$$

holds for every n .

We can expand the definition of trace as using basis elements as

$$\begin{aligned}
 \text{Tr}(\Lambda^n(T)) &= \sum_{1 \leq i_1 < \dots < i_n \leq m} \left(\bigwedge_{k=1}^n e_{i_k} \right)^\vee \left(\Lambda^n(T) \left(\bigwedge_{k=1}^n e_{i_k} \right) \right) \\
 &= \sum_{1 \leq i_1 < \dots < i_n \leq m} \left(\bigwedge_{k=1}^n e_{i_k} \right)^\vee \left(\bigwedge_{k=1}^n T(e_{i_k}) \right) \\
 &= \sum_{1 \leq i_1 < \dots < i_n \leq m} \left(\bigwedge_{k=1}^n e_{i_k} \right)^\vee \left(\bigwedge_{k=1}^n \left(\sum_{j=1}^m x_{i_k j} e_j \right) \right) \\
 &= \sum_{1 \leq i_1 < \dots < i_n \leq m} \sum_{\pi \in S_n} \text{sign}(\pi) \prod_{k=1}^n x_{i_{\pi(k)} k} \\
 &= \sum_{\substack{X \subseteq \{1, \dots, m\} \\ |X|=n}} \sum_{\pi \in S_X} \text{sign}(\pi) \prod_{i \in X} x_{t\pi(i)}
 \end{aligned}$$

Hence it remains to show that the permutations over X are in bijection with the permutations over S_m which fix $\{1, \dots, m\} - X$, which is clear, and moreover, the signs clearly coincide.

8C. The answer is

$$\begin{bmatrix} t & t \\ t & t \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} -3t & -t \\ t & 3t \end{bmatrix}$$

for $t \in \mathbb{R}$. These work by taking $k = 3$.

Now to see these are the only ones, consider an arithmetic matrix

$$M = \begin{bmatrix} a & a+e \\ a+2e & a+3e \end{bmatrix}.$$

with $e \neq 0$. Its characteristic polynomial is $t^2 - (2a+3e)t - 2e^2$, with discriminant $(2a+3e)^2 + 8e^2$, so it has two distinct real roots; moreover, since $-2e^2 \leq 0$ either one of the roots is zero or they are of opposite signs. Now we can diagonalize M by writing

$$M = \begin{bmatrix} s & -q \\ -r & p \end{bmatrix} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} \begin{bmatrix} p & q \\ r & s \end{bmatrix} = \begin{bmatrix} ps\lambda_1 - qr\lambda_2 & qs(\lambda_1 - \lambda_2) \\ pr(\lambda_2 - \lambda_1) & ps\lambda_2 - qr\lambda_1 \end{bmatrix}$$

where $ps - qr = 1$. By using the fact the diagonal entries have sum equalling the off-diagonal entries, we obtain that

$$(ps - qr)(\lambda_1 + \lambda_2) = (qs - pr)(\lambda_1 - \lambda_2) \implies qs - pr = \frac{\lambda_1 + \lambda_2}{\lambda_1 - \lambda_2}.$$

Now if $M^k \in S$ too then the same calculation gives

$$qs - pr = \frac{\lambda_1^k + \lambda_2^k}{\lambda_1^k - \lambda_2^k}.$$

Let $x = \lambda_1/\lambda_2 < 0$ (since $-2e^2 < 0$). We appropriately get

$$\frac{x+1}{x-1} = \frac{x^k+1}{x^k-1} \implies \frac{2}{x-1} = \frac{2}{x^k-1} \implies x = x^k \implies x = -1 \text{ or } x = 0$$

and k odd. If $x = 0$ we get $e = 0$ and if $x = -1$ we get $2a + 3e = 0$, which gives the curve of solutions that we claimed.

A slicker approach is to note that by Cayley-Hamilton. Assume that $e \neq 0$, so M has two distinct real eigenvalues as above. We have $M^k = cM + d$ for some constants c and d (since M satisfies some quadratic polynomial). Since $M \in S$, $M^k \in S$ we obtain $d = 0$. Thus $M^k = cM$, so it follows the eigenvalues of M are negatives of each other. That means $\text{Tr } M = 0$, and the rest is clear.

9B. Interpret clubs as vectors in the vector space $\mathbb{F}_2^n$. Consider a “dot product” to show that all k vectors are linearly independent. Thus $k \leq \dim \mathbb{F}_2^n = n$.

9C*. The inner form given by

$$\langle v_1 \otimes w_1, v_2 \otimes w_2 \rangle_{V \otimes W} = \langle v_1, v_2 \rangle_V \langle w_1, w_2 \rangle_W$$

on pure tensors, then extending linearly. For (b) take $e_i \otimes f_j$ for $1 \leq i \leq n$, $1 \leq j \leq m$.

11B. If not, there exists eight Sylow 7-groups and seven Sylow 8-groups (since $n_7 \equiv 1 \pmod{7}$, $n_8 \mid 7$, and $n_7, n_8 > 1$). But no two of these $8 + 7 = 15$ groups can intersect at an element other than 1_G , which is clearly absurd.

11C. One example is upper triangular matrices with entries in mod 3.

11D. Let G be said group. If G is abelian then all subgroups are normal, and since G is simple, G can't have any subgroups at all. We can clearly find an element of order p , hence G has a subgroup of order p , which can only happen if $n = 1$, $G \cong \mathbb{Z}_p$.

Thus it suffices to show G can't be abelian. For this, we can use the class equation, but let's avoid that and do it directly:

Assume not and let $Z(G) = \{g \in G \mid xg = gx \forall x \in G\}$ be the center of the group. Since $Z(G)$ is normal in G , and G is simple, we see $Z(G) = \{1_G\}$. But then let G act on itself by conjugation: $g \cdot x = gxg^{-1}$. This breaks G into a bunch of orbits $\mathcal{O}_0 = \{1_G\}$, $\mathcal{O}_1$, $\mathcal{O}_2$, $\dots$, and since 1_G is the only fixed point by definition, all other orbits have size greater than 1. The Orbit-stabilizer theorem says that each orbit now has size dividing p^n , so they must all have size zero mod p .

But then summing across all orbits (which partition G), we obtain $|G| \equiv 1 \pmod{p}$, which is a contradiction.

12B*. For part (b), look at the poset of *proper* ideals. Apply Zorn's lemma (again using a union trick to verify the condition; be sure to verify that the union is proper!). In part (a) we are given no ascending infinite chains, so no need to use Zorn's lemma.

12E. Nope! Pick

$$\begin{aligned} A &= \mathbb{Z}[x_1, x_2, \dots] \\ B &= \mathbb{Z}[x_1, x_2, \dots, \varepsilon x_1, \varepsilon x_2, \dots] \\ C &= \mathbb{Z}[x_1, x_2, \dots, \varepsilon]. \end{aligned}$$

where $\varepsilon \neq 0$ but $\varepsilon^2 = 0$. I think the result is true if you add the assumption A is Noetherian.

- 14B***. Proceed by contradiction, meaning there exists a sequence $z_1, z_2, \dots \rightarrow z$ where $0 = f(z_1) = f(z_2) = \dots$ all distinct. WLOG set $z = 0$. Look at the Taylor series of f around $z = 0$. Since it isn't uniformly zero by assumption, write it as $a_N z^N + a_{N+1} z^{N+1} + \dots$, $a_N \neq 0$. But by continuity of $h(z) = a_N + a_{N+1} z + \dots$ there is some neighborhood of zero where $h(z) \neq 0$.
- 14C***. Let S be the interior of the points satisfying $f = g$. By definition S is open. By the previous part, S is closed: if $z_i \rightarrow z$ and $z_i \in S$, then $f = g$ in some neighborhood of z , hence $z \in S$. Since S is clopen and nonempty, $S = U$.
- 14E**. Suppose we want to show that there's a point in the image within ε of a given a point $w \in \mathbb{C}$. Look at $\frac{1}{f(z)-w}$ and use Liouville's theorem.
- 17A**. By a straightforward computation, we have $|\Psi_-\rangle = -\frac{1}{\sqrt{2}}(|\rightarrow\rangle_A \otimes |\leftarrow\rangle_B - |\leftarrow\rangle_A \otimes |\rightarrow\rangle_B)$. Now, $|\rightarrow\rangle_A \otimes |\rightarrow\rangle_B, |\rightarrow\rangle_A \otimes |\leftarrow\rangle_B$ span one eigenspace of $\sigma_x^A \otimes \text{id}_B$, and $|\leftarrow\rangle_A \otimes |\rightarrow\rangle_B, |\leftarrow\rangle_A \otimes |\leftarrow\rangle_B$ span the other. So this is the same as before: $+1$ gives $|\leftarrow\rangle_B$ and -1 gives $|\leftarrow\rangle_A$.
- 18A**. To show the Fredkin gate is universal it suffices to reversibly create a CCNOT gate with it. We write the system

$$\begin{aligned}(z, \neg z, -) &= \text{Fred}(z, 1, 0) \\ (x, a, -) &= \text{Fred}(x, 1, 0) \\ (y, b, -) &= \text{Fred}(y, a, 0) \\ (-, c, -) &= \text{Fred}(b, 0, 1) \\ (-, d, -) &= \text{Fred}(c, z, \neg z).\end{aligned}$$

Direct computation shows that $d = z + xy \pmod{2}$.

