Cement2: Temporal Hardware Transactions for High-Level and Efficient FPGA Programming

Youwei Xiao
Peking University
China

Zizhang Luo
Peking University
China

Weijie Peng
Peking University
China

Yuyang Zou
Peking University
China

Yun Liang
Peking University
China

Abstract

Hardware design faces a fundamental challenge: raising abstraction to improve productivity while maintaining control over low-level details like cycle accuracy. Traditional RTL design in languages like SystemVerilog composes modules through wiring-style connections that provide weak guarantees for behavioral correctness. While high-level synthesis (HLS) and emerging abstractions attempt to address this, they either introduce unpredictable overhead or restrict design generality. Although transactional HDLs provide a promising foundation by lifting design abstraction to atomic and composable rules, they solely model intra-cycle behavior and do not reflect the native temporal design characteristics, hindering applicability and productivity for FPGA programming scenarios.

We propose temporal hardware transactions, a new abstraction that brings cycle-level timing awareness to designers at the transactional language level. Our approach models temporal relationships between rules and supports the description of rules whose actions span multiple clock cycles, providing intuitive abstraction to describe multi-cycle architectural behavior. We implement this in `Cement2`, a transactional HDL embedded in Rust, enabling programming hardware constructors to build both intra-cycle and temporal transactions. `Cement2`'s synthesis framework lowers description abstraction through multiple analysis and optimization phases, generating efficient hardware. With `Cement2`'s abstraction, we program a RISC-V soft-core processor, custom CPU instructions, linear algebra kernels, and systolic array accelerators, leveraging the high-level abstraction for boosted productivity. Evaluation shows that `Cement2` does not sacrifice performance and resources compared to hand-coded RTL designs, demonstrating the high applicability for general FPGA design tasks.

1 Introduction

Hardware design at the register transfer level (RTL) is becoming increasingly challenging as architectural innovations demand more complex implementations [6, 46, 49]. Traditional RTL design in languages like SystemVerilog and VHDL composes larger modules by connecting ports of smaller ones, but this wiring-style composition provides weak guarantees. Behavioral correctness cannot be guaranteed, causing data communication errors like producer-consumer mismatches [37]. For FPGA programming, this fundamental limitation causes a significant gap between architecture design and hardware implementation. Designers often need to describe their architecture design in two models: higher-level simulators [35, 45]

for fast idea validation, and tedious RTL design for detailed implementation on FPGAs. The potentially inaccurate simulation results and missing area and power information can lead to wrong architectural decisions and impede iteration speed. Besides, designers must invest significant cognitive load and development cost to ensure correct implementation aligned with design intent.

Raising hardware design abstraction is necessary to address these challenges. Describing hardware at a higher level can avoid the tedious and error-prone composition. However, choosing the appropriate abstraction level presents a non-trivial tradeoff. High-level synthesis (HLS) [5, 21, 30] takes untimed software descriptions and generates hardware designs, boosting productivity. However, it often produces unpredictable performance and resource overheads [39], since it loses control over low-level details, such as cycle accuracy and resource overheads, limiting its applicability. Accelerator design languages [32, 38, 39, 41, 57] are also too specific. Other emerging FPGA programming approaches, including Cement [52] and others [19, 40, 47, 58], add latency-sensitive/-insensitive information but put more constraints on design architecture and hardware composition manners.

An ideal FPGA programming methodology should raise the design abstraction level while maintaining control over clock-cycle timing and low-level hardware details. Transactional HDLs [11, 14, 42] provide a promising foundation by abstracting hardware design as behavioral rules, composable logic units with execution atomicity. However, prior rule-based works are restricted to intra-cycle logic and cannot reflect hardware's temporal behavior or architecture design intent. In practice, designers either adopt a latency-insensitive design style to eliminate temporal concerns with extra hardware overheads or manually coordinate intra-cycle rules for efficient latency-sensitive design at the expense of productivity.