- 18C**. Put $|\leftarrow\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$. Then we have that U_f sends

$$|x_1\rangle \dots |x_m\rangle |0\rangle - |x_1\rangle \dots |x_m\rangle |1\rangle \xrightarrow{U_f} \pm |x_1\rangle \dots |x_m\rangle |0\rangle \mp |x_1\rangle \dots |x_m\rangle |1\rangle$$

the sign being $+, -$ exactly when $f(x_1, \dots, x_m) = 1$.

Now, upon inputting $|0\rangle \dots |0\rangle |1\rangle$, we find that $H^{\otimes m+1}$ maps it to

$$2^{-n/2} \sum_{x_1, \dots, x_n} |x_1\rangle \dots |x_n\rangle |\leftarrow\rangle.$$

Then the image under U_f is

$$2^{-n/2} \sum_{x_1, \dots, x_n} (-1)^{f(x_1, \dots, x_n)} |x_1\rangle \dots |x_n\rangle |\leftarrow\rangle.$$

We now discard the last qubit, leaving us with

$$2^{-n/2} \sum_{x_1, \dots, x_n} (-1)^{f(x_1, \dots, x_n)} |x_1\rangle \dots |x_n\rangle.$$

Applying $H^{\otimes m}$ to this, we get

$$2^{-n/2} \sum_{x_1, \dots, x_n} (-1)^{f(x_1, \dots, x_n)} \cdot \left(2^{-n/2} \sum_{y_1, \dots, y_n} (-1)^{x_1 y_1 + \dots + x_n y_n} |y_1\rangle |y_2\rangle \dots |y_n\rangle \right)$$

since $H|0\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ while $H|1\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$, so minus signs arise exactly if $x_i = 0$ and $y_i = 0$ simultaneously, hence the term $(-1)^{x_1 y_1 + \dots + x_n y_n}$. Swapping the order of summation, we get

$$2^{-n} \sum_{y_1, \dots, y_n} C(y_1, \dots, y_n) |y_1\rangle |y_2\rangle \dots |y_n\rangle$$

where $C_{y_1, \dots, y_n} = \sum_{x_1, \dots, x_n} (-1)^{f(x_1, \dots, x_n) + x_1 y_1 + \dots + x_n y_n}$. Now, we finally consider two cases.

- If f is the constant function, then we find that

$$C(y_1, \dots, y_n) = \begin{cases} \pm 1 & y_1 = \dots = y_n = 0 \\ 0 & \text{otherwise.} \end{cases}$$

To see this, note that the result is clear for $y_1 = \dots = y_n = 0$; otherwise, if WLOG $y_1 = 1$, then the terms for $x_1 = 0$ exactly cancel the terms for $x_1 = 1$, pair by pair. Thus in this state, the measurements all result in $|0\rangle \dots |0\rangle$.

- On the other hand if f is balanced, we derive that

$$C(0, \dots, 0) = 0.$$

Thus *no* measurements result in $|0\rangle \dots |0\rangle$.

In this way, we can tell whether f is balanced or not.

- 24A.** The main observation is that in $\mathcal{A} \times \mathbf{2}$, you have the arrows in $\mathcal{A}$ (of the form (f, id_2)), and then the arrows crossing the two copies of $\mathcal{A}$ (of the form $(\text{id}_A, 0 \leq 1)$). But there are some more arrows $(f, 0 \leq 1)$: nonetheless, they can be thought of as compositions

$$(f, 0 \leq 1) = (f, \text{id}_2) \circ (\text{id}_A, 0 \leq 1) = (\text{id}_A, 0 \leq 1) \circ (f, \text{id}_2).$$

Now we want to specify a functor $\alpha : \mathcal{A} \times \mathbf{2}$, we only have to specify where each of these two more basic things goes. The conditions on α already tells us that (f, id_2) should be mapped to $F(f)$ or $G(f)$ (depending on whether the arrow above is in $\mathcal{A} \times \{0\}$ or $\mathcal{A} \times \{1\}$), and specifying the arrow $(\text{id}_A, 0 \leq 1)$ amounts to specifying the A th component. Where does naturality come in?

The above discussion transfers to products of categories in general: you really only have to think about (f, id) and (id, g) arrows to get the general arrow $(f, g) = (f, \text{id}) \circ (\text{id}, g) = (\text{id}, g) \circ (f, \text{id})$.

- 30A.** Applying the functor H_{n-1} we get that the composition $\mathbb{Z} \rightarrow 0 \rightarrow \mathbb{Z}$ is the identity which is clearly not possible.
- 31B.** The answer is $\tilde{H}_{n-1}(X) \cong \mathbb{Z}^{\oplus p}$, with all other groups vanishing. For $p = 1$, $\mathbb{R}^n - \{*\} \cong S^{n-1}$ so we're done. For all other p , draw a hyperplane dividing the p points into two halves with a points on one side and b points on the other (so $a + b = p$). Set U and V and use induction.

Alternatively, let U be the desired space and let V be the union of p disjoint balls, one around every point. Then $U \cup V = \mathbb{R}^n$ has all reduced homology groups trivial. From the Mayer-Vietoris sequence we can read $\tilde{H}_k(U \cap V) \cong \tilde{H}_k(U) \cap \tilde{H}_k(V)$. Then $U \cap V$ is p punctured balls, which are each the same as S^{n-1} . One can read the conclusion from here.

31C*. It is $\mathbb{Z}$ for $k = n$ and 0 otherwise.

31D*. Use the short exact sequence $0 \rightarrow C_n(A, B) \rightarrow C_n(X, B) \rightarrow C_n(X, A) \rightarrow 0$.

32B. We have an exact sequence

$$\underbrace{\tilde{H}_1(\mathbb{R})}_{=0} \rightarrow \tilde{H}_1(\mathbb{R}, \mathbb{Q}) \rightarrow \tilde{H}_0(\mathbb{Q}) \rightarrow \underbrace{\tilde{H}_0(\mathbb{R})}_{=0}.$$

Now, since $\mathbb{Q}$ is path-disconnected (i.e. no two of its points are path-connected) it follows that $\tilde{H}_0(\mathbb{Q})$ consists of countably infinitely many copies of $\mathbb{Z}$.

33D. For concreteness, let's just look at the homology at $H_2(X^2, X^1)$ and show it's isomorphic to $H_2(X)$. According to the diagram

$$\begin{aligned} H_2(X) &\cong H_2(X^3) \\ &\cong H_2(X^2) / \ker [H_2(X^2) \rightarrow H_2(X^3)] \\ &\cong H_2(X^2) / \operatorname{im} \partial_3 \\ &\cong \operatorname{im} [H_2(X^2) \hookrightarrow H_2(X^2, X^1)] / \operatorname{im} \partial_3 \\ &\cong \ker(\partial_2) / \operatorname{im} \partial_3 \\ &\cong \ker d_2 / \operatorname{im} d_3. \end{aligned}$$

35D. See [Ma13a, Example 3.3.14, pages 68-69].

36F. Pick a $\mathbb{Q}$ -basis of $\alpha_1, \dots, \alpha_n$ of K and WLOG the α_i are in $\mathcal{O}_K$ by scaling. Pick $\alpha \in \mathcal{O}_K$. Look at $N(\alpha) = N(c_1\alpha_1 + \dots + c_n\alpha_n)$. If we do a giant norm computation, we find that $N(\alpha)$ is in the c_i with fixed coefficients. (For example, $N(c_1 + c_2\sqrt{2}) = c_1^2 - 2c_2^2$, say.) But $N(\alpha)$ is an *integer*, so the denominators of the c_i have to be bounded by some very large integer N . Thus

$$\bigoplus_i \mathbb{Z} \cdot \alpha_i \subseteq \mathcal{O}_K \subseteq \frac{1}{N} \bigoplus_i \mathbb{Z} \cdot \alpha_i.$$

The latter inclusion shows that $\mathcal{O}_K$ is a subgroup of a free group, and hence it is itself free. On the other hand, the first inclusion shows it's rank n .

37C. If $\alpha \equiv 0 \pmod{\mathfrak{p}}$ it's clear, so assume this isn't the case. Then $\mathcal{O}_K/\mathfrak{p}$ is a finite field with $N(\mathfrak{p})$ elements. Looking at $(\mathcal{O}_K/\mathfrak{p})^*$, it's a multiplicative group with $N(\mathfrak{p}) - 1$ elements, so $\alpha^{N(\mathfrak{p})-1} \equiv 1 \pmod{\mathfrak{p}}$, as desired.

37D[†]. Suppose it's generated by some elements in K ; we can write them as $\frac{\beta_i}{\alpha_i}$ for $\alpha_i, \beta_i \in \mathcal{A}$. Hence

$$J = \left\{ \sum_i \gamma_i \cdot \frac{\beta_i}{\alpha_i} \mid \alpha_i, \beta_i, \gamma_i \in \mathcal{O}_K \right\}.$$

Now “clear denominators”. Set $\alpha = \alpha_1 \dots \alpha_n$, and show that αI is an integral ideal.

37E. For part (a), note that the $\mathfrak{p}_i$ are prime just because

$$\mathcal{O}_K/\mathfrak{p}_i \cong (\mathbb{Z}[x]/f)/(p, f_i) \cong \mathbb{F}_p[x]/(f_i)$$

is a field, since the f_i are irreducible.

We check (b). Computing the product modulo p yields¹

$$\prod_{i=1}^r (f_i(\alpha))^{e_i} \equiv (f(\alpha)) \equiv 0 \pmod{p}$$

so we've shown that $I \subseteq (p)$.

Finally, we prove (c) with a size argument. The idea is that I and (p) really should have the same size; to nail this down we'll use the ideal norm. Since (p) divides I , we can write $(p) = \prod_{i=1}^r \mathfrak{p}_i^{e'_i}$ where $e'_i \leq e_i$ for each i . Remark $\mathcal{O}_K/(p) \cong \mathbb{Z}_p[x]/(f)$ has size $p^{\deg f}$. Similarly, $\mathcal{O}_K/(f_i)$ has degree $p^{\deg f_i}$ for each i . Compute $N((p))$ using the e'_i now and compare the results.

38F. Let $K = \mathbb{Q}(\sqrt{-5})$. Check that Cl_K has order two using the Minkowski bound; moreover $\Delta_K = 20$. Now note that $\mathcal{O}_K = \mathbb{Z}[\sqrt{-5}]$, and $x^2 + 5$ factors mod p as $(x+k)(x-k)$; hence in $\mathcal{O}_K$ we have $(p) = (p, \sqrt{-5} + k)(p, \sqrt{-5} - k) = \mathfrak{p}_1 \mathfrak{p}_2$, say. For $p > 5$ the prime p does not ramify and we have $\mathfrak{p}_1 \neq \mathfrak{p}_2$, since $\Delta_K = 20$.

Then $(p^2) = \mathfrak{p}_1^2 \cdot \mathfrak{p}_2^2$. Because the class group has order two, both ideals on the right hand side are principal, and because $\mathfrak{p}_1 \neq \mathfrak{p}_2$ they are distinct. Thus p^2 is a nontrivial product of two elements of $\mathcal{O}_K$; from this we can extract the desired factorization.

41A*. It's just $\mathbb{Z}_{p-1}$, since ζ_p needs to get sent to one (any) of the $p-1$ primitive roots of unity.

41E. <http://www.math.cornell.edu/~kbrown/6310/primitive.pdf>

44A. It is still true that

$$\left(\frac{2}{q}\right) = 1 \iff \sigma_2 \in H \iff 2 \text{ splits in } \mathbb{Z}\left[\frac{1}{2}(1 + \sqrt{q^*})\right].$$

Now, 2 splits in the ring if and only if $t^2 - t - \frac{1}{4}(1 - q^*)$ factors mod 2. This happens if and only if $q^* \equiv 1 \pmod{8}$. One can check this is exactly if $q \equiv \pm 1 \pmod{8}$, which gives the conclusion.

44C†. Let $K = \text{Gal}(\mathbb{Q}(\zeta_m)/\mathbb{Q})$. One can show that $\text{Gal}(K/\mathbb{Q}) \cong \mathbb{Z}_m^\times$ exactly as before. In particular, $\text{Gal}(K/\mathbb{Q})$ is abelian and therefore its conjugacy classes are singleton sets; there are $\phi(m)$ of them.

As long as p is sufficiently large, it is unramified and $\sigma_p = \text{Frob}_{\mathfrak{p}}$ for any $\mathfrak{p}$ above p (as m th roots of unity will be distinct modulo p ; differentiate $x^m - 1 \pmod{p}$ again).

44E. This solution is by David Corwin. By primitive roots, it's the same as the action of $\times 3$ on $\mathbb{Z}_{p-1}$. Let ζ be a $(p-1)$ st root of unity.

Consider

$$d = \prod_{0 \leq i < j < p-1} (\zeta^i - \zeta^j).$$

¹For example, suppose we want to know that $(3, 1 + \sqrt{7})(3, 1 - \sqrt{7})$ is contained in (3) . We could do the full computation and get $(9, 3 + 3\sqrt{7}, 3 - 3\sqrt{7}, 6)$. But if all we care about is that every element is divisible by 3, we could have just taken “mod 3” at the beginning and looked at just $(1 + \sqrt{7})(1 - \sqrt{7}) = (6)$; all the other products we get will obviously have factors of 3.

This is the square root of the discriminant of the polynomial $X^{p-1} - 1$; in other words $d^2 \in \mathbb{Z}$. In fact, by elementary methods one can compute

$$(-1)^{\binom{p-1}{2}} d^2 = -(p-1)^{p-1}$$

Now take the extension $K = \mathbb{Q}(d)$, noting that

- If $p \equiv 3 \pmod{4}$, then $d = (p-1)^{\frac{1}{2}(p-1)}$, so $K = \mathbb{Q}$.
- If $p \equiv 1 \pmod{4}$, then $d = i(p-1)^{\frac{1}{2}(p-1)}$, so $K = \mathbb{Q}(i)$.

Either way, in $\mathcal{O}_K$, let $\mathfrak{p}$ be a prime ideal above $(3) \subseteq \mathcal{O}_K$. Let $\sigma = \text{Frob}_{\mathfrak{p}}$ then be the unique element such that $\sigma(x) = x^3 \pmod{\mathfrak{p}}$ for all x . Then, we observe that

$$\sigma(d) \equiv \prod_{0 \leq i < j < p-1} (\zeta^{3i} - \zeta^{3j}) \equiv \begin{cases} +d & \text{if } \pi \text{ is even} \\ -d & \text{if } \pi \text{ is odd} \end{cases} \pmod{\mathfrak{p}}.$$

Now if $K = \mathbb{Q}$, then σ is the identity, thus σ even. Conversely, if $K = \mathbb{Q}(i)$, then 3 does not split, so $\sigma(d) = -d$ (actually σ is complex conjugation) thus π is odd.

Note the condition that $p \equiv 2 \pmod{3}$ is used only to guarantee that π is actually a permutation (and thus $d \neq 0$); it does not play any substantial role in the solution.

45A[†]. Suppose $\mathfrak{f}(L/\mathbb{Q}) \mid m\infty$ for some m . Then by the example from earlier we have the chain

$$P_{\mathbb{Q}}(m\infty) = H(\mathbb{Q}(\zeta)/\mathbb{Q}, m\infty) \subseteq H(L/\mathbb{Q}, m) \subseteq I_{\mathbb{Q}}(m\infty).$$

So by inclusion reversal we're done.