We propose a new abstraction, *temporal hardware transactions*, which brings cycle-level timing awareness to designers at the language level, for high-level FPGA programming. We model temporal relationships between rules, building an intuitive temporal view of rules' execution across cycles. To further boost productivity in describing multi-cycle behaviors, we introduce multi-cycle rules whose actions span multiple clock cycles under specified constraints. Our abstraction's temporal behavior modeling enables comprehensive inter-cycle hardware analysis, checking, and optimizations, leading to compiler-enforced correctness and efficiency. Moreover, it employs a multi-phase synthesis flow to transition from high-level abstraction to low-level, while avoiding the introduction of unnecessary overhead, resulting in efficient hardware implementation. Both (micro)architecture and hardware design phases of

FPGA programming benefit from temporal hardware transactions. Architects need a precise estimation of performance and hardware overheads to guide (micro-) architectural optimization. With our abstractions, architects can accurately model and implement hardware using a single behavioral interface, rather than multiple models across behavioral and structural levels [33, 35, 45]. This greatly improves the development productivity. For hardware designers, our abstraction provides a more productive design methodology. It allows behavior description with temporal intuition and enables the compiler to conduct rich behavioral checking for early error detection. The powerful synthesis flow not only handles hybrid latency-sensitive/-insensitive but also performs retiming on multi-cycle rules, generating an efficient RTL implementation. Specifically, for HLS designers who care about performance, our abstraction provides more accurate control over details while retaining synthesis features; for RTL designers, the abstraction helps to boost productivity. Moreover, the compiler can detect many design errors early, greatly reducing the cost and effort required for fixing them later (e.g., manually writing test benches, debugging, etc.)

We implement the novel abstraction in the `Cement2` framework, which is abbreviated as `CMT2`. The frontend language `CMT2-rs` is a transactional HDL embedded in Rust that enables creating hardware constructors to describe hardware transactions with temporal behavior modeling. `Cement2` represents constructed rules in `CTIR` (Cement2 Transaction Intermediate Representation), providing a unified representation for both intra-cycle and temporal hardware transactions. The synthesis framework conducts temporal scheduling, temporal partitioning, and temporal implementation, generating a high-quality RTL implementation for FPGA deployment.

Our contributions are:

- The novel temporal hardware transactions that raise abstraction while maintaining control over FPGA programming details;
- The open-source `Cement2` framework including a Rust frontend and `CTIR` implementing temporal hardware transactions;
- A synthesis flow that efficiently generates high-quality RTL implementation for FPGA deployment.

We evaluate `Cement2` through four case studies on FPGA: building a RISC-V soft core, extending the core with custom instructions, describing linear algebra kernels, and designing systolic array accelerators. Experiments show that `Cement2`'s soft core achieves a higher frequency at 377MHz with lower resource usage compared to the Sodor [3] core designed in Chisel [7]. For other case studies, `Cement2` boosts design productivity and achieves comparable or better performance and hardware quality than human-crafted RTL designs. The evaluation demonstrates that temporal hardware transactions facilitate various design scenarios, and `Cement2` is a productive solution for general FPGA programming tasks.

2 Background and Motivation

This section discusses abstraction levels available for FPGA programming and analyzes their tradeoffs, as summarized in Table 1.

2.1 Hardware Abstractions and Tradeoffs

Traditional hardware design at the register-transfer level (RTL) [7, 18] provides a structural description of hardware components. In
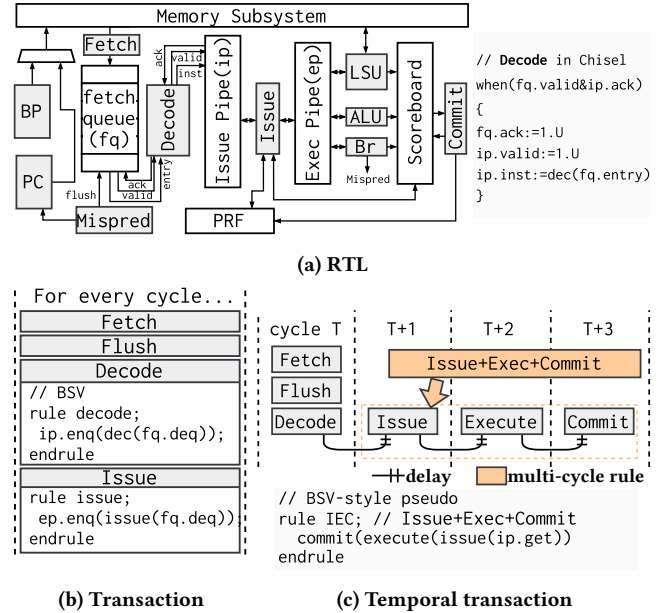