45B[†]. Apply the Takagi existence theorem with $\mathfrak{m} = 1$ to obtain an unramified extension E/K such that $H(E/K, 1) = P_K(1)$. We claim this works:

- To see it is maximal by inclusion, note that any other extension M/K with this property has conductor 1 (no primes divide the conductor), and then we have $P_K(1) = H(E/K, 1) \subseteq H(M/K, 1) \subseteq I_K(1)$, so inclusion reversal gives $M \subseteq E$.
- We have $\text{Gal}(L/K) \cong I_K(1)/P_K(1) = C_K(1)$ the class group.
- The isomorphism in the previous part is given by the Artin symbol. So $\mathfrak{p}$ splits completely if and only if $\left(\frac{L/K}{\mathfrak{p}}\right) = \text{id}$ if and only if $\mathfrak{p}$ is principal (trivial in $C_K(1)$).

This completes the proof.

46D^{*}. The operators are those of the form $T(a) = ab$ for some fixed $b \in A$. One can check these work, since for $c \in A$ we have $T(c \cdot a) = cab = c \cdot T(a)$. To see they are the only ones, note that $T(a) = T(a \cdot 1_A) = a \cdot T(1_A)$ for any $a \in A$.

47C. Pick any $v \in V$, then the subspace spanned by elements $g \cdot v$ for $v \in V$ is G -invariant; this is a finite-dimensional subspace, so it must equal all of V .

48B. $\mathbb{C}_{\text{sign}} \oplus \mathbb{C}^2 \oplus \text{refl}_0 \oplus (\text{refl}_0 \otimes \mathbb{C}_{\text{sign}})$.

48C. First, observe that $|\chi_W(g)| = 1$ for all $g \in G$.

$$\begin{aligned}\langle \chi_{V \otimes W}, \chi_{V \otimes W} \rangle &= \langle \chi_V \chi_W, \chi_V \chi_W \rangle \\ &= \frac{1}{|G|} \sum_{g \in G} |\chi_V(g)|^2 |\chi_W(g)|^2 \\ &= \frac{1}{|G|} \sum_{g \in G} |\chi_V(g)|^2 \\ &= \langle \chi_V, \chi_V \rangle = 1.\end{aligned}$$

48D. The table is given by

Q_8	1	-1	$\pm i$	$\pm j$	$\pm k$
$\mathbb{C}_{\text{triv}}$	1	1	1	1	1
$\mathbb{C}_i$	1	1	1	-1	-1
$\mathbb{C}_j$	1	1	-1	1	-1
$\mathbb{C}_k$	1	1	-1	-1	1
$\mathbb{C}^2$	2	-2	0	0	0

The one-dimensional representations (first four rows) follows by considering the homomorphism $Q_8 \rightarrow \mathbb{C}^\times$. The last row is two-dimensional and can be recovered by using the orthogonality formula.

50A. If $V = \mathcal{V}(I)$ with $I = (f_1, \dots, f_m)$ (as usual there are finitely many polynomials since $\mathbb{R}[x_1, \dots, x_n]$ is Noetherian) then we can take $f = f_1^2 + \dots + f_m^2$.

50B. The point is to check that if f vanishes on all of $\mathcal{V}(I)$, then $f \in \sqrt{I}$.

Take a set of generators $f_1, \dots, f_m$, in the original ring $\mathbb{C}[x_1, \dots, x_n]$; we may assume it's finite by the Hilbert basis theorem.

We're going to do a trick now: consider $S = \mathbb{C}[x_1, \dots, x_n, x_{n+1}]$ instead. Consider the ideal $I' \subseteq S$ in the bigger ring generated by $\{f_1, \dots, f_m\}$ and the polynomial $x_{n+1}f - 1$. The point of the last guy is that its zero locus does not touch our copy $x_{n+1} = 0$ of $\mathbb{A}^n$ nor any point in the "projection" of f through $\mathbb{A}^{n+1}$ (one can think of this as $\mathcal{V}(I)$ in the smaller ring direct multiplied with $\mathbb{C}$). Thus $\mathcal{V}(I') = \emptyset$, and by the weak Nullstellensatz we in fact have $I' = \mathbb{C}[x_1, \dots, x_{n+1}]$. So

$$1 = g_1 f_1 + \dots + g_m f_m + g_{m+1} (x_{n+1} f - 1).$$

Now the hack: **replace every instance of x_{n+1} by $\frac{1}{f}$** , and then clear all denominators. Thus for some large enough integer N we can get

$$f^N = f^N (g_1 f_1 + \dots + g_m f_m)$$

which eliminates any fractional powers of f in the right-hand side. It follows that $f^N \in I$.

53A. From the exactness, $h_I(d) = h_I(d - k) + h_{I+(f)}(d)$, and it follows that

$$\chi_{I+(f)}(d) = \chi_I(d) - \chi_I(d - k).$$

Let $m = \dim \mathcal{V}_{\text{pr}}(I) \geq 1$. Now $\dim \mathcal{V}_{\text{pr}}(I + (f)) = m - 1$, so and $c_{\text{new}} = \deg I + (f)$ then we have

$$\frac{\deg(I + (f))d^{m-1} + \dots}{(m-1)!} = \frac{1}{m!} (\deg I(d^m - (d-k)^m) + \text{lower order terms})$$

from which we read off

$$\deg(I + (f)) = \frac{(m-1)!}{m!} \cdot k \binom{m}{1} \deg I = k \deg I$$

as needed.

- 53B.** In complex numbers with ABC the unit circle, it is equivalent to solving the two cubic equations

$$(p-a)(p-b)(p-c) = (abc)^2(q-1/a)(q-1/b)(q-1/c)$$

$$0 = \prod_{\text{cyc}}(p+c-b-bcq) + \prod_{\text{cyc}}(p+b-c-bcq)$$

in p and $q = \bar{p}$. Viewing this as two cubic curves in $(p, q) \in \mathbb{C}^2$, by Bezout's theorem it follows there are at most nine solutions (unless both curves are not irreducible, but it's easy to check the first one cannot be factored). Moreover it is easy to name nine solutions (for ABC scalene): the three vertices, the three excenters, and I , O , H . Hence the answer is just those three triangle centers I , O and H .

- 54C.** If they were isomorphic, we would have $\mathcal{O}_V(V) \cong \mathcal{O}_W(W)$. For irreducible projective varieties, $\mathcal{O}_W(W) \cong \mathbb{C}$, while for affine varieties $\mathcal{O}_V(V) \cong \mathbb{C}[V]$. Thus we conclude V must be a single point.

- 56C.** $\mathfrak{p}$ gets sent to the characteristic of the field $\mathcal{O}_{X,\mathfrak{p}}/\mathfrak{m}_{X,\mathfrak{p}}$.

- 56D.** This is the proposition on the second page of http://www.acritch.com/media/math/Stalk-local_detection_of_irreducibility.pdf

- 61C[†].** For a sentence ϕ let

$$f_\phi : \kappa \rightarrow \kappa$$

send α to the least $\beta < \kappa$ such that for all $\vec{b} \in V_\alpha$, if there exists $a \in V_\kappa$ such that $V_\kappa \models \phi[a, \vec{b}]$ then $\exists a \in V_\beta$ such that $V_\kappa \models \phi[a, \vec{b}]$.

We claim this is well-defined. There are only $|V_\alpha|^n$ many possible choices of $\vec{b}$, and in particular there are fewer than κ of these (since we know that $|V_\alpha| < \kappa$; compare Problem 60C^{*}). Otherwise, we can construct a cofinal map from $|V_\alpha|^n$ into κ by mapping each vector $\vec{b}$ into a β for which the proposition fails. And that's impossible since κ is regular!

In other words, what we've done is fix ϕ and then use Tarski-Vaught on all the $\vec{b} \in V_\alpha^n$. Now let $g : \kappa \rightarrow \kappa$ be defined by

$$\alpha \mapsto \sup f_\phi(\alpha).$$

Since κ is regular and there are only countably many formulas, $g(\alpha)$ is well-defined.

Check that if α has the property that g maps α into itself (in other words, α is closed under g), then by the Tarski-Vaught test, we have $V_\alpha \prec V_\kappa$.

So it suffices to show there are arbitrarily large $\alpha < \kappa$ which are closed under g . Fix α_0 . Let $\alpha_1 = g(\alpha_0)$, et cetera and define

$$\alpha = \sup_{n < \omega} \alpha_n.$$

This α is closed under g , and by making α_0 arbitrarily large we can make α as large as we like.

62B. Since M is countable, there are only countably many dense sets (they live in M !), say

$$D_1, D_2, \dots, D_n, \dots \in M.$$

Using Choice, let $p_1 \in D_1$, and then let $p_2 \leq p_1$ such that $p_2 \in D_2$ (this is possible since D_2 is dense), and so on. In this way we can inductively exhibit a chain

$$p_1 \geq p_2 \geq p_3 \geq \dots$$

with $p_i \in D_i$ for every i .

Hence, we want to generate a filter from the $\{p_i\}$. Just take the upwards closure – let G be the set of $q \in \mathbb{P}$ such that $q \geq p_n$ for some n . By construction, G is a filter (this is actually trivial). Moreover, G intersects all the dense sets by construction.

63A. It suffices to show that $\mathbb{P}$ preserves regularity greater than or equal to κ . Consider $\lambda > \kappa$ which is regular in M , and suppose for contradiction that λ is not regular in $M[G]$. That's the same as saying that there is a function $f \in M[G]$, $f: \bar{\lambda} \rightarrow \lambda$ cofinal, with $\bar{\lambda} < \lambda$. Then by the Possible Values Argument, there exists a function $F \in M$ from $\bar{\lambda} \rightarrow \mathcal{P}(\lambda)$ such that $f(\alpha) \in F(\alpha)$ and $|F(\alpha)|^M < \kappa$ for every α .

Now we work in M again. Note for each $\alpha \in \bar{\lambda}$, $F(\alpha)$ is bounded in λ since λ is regular in M and greater than $|F(\alpha)|$. Now look at the function $\bar{\lambda} \rightarrow \lambda$ in M by just

$$\alpha \mapsto \cup F(\alpha) < \lambda.$$

This is cofinal in M , contradiction.

D Glossary of notations

§D.1 General

- $\forall$: for all
- $\exists$: there exists
- $\text{sign}(\sigma)$: sign of permutation σ
- $X \implies Y$: X implies Y

§D.2 Functions and sets

- $f(S)$ is the image of $f : X \rightarrow Y$ for $S \subseteq X$.
- $f^{-1}(y)$ is the inverse for $f : X \rightarrow Y$ when $y \in Y$.
- $f^{\text{pre}}(T)$ is the pre-image for $f : X \rightarrow Y$ when $T \subseteq Y$.
- $f|_S$ is the restriction of $f : X \rightarrow Y$ to $S \subseteq X$.
- f^n is the function f applied n times

Below are some common sets. These may also be thought of as groups, rings, fields etc. in the obvious way.

- $\mathbb{C}$: set of complex numbers
- $\mathbb{R}$: set of real numbers
- $\mathbb{N}$: set of positive integers
- $\mathbb{Q}$: set of rational numbers
- $\mathbb{Z}$: set of integers
- $\emptyset$: empty set

Some common notation with sets:

- $A \subset B$: A is any subset of B
- $A \subseteq B$: A is any subset of B
- $A \subsetneq B$: A is a *proper* subset of B
- $S \times T$: Cartesian product of sets S and T
- $S \setminus T$: difference of sets S and T
- $S \cup T$: set union of S and T
- $S \cap T$: set intersection of S and T

- $S \sqcup T$: disjoint union of S and T
- $|S|$: cardinality of S
- $S/\sim$: if $\sim$ is an equivalence relation on S , this is the set of equivalence classes
- $x + S$: denotes the set $\{x + s \mid s \in S\}$.
- xS : denotes the set $\{xs \mid s \in S\}$.

§D.3 Abstract and linear algebra

Some common groups/rings/fields:

- $\mathbb{Z}_n$: cyclic group of order n
- $\mathbb{Z}_n^\times$: set of units of $\mathbb{Z}_n$.
- S_n : symmetric group on $\{1, \dots, n\}$
- D_{2n} : dihedral group of order $2n$.
- $0, 1$: trivial group (depending on context)
- $\mathbb{F}_p$: integers modulo p

Notation with groups:

- 1_G : identity element of the group G
- $N \trianglelefteq G$: subgroup N is normal in G .
- G/N : quotient group of G by the normal subgroup N
- $Z(G)$: center of group G
- $N_G(H)$: normalizer of the subgroup H of G
- $G \times H$: product group of G and H
- $G \oplus H$: also product group, but often used when G and H are abelian (and hence we can think of them as $\mathbb{Z}$ -modules)
- $\text{Stab}_G(x)$: the stabilizer of $x \in X$, if X is acted on by G
- $\text{FixPt } g$, the set of fixed points by $g \in G$ (under a group action)

Notation with rings:

- R/I : quotient of ring R by ideal I
- $(a_1, \dots, a_n)$: ideal generated by the a_i
- $R^\times$: the group of units of R
- $R[x_1, \dots, x_n]$: polynomial ring in x_i , or ring obtained by adjoining the x_i to R
- $F(x_1, \dots, x_n)$: field obtained by adjoining x_i to F
- R^d : d th graded part of a graded (pseudo)ring R

Linear algebra:

- $V \oplus W$: direct sum
- $V^{\oplus n}$: direct sum of V , n times
- $V \otimes W$: tensor product
- $V^{\otimes n}$: tensor product of V , n times
- $V^\vee$: dual space
- $T^\vee$: dual map (for T a vector space)
- $T^\dagger$: conjugate transpose (for T a vector space)
- $\langle -, - \rangle$: a bilinear form
- $\text{Mat}(V)$: endomorphisms of V , i.e. $\text{Hom}_k(V, V)$
- $\mathbf{e}_1, \dots, \mathbf{e}_n$: the “standard basis” of $k^{\oplus n}$

§D.4 Quantum computation

- $|\psi\rangle$: a vector in some vector space H
- $\langle\psi|$: a vector in some vector space $H^\vee$, dual to $|\psi\rangle$.
- $\langle\phi|\psi\rangle$: evaluation of an element $\langle\phi| \in H^\vee$ at $|\psi\rangle \in H$.
- $|\uparrow\rangle, |\downarrow\rangle$: spin z -up, spin z -down
- $|\rightarrow\rangle, |\leftarrow\rangle$: spin x -up, spin x -down
- $|\otimes\rangle, |\odot\rangle$: spin y -up, spin y -down

§D.5 Topology and (complex) analysis

Common topological spaces:

- S^1 : the unit circle
- S^n : surface of an n -sphere (in $\mathbb{R}^{n+1}$)
- D^{n+1} : closed $n + 1$ dimensional ball (in $\mathbb{R}^{n+1}$)
- $\mathbb{RP}^n$: real projective n -space
- $\mathbb{CP}^n$: complex projective n -space

Some topological notation:

- ∂Y : boundary of a set Y (in some topological space)
- X/S : quotient topology of X by $S \subseteq X$
- $X \times Y$: product topology of spaces X and Y
- $X \amalg Y$: disjoint union of spaces X and Y

- $X \vee Y$: wedge product of (pointed) spaces X and Y

Complex analysis:

- $\int_{\alpha} f dz$: contour integral of f along path α
- $\text{Res}(f; p)$: the residue of a meromorphic function f at point p
- $\mathbf{I}(\gamma, p)$: winding number of γ around p .