(a) RTL

(b) Transaction   **(c) Temporal transaction**

Figure 1: Motivating example: illustrating the 5-stage CPU core pipeline described in different abstraction levels.

RTL, data flows between registers, and logical operations are described through instantiation and structural wires. For example, in a processor pipeline as shown in Figure 1a, pipeline stages are divided by register-based stateful components like `fetch queue` and `issue pipeline register`, whose ports are connected by combinational logic blocks like `Decode` and `Issue`. The code block on the right presents Chisel [7] description of `Decode`, that structurally connects all wires from `fetch queue` and `issue pipeline register`. While this low-level description gives designers precise control over details, it has significant drawbacks: *structural composition requires designers to manually operate wires without behavior promises provided by the language abstraction*, making large designs verbose, error-prone, and hard to debug during RTL simulation [50] or FPGA running [62].

Although veteran RTL designers can avoid certain language pitfalls with the help of lint tools [48], they still can make incorrect designs. For instance, `fetch queue` is structurally connected to logic blocks including `Fetch`, `Decode`, and `Mispred`, with latency-insensitive protocol signals (`valid`s and `ack`s) exposed. One potential design mistake is that `fetch queue`'s `ack` signal for `Fetch` is not disabled by the high `flush` signal. When a misprediction is detected, the `Mispred` block will raise the `flush` signal to discard existing instructions in `fetch queue`, while the `Fetch` block may still fetch a new instruction from the wrong branch and enqueue it to `fetch queue` at the same cycle, causing flush failure and execution of wrong instructions. However, such errors cannot be checked and located by RTL tools. The reason is that *RTL description does not model hardware in the behavioral manner and is unaware of conflicts among behaviors that manipulate shared states.* From a behavioral view, both `Mispred`'s *flush* action and `Fetch`'s *enque* action try to update the `fetch queue`, and the *flush* action should prevent the

Table 1: Comparison of hardware and architecture design approaches. The table presents: (1) composition description manner, (2) inter-cycle behavior description support, (3) low-level control over details ("Cycle" for clock-cycle timing, and "Register" for register instantiation and access), (4) hardware overheads (performance and resources), and (5) design generality.

Approaches	Composition	Inter-cycle behavior	Control over...		Hardware overhead...		Generality
			Cycle	Register	Performance	Resources	
RTL [7, 18]	Wiring	✗	✓	✓	Low	Low	✓
Calyx [41]	Wiring	✓	✓	✓	Low	Low	✓✗
Filament [40]	Wiring	✓	✓	✓	Low	Low	✗
Cement [52]	Wiring	✓	✓	✓	Low	Low	✓✗
HLS (static [5], dynamic [30])	Call	✓	✗	✗	High	High	✗
PDL [58]	Call	✓	✓	✗	High	High	✗
Transaction [11, 42]	Call	✗	✓	✓	Low	Low	✓
Simulation models [35, 45]	Call	✓	✓✗	✗	Unavailable	Unavailable	✓
Cement2 (this work)	Call	✓	✓	✓	Low	Low	✓

enque action at the same cycle. Similar issues are common among CPU pipeline components, such as Scoreboard [59].