Algebraic topology:

- $\alpha \simeq \beta$: for paths, this indicates path homotopy
- $*$: path concatenation
- $\pi_1(X) = \pi_1(X, x_0)$: the fundamental group of (pointed) space X
- $\pi_n(X) = \pi_n(X, x_0)$: the n th homotopy group of (pointed) space X
- $f_{\#}$: the induced map $\pi_1(X) \rightarrow \pi_1(Y)$ of $f : X \rightarrow Y$
- Δ^n : the standard n -simplex
- $\partial\sigma$: the boundary of a singular n -simplex σ
- $H_n(A_{\bullet})$: the n th homology group of the chain complex $A_{\bullet}$
- $H_n(X)$: the n th homology group of a space X
- $\tilde{H}_n(X)$: the n th reduced homology group of X
- $H_n(X, A)$: the n th relative homology group of X and $A \subseteq X$
- f_* : the induced map on $H_n(A_{\bullet}) \rightarrow H_n(B_{\bullet})$ of $f : A_{\bullet} \rightarrow B_{\bullet}$, or $H_n(X) \rightarrow H_n(Y)$ for $f : X \rightarrow Y$
- $\chi(X)$: Euler characteristic of a space X
- $H^n(A^{\bullet})$: the n th cohomology group of a cochain complex $A^{\bullet}$
- $H^n(A_{\bullet}; G)$: the n th cohomology group of the cochain complex obtained by applying $\text{Hom}(-, G)$ to $A_{\bullet}$
- $H^n(X; G)$: the n th cohomology group/ring of X with G -coefficients
- $\tilde{H}^n(X; G)$: the n th reduced cohomology group/ring of X with G -coefficients
- $H^n(X, A; G)$: the n th relative cohomology group/ring of X and $A \subset X$ with G -coefficients
- $f^{\sharp}$: the induced map on $H^n(A^{\bullet}) \rightarrow H^n(B^{\bullet})$ of $f : A^{\bullet} \rightarrow B^{\bullet}$, or $H^n(X) \rightarrow H^n(Y)$ for $f : X \rightarrow Y$
- $\text{Ext}(-, -)$: the Ext functor
- $\phi \smile \psi$: cup product of cochains ϕ and ψ

§D.6 Category theory

Some common categories (in alphabetical order):

- $\mathbf{Grp}$: category of groups
- $\mathbf{CRing}$: category of commutative rings
- $\mathbf{Top}$: category of topological spaces
- $\mathbf{Top}_*$: category of pointed topological spaces
- $\mathbf{Vect}_k$: category of k -vector spaces
- $\mathbf{FDVect}_k$: category of finite-dimensional vector spaces
- $\mathbf{Set}$: category of sets
- $\mathbf{hTop}$: category of topological spaces, whose morphisms are homotopy classes of maps
- $\mathbf{hTop}_*$: pointed version of $\mathbf{hTop}$
- $\mathbf{hPairTop}$: category of pairs (X, A) with morphisms being pair-homotopy equivalence classes
- $\mathbf{OpenSets}(X)$: the category of open sets of X , as a poset

Operations with categories:

- $\mathbf{obj} \mathcal{A}$: objects of the category $\mathcal{A}$
- $\mathcal{A}^{\text{op}}$: opposite category
- $\mathcal{A} \times \mathcal{B}$: product category
- $[\mathcal{A}, \mathcal{B}]$: category of functors from $\mathcal{A}$ to $\mathcal{B}$
- $\ker f : \text{Ker } f \rightarrow B$: for $f : A \rightarrow B$, categorical kernel
- $\text{coker } f : A \rightarrow \text{Coker } f$: for $f : A \rightarrow B$, categorical cokernel
- $\text{im } f : A \rightarrow \text{Im } f$: for $f : A \rightarrow B$, categorical image

§D.7 Differential geometry

- Df : total derivative of f
- $(Df)_p$: total derivative of f at point p
- $\frac{\partial f}{\partial e_i}$: i^{th} partial derivative
- α_p : evaluating a k -form α at p
- $\int_c \alpha$: integration of the differential form α over a cell c
- $d\alpha$: exterior derivative of a k -form α
- $\phi^* \alpha$: pullback of k -form α by ϕ

§D.8 Algebraic number theory

- $\overline{\mathbb{Q}}$: ring of algebraic numbers
- $\overline{\mathbb{Z}}$: ring of algebraic integers
- $\overline{F}$: algebraic closure of a field F
- $N_{K/\mathbb{Q}}(\alpha)$: the norm of α in extension $K/\mathbb{Q}$
- $\text{Tr}_{K/\mathbb{Q}}(\alpha)$: the trace of α in extension $K/\mathbb{Q}$
- $\mathcal{O}_K$: ring of integers in K
- $\mathfrak{a} + \mathfrak{b}$: sum of two ideals $\mathfrak{a}$ and $\mathfrak{b}$
- $\mathfrak{a}\mathfrak{b}$: ideal generated by products of elements in ideals $\mathfrak{a}$ and $\mathfrak{b}$
- $\mathfrak{a} \mid \mathfrak{b}$: ideal $\mathfrak{a}$ divides ideal $\mathfrak{b}$
- $\mathfrak{a}^{-1}$: the inverse of $\mathfrak{a}$ in the ideal group
- $N(I)$: ideal norm
- Cl_K : class group of K
- Δ_K : discriminant of number field K
- $\mu(\mathcal{O}_K)$: set of roots of unity contained in $\mathcal{O}_K$
- $[K : F]$: degree of a field extension
- $\text{Aut}(K/F)$: set of field automorphisms of K fixing F
- $\text{Gal}(K/F)$: Galois group of K/F
- $D_{\mathfrak{p}}$: decomposition group of prime ideal $\mathfrak{p}$
- $I_{\mathfrak{p}}$: inertia group of prime ideal $\mathfrak{p}$
- $\text{Frob}_{\mathfrak{p}}$: Frobenius element of $\mathfrak{p}$ (element of $\text{Gal}(K/\mathbb{Q})$)
- $P_K(\mathfrak{m})$: ray of principal ideals of a modulus $\mathfrak{m}$
- $I_K(\mathfrak{m})$: fractional ideals of a modulus $\mathfrak{m}$
- $C_K(\mathfrak{m})$: ray class group of a modulus $\mathfrak{m}$
- $\left(\frac{L/K}{\bullet}\right)$: the Artin symbol
- $\text{Ram}(L/K)$: primes of K ramifying in L
- $\mathfrak{f}(L/K)$: the conductor of L/K

§D.9 Representation theory

- $k[G]$: group algebra
- $V \oplus W$: direct sum of representations $V = (V, \rho_V)$ and $W = (W, \rho_W)$ of an algebra A
- $V^\vee$: dual representation of a representation $V = (V, \rho_V)$
- $\text{Reg}(A)$: regular representation of an algebra A
- $\text{Hom}_{\text{rep}}(V, W)$: algebra of morphisms $V \rightarrow W$ of representations
- χ_V : the character $A \rightarrow k$ attached to an A -representation V
- $\text{Classes}(G)$: set of conjugacy classes of G
- $\text{Fun}_{\text{class}}(G)$: the complex vector space of functions $\text{Classes}(G) \rightarrow \mathbb{C}$
- $V \otimes W$: tensor product of representations $V = (V, \rho_V)$ and $W = (W, \rho_W)$ of a group G (rather than an algebra)
- $\mathbb{C}_{\text{triv}}$: the trivial representation
- $\mathbb{C}_{\text{sign}}$: the sign representation

§D.10 Algebraic geometry

- $\mathcal{V}(-)$: vanishing locus of a set or ideal
- $\mathbb{A}^n$: n -dimensional (complex) affine space
- $\sqrt{I}$: radical of an ideal I
- $\mathbb{C}[V]$: coordinate ring of an affine variety V
- $\mathcal{O}_V(U)$: ring of rational functions on U
- $D(f)$: distinguished open set
- $\mathbb{CP}^n$: complex projective n -space (ambient space for projective varieties)
- $(x_0 : \cdots : x_n)$: coordinates of projective space
- U_i : standard affine charts
- $\mathcal{V}_{\text{pr}}(-)$: projective vanishing locus.
- h_I, h_V : Hilbert function of an ideal I or projective variety V
- f^* : the pullback $\mathcal{O}_Y \rightarrow \mathcal{O}_X(f^{\text{pre}}(U))$ obtained from $f : X \rightarrow Y$
- $\mathcal{F}_p$: the stalk of a (pre-)sheaf $\mathcal{F}$ at a point p
- $[s]_p$: the germ of $s \in \mathcal{F}(U)$ at the point p
- $\mathcal{O}_{X,p}$: shorthand for $(\mathcal{O}_X)_p$.
- $\mathcal{F}^{\text{sh}}$: sheafification of pre-sheaf $\mathcal{F}$

- $\alpha_p : \mathcal{F}_p \rightarrow \mathcal{G}_p$: morphism of stalks obtained from $\alpha : \mathcal{F} \rightarrow \mathcal{G}$
- $\mathfrak{m}_{X,p}$: the maximal ideal of $\mathcal{O}_{X,p}$
- $\text{Spec } R$: the spectrum of a ring R
- $S^{-1}R$: localization of ring R at a set S
- $\text{Proj } R$: the projective scheme of a graded ring S

§D.11 Set theory

- ZFC: standard theory of ZFC
- ZFC^+ : standard theory of ZFC, plus the sentence “there exists a strongly inaccessible cardinal”
- $\mathcal{P}(S)$: power set of S
- $A \wedge B$: A and B
- $A \vee B$: A or B
- $\neg A$: not A
- V : class of all sets (von Neumann universe)
- ω : the first infinite ordinal, also the set of nonnegative integers
- V_α : level of the von Neumann universe
- On : class of ordinals
- $\bigcup A$: the union of elements inside A
- $A \approx B$: sets A and B are equinumerous
- $\aleph_\alpha$: the aleph numbers
- $\text{cof } \lambda$: the cofinality of λ
- $\mathcal{M} \models \phi[b_1, \dots, b_n]$: model $\mathcal{M}$ satisfies sentence ϕ with parameters $b_1, \dots, b_n$
- $\Delta_n, \Sigma_n, \Pi_n$: levels of the Levy hierarchy
- $\mathcal{M}_1 \subseteq \mathcal{M}_2$: $\mathcal{M}_1$ is a substructure of $\mathcal{M}_2$
- $\mathcal{M}_1 \prec \mathcal{M}_2$: $\mathcal{M}_1$ is an elementary substructure of $\mathcal{M}_2$
- $p \parallel q$: elements p and q of a poset $\mathbb{P}$ are compatible
- $p \perp q$: elements p and q of a poset $\mathbb{P}$ are incompatible
- Name_α : the hierarchy of $\mathbb{P}$ -names
- τ^G : interpretation of a name τ by filter G
- $M[G]$: the model obtained from a forcing poset $G \subseteq \mathbb{P}$
- $p \Vdash \varphi(\sigma_1, \dots, \sigma_n)$: $p \in \mathbb{P}$ forces the sentence φ

- $\check{x}$: the name giving an $x \in M$ when interpreted
- $\dot{G}$: the name giving G when interpreted

DRAFT (Evan Chen)
Updated February 15, 2017

E Terminology on sets and functions

This appendix will cover some notions on sets and functions such as “bijections”, “equivalence classes”, and so on.

Remark for experts: I am not dealing with foundational issues in this chapter. See Chapter 58 (and onwards) if that’s what you’re interested in. Consequently I will not prove most assertions.

§E.1 Sets

A **set** for us will just be a collection of elements (whatever they may be). For example, the set $\mathbb{N} = \{1, 2, 3, 4, \dots\}$ is the positive integers, and $\mathbb{Z} = \{\dots, -2, -1, 0, 1, 2, \dots\}$ is the set of all integers. As another example, we have a set of humans:

$$H = \{x \mid x \text{ is a featherless biped}\}.$$

(Here the “ $\mid$ ” means “such that”.)

There’s also a set with no elements, which we call the **empty set**. It’s denoted by $\emptyset$.

It’s conventional to use capital letters for sets (like H), and lowercase letters for elements of sets (like x).

Definition E.1.1. We write $x \in S$ to mean “ x is in S ”, for example $3 \in \mathbb{N}$.

Definition E.1.2. If every element of a set A is also in a set B , then we say A is a **subset** of B , and denote this by $A \subseteq B$. If moreover $A \neq B$, we say A is a **proper subset** and write $A \subsetneq B$. (This is analogous to $\leq$ and $<$.)

Example E.1.3 (Examples of subsets)

- (a) $\{1, 2, 3\} \subseteq \mathbb{N} \subseteq \mathbb{Z}$.
- (b) $\emptyset \subseteq A$ for any set A . (Why?)
- (c) $A \subseteq A$ for any set A .

Definition E.1.4. We write

- $A \cup B$ for the set of elements in *either* A or B (possibly both), called the **union** of A and B .
- $A \cap B$ for the set of elements in *both* A and B , and called the **intersection** of A and B .
- $A \setminus B$ for the set of elements in A but *not* in B .

Example E.1.5 (Examples of set operations)

Let $A = \{1, 2, 3\}$ and $B = \{3, 4, 5\}$. Then

$$A \cup B = \{1, 2, 3, 4, 5\}$$

$$A \cap B = \{3\}$$

$$A \setminus B = \{1, 2\}.$$

Exercise E.1.6. Convince yourself: for any sets A and B , we have $A \cap B \subseteq A \subseteq A \cup B$.

Here are some commonly recurring sets:

- $\mathbb{C}$ is the set of complex numbers, like $3.2 + \sqrt{2}i$.
- $\mathbb{R}$ is the set of real numbers, like $\sqrt{2}$ or π .
- $\mathbb{N}$ is the set of positive integers, like 5 or 9.
- $\mathbb{Q}$ is the set of rational numbers, like $7/3$.
- $\mathbb{Z}$ is the set of integers, like -2 or 8 .

(These are pronounced in the way you would expect: “see”, “are”, “en”, “cue”, “zed”.)

§E.2 Functions

Given two sets A and B , a **function** f from A to B is a mapping of every element of A to some element of B .

We call A the **domain** of f , and B the **codomain**. We write this as $f : A \rightarrow B$ or $A \xrightarrow{f} B$.

Abuse of Notation E.2.1. If the name f is not important, we will often just write $A \rightarrow B$.

We write $f(a) = b$ or $a \mapsto b$ to signal that f takes a to b .

If B has 0 as an element and $f(a) = 0$, we often say a is a **root** or **zero** of f , and that f **vanishes** at a .

E.2.1 Injective / surjective / bijective functions

Definition E.2.2. A function f is:

- **injective** if whenever $a_1 \neq a_2$, we have $f(a_1) \neq f(a_2)$. Often, we will write $f : A \hookrightarrow B$ to emphasize this.
- **surjective** if whenever $b \in B$, there exists some $a \in A$ such that $f(a) = b$. Often, we will write $f : A \twoheadrightarrow B$ to emphasize this.
- **bijection** if both, i.e. f is a “one-to-one correspondence”: for every $b \in B$ there is a *unique* $a \in A$ such that $f(a) = b$.

Example E.2.3 (Examples of functions)

- (a) There's a function taking every human to their age in years (rounded to the nearest integer). This function is **not injective**, because for example there are many people with age 20. This function is also **not surjective**: no one has age 10000.
- (b) There's a function taking every USA citizen to their social security number. This is also **not surjective** (no one has SSN equal to 3), but at least it **is injective** (no two people have the same SSN).

Example E.2.4 (Examples of bijections)

- (a) Let $A = \{1, 2, 3, 4, 5\}$ and $B = \{6, 7, 8, 9, 10\}$. Then the function $f : A \rightarrow B$ by $a \mapsto a + 5$ is a bijection.
- (b) In a classroom with 30 seats, there is exactly one student in every seat. Thus the function taking each student to the seat they're in is a bijection; in particular, there are exactly 30 students.

Remark E.2.5. Assume for convenience that A and B are finite sets. Then:

- If $f : A \hookrightarrow B$ is injective, then the size of A is at most the size of B .
- If $f : A \twoheadrightarrow B$ is surjective, then the size of A is at least the size of B .
- If $f : A \rightarrow B$ is a bijection, then the size of A equals the size of B .

Now, notice that if $f : A \rightarrow B$ is a bijection, then we can “apply f backwards”: (for example, rather than mapping each student to the seat they're in, we map each seat to the student sitting in it). This is called an **inverse function**; we denote it $f^{-1} : B \rightarrow A$.

E.2.2 Images and pre-images

Let $X \xrightarrow{f} Y$ be a function.

Definition E.2.6. Suppose $T \subseteq Y$. The **pre-image** $f^{\text{pre}}(T)$ is the set of all $x \in X$ such that $f(x) \in T$. Thus, $f^{\text{pre}}(T)$ is a subset of X .

Example E.2.7 (Examples of pre-image)

Let $f : H \rightarrow \mathbb{Z}$ be the age function from earlier. Then

- (a) $f^{\text{pre}}(\{13, 14, 15, 16, 17, 18, 19\})$ is the set of teenagers.
- (b) $f^{\text{pre}}(\{0\})$ is the set of newborns.
- (c) $f^{\text{pre}}(\{1000, 1001, 1002, \dots\}) = \emptyset$, as I don't think anyone is that old.

Abuse of Notation E.2.8. By abuse of notation, we may abbreviate $f^{\text{pre}}(\{y\})$ to $f^{\text{pre}}(y)$. So for example, $f^{\text{pre}}(\{0\})$ above becomes shortened to $f^{\text{pre}}(0)$.

The dual notion is:

Definition E.2.9. Suppose $S \subseteq X$. The **image** $f“(S)$ is the set of all things of the form $f(s)$.

Example E.2.10 (Examples of images)

Let $A = \{1, 2, 3, 4, 5\}$ and $B = \mathbb{Z}$. Consider a function $f : A \rightarrow B$ given by

$$f(1) = 17 \quad f(2) = 17 \quad f(3) = 19 \quad f(4) = 30 \quad f(5) = 234.$$

- (a) The image $f“(\{1, 2, 3\})$ is the set $\{17, 19\}$.
- (b) The image $f“(A)$ is the set $\{17, 19, 30, 234\}$.

Question E.2.11. Suppose $f : A \rightarrow B$ is surjective. What is $f“(A)$?

§E.3 Equivalence relations

Let X be a fixed set now. A binary relation $\sim$ on X assigns a truth value “true” or “false” to $x \sim y$ for each x or y . Now an **equivalence relation** $\sim$ on X is a binary relation which satisfies the following axioms:

- Reflexive: we have $x \sim x$.
- Symmetric: if $x \sim y$ then $y \sim x$.
- Transitive: if $x \sim y$ and $y \sim z$ then $x \sim z$.

An **equivalence class** is then a set of all things equivalent to each other. One can show that X becomes partitioned by these equivalence classes:

Example E.3.1 (Example of an equivalence relation)

Let $\mathbb{N}$ denote the set of positive integers. Then suppose we declare $a \sim b$ if a and b have the same last digit, for example $131 \sim 211$, $45 \sim 125$, and so on.

Then $\sim$ is an equivalence relation. It partitions $\mathbb{N}$ into ten equivalence classes, one for each trailing digit.

Often, the set of equivalence classes will be denoted $X/\sim$ (pronounced “ X mod sim”).

Image Attributions

- [ca] Pop-up casket. *Omega exp*. Public domain. URL: <https://commons.wikimedia.org/wiki/File:Omega-exp-omega-labeled.svg>.
- [Ee] Eeyore22. *Weierstrass function*. Public domain. URL: <https://commons.wikimedia.org/wiki/File:WeierstrassFunction.svg>.
- [Fr] Fropuff. *Klein bottle*. Public domain. URL: <https://en.wikipedia.org/wiki/File:KleinBottle-01.png>.
- [Ge] Topological Girl's Generation. *Topological Girl's Generation*. URL: <http://topologicalgirlsgeneration.tumblr.com/>.
- [gk] g.kov. *Normal surface vector*. URL: <http://tex.stackexchange.com/a/235142/76888>.
- [Go08] Abstruse Goose. *Math Text*. CC 3.0. 2008. URL: <http://abstrusegoose.com/12>.
- [Go09] Abstruse Goose. *Zornaholic*. CC 3.0. 2009. URL: <http://abstrusegoose.com/133>.
- [Ho] George Hodan. *Apple*. Public domain. URL: <http://www.publicdomainpictures.net/view-image.php?image=20117>.
- [In] Inductiveload. *Klein Bottle Folding*. Public domain. URL: https://commons.wikimedia.org/wiki/File:Klein_Bottle_Folding_1.svg.
- [Kr] Petr Kratochvil. *Velociraptor*. Public domain. URL: <http://www.publicdomainpictures.net/view-image.php?image=93881>.
- [Mu] Randall Munroe. *Fundamental Forces*. CC 2.5. URL: <https://xkcd.com/1489/>.
- [Na] Krish Navedala. *Stokes patch*. Public domain. URL: https://en.wikipedia.org/wiki/File:Stokes_patch.svg.
- [Or] Ben Orlin. *The Math Major Who Never Reads Math*. URL: <http://mathwithbaddrawings.com/2015/03/17/the-math-major-who-never-reads-math/>.
- [To] Toahigherlevel. *Projection color torus*. Public domain. URL: https://en.wikipedia.org/wiki/File:Projection_color_torus.jpg.
- [Wo] Wordslaugh. *Covering space diagram*. CC 3.0. URL: https://commons.wikimedia.org/wiki/File:Covering_space_diagram.svg.

Bibliography

- [Ax97] Sheldon Axler. *Linear algebra done right*. New York: Springer, 1997. ISBN: 978-0387982588.
- [Ba10] Joseph Bak. *Complex analysis*. New York: Springer Science+Business Media, LLC, 2010. ISBN: 978-1441972873.
- [Et11] Pavel Etingof. “Introduction to Representation Theory”. 2011. URL: <http://math.mit.edu/~etingof/relect.pdf>.
- [Ga03] Andreas Gathmann. “Algebraic Geometry”. 2003. URL: <http://www.mathematik.uni-kl.de/~gathmann/class/alggeom-2002/main.pdf>.
- [Ga14] Dennis Gaitsgory. “Math 55a: Honors Abstract and Linear Algebra”. 2014. URL: <http://www.mit.edu/~evanchen/coursework.html>.
- [Ga15] Dennis Gaitsgory. “Math 55b: Honors Real and Complex Analysis”. 2015. URL: <http://www.mit.edu/~evanchen/coursework.html>.
- [Go11] Timothy Gowers. “Normal subgroups and quotient groups”. 2011. URL: <https://gowers.wordpress.com/2011/11/20/normal-subgroups-and-quotient-groups/>.
- [Ha02] Allen Hatcher. *Algebraic topology*. Cambridge, New York: Cambridge University Press, 2002. ISBN: 0-521-79160-X. URL: <http://opac.inria.fr/record=b1122188>.
- [Hi13] A. J. Hildebrand. “Introduction to Analytic Number Theory”. 2013. URL: <http://www.math.illinois.edu/~hildebr/ant/>.
- [Ko14] Peter Koellner. “Math 145a: Set Theory I”. 2014. URL: <http://www.mit.edu/~evanchen/coursework.html>.
- [Le] Holden Lee. “Number Theory”. URL: <https://github.com/holdenlee/number-theory>.
- [Le02] Hendrik Lestra. “The Chebotarev Density Theorem”. 2002. URL: <http://websites.math.leidenuniv.nl/algebra/>.
- [Le14] Tom Leinster. *Basic category theory*. Cambridge: Cambridge University Press, 2014. ISBN: 978-1107044241. URL: <https://arxiv.org/abs/1612.09375>.
- [Li15] Seth Lloyd. “18.435J: Quantum Computation”. 2015. URL: <http://www.mit.edu/~evanchen/coursework.html>.
- [Ma13a] Laurentiu Maxim. “Math 752 Topology Lecture Notes”. 2013. URL: <https://www.math.wisc.edu/~maxim/752notes.pdf>.
- [Ma13b] Maxima. “Burnside’s Lemma, post 6”. 2013. URL: <http://www.aops.com/Forum/viewtopic.php?p=3089768#p3089768>.
- [Mi14] Alexandre Miquel. “An Axiomatic Presentation of the Method of Forcing”. 2014. URL: <http://www.fing.edu.uy/~amiquel/forcing/>.
- [Mu00] James Munkres. *Topology*. 2nd. Prentice-Hall, Inc., Jan. 2000. ISBN: 9788120320468. URL: <http://amazon.com/o/ASIN/8120320468/>.
- [Og10] Frederique Oggier. “Algebraic Number Theory”. 2010. URL: <http://www1.spms.ntu.edu.sg/~frederique/ANT10.pdf>.

- [Pu02] C. C. Pugh. *Real mathematical analysis*. New York: Springer, 2002. ISBN: 978-0387952970.
- [Sj05] Reyer Sjamaar. “Manifolds and Differential Forms”. 2005. URL: <http://www.math.cornell.edu/~sjamaar/classes/3210/notes.html>.
- [U108] Brooke Ullery. “Minkowski Theory and the Class Number”. 2008. URL: <http://www.math.uchicago.edu/~may/VIGRE/VIGRE2008/REUPapers/Ullery.pdf>.
- [Va15] Ravi Vakil. “The Rising Sea: Foundations of Algebraic Geometry”. 2015. URL: <http://math.stanford.edu/~vakil/216blog/>.
- [Ya12] Andrew Yang. “Math 43: Complex Analysis”. 2012. URL: <https://math.dartmouth.edu/~m43s12/syllabus.html>.

DRAFT (Evan Chen)
Updated February 15, 2017