To overcome RTL's drawbacks, various approaches have been proposed to raise the abstraction level of hardware design. As shown in Table 1, high-level synthesis (HLS) [5, 30] takes untimed software descriptions where modules are composed through function calls. While this provides the highest abstraction level and productivity, it often produces unpredictable performance since it highly relies on heuristics and lacks detailed timing control. For example, HLS cannot synthesize a CPU pipeline with data forwarding and branch prediction since those features cause unresolved inter-iteration dependencies in the source software loop. Other approaches target specific hardware design patterns. PDL [58] generates pipelines from behavioral descriptions with CPU-specific features, but it cannot describe general out-of-order execution. Calyx [41] and Filament [40] generate efficient multi-cycle or pipelined designs from structural descriptions enhanced with control flow language or timeline type system, but they only target specific accelerator designs and cannot model the pipelined processor's complex features. Cement [52] also combines RTL description and control flow language, providing control over cycle accuracy for deterministic FPGA programming. Notably, both Calyx and Cement can be used without their high-level features, falling back to RTL design. Besides, they still suffer from RTL's error-prone features and cannot detect conflicts among behaviors that manipulate shared states. Overall, none of the approaches support general design tasks due to their over-specialized abstractions.

2.2 Transactional Hardware Design

Transactional hardware design approaches [11, 14, 42] organize hardware as collections of atomic transactions. We provide a formal definition of a transactional module:

Definition 2.1 (Transactional Hardware Module). A transactional hardware module is denoted as $M = \langle I, R, S \rangle$, where I is a set of instances, R is a set of rules, and S is a set of binary relations for scheduling priority. Each rule $\langle id, g, f = a \mid \lambda x.a \rangle \in R$ is defined as a guarded atomic action. Here, g is the rule guard, a *side-effect-free* boolean predicate for explicit rule fire conditions, and f is the rule fire logic. There are two types of rules: *always* ($f = a$) executes proactively, and *method* ($f = \lambda x.a$) executes when called with arguments x and returns values. Action a is a set of expressions and

method calls, including register read and write. Each precedence relation $(\prec) \langle r_i, r_j \rangle \in S$ specifies that the rule r_i must execute earlier than rule r_j in every clock cycle.

Figure 1b illustrates the CPU pipeline described as hardware transactions and presents the BSV [42] description, for example. Specifically, Decode and Issue logic are described as two *always* rules, both of which operate stateful components, such as dequeueing fetch queue, and enqueueing issue pipeline register. In this way, hardware transactions provide a *behavioral* description.

Transactional HDLs' execution model guarantees *guarded atomicity* of rules: *one rule can execute only when its guard holds and all method calls in the fire logic can execute, and they are executed atomically.* Besides, it resolves *conflicts* between rule behaviors that manipulate shared states. Specifically, for module M, it generates a *scheduler* according to the precedence relation set S with a total order to execute the rules in every cycle. For example, we assume that fetch queue's rules have precedence relations: deq≺enq≺flush, then for a reasonable CPU pipeline schedule order [Flush, Decode, Fetch], the execution of Flush, which calls flush, at an earlier place will block both Decode and Fetch due to the *precedence violations* of deq≺flush and enq≺flush. This mechanism helps hardware designers avoid the design mistakes discussed in Section 2.1.

However, current transactional HDLs are limited by their intra-cycle semantics. They lack a temporal view, as all rules stand side by side and try to fire within every clock cycle. This causes the following drawbacks: (1) tedious human efforts for coordinating multi-cycle behaviors, such as describing individual rules for every logic block in Figure 1b and composing them into a complete pipeline by describing rules' manipulation on pre-instantiated stateful components, (2) inability of analysis, checking and optimization for inter-cycle hardware behaviors, such as latency sensitivity, (3) limited synthesis capabilities, such as forcing designers to manually split a rule on critical path into multiple rules to meet timing-closure constraints through a trial-and-error process, and (4) anti-intuitive design experience, showing no pipeline stages with temporal relationships in Figure 1b, misaligned with human intuition.

2.3 Motivation

The fundamental challenge in advanced FPGA programming is striking the right balance between abstraction and control. We also analyze features of architecture simulators [35, 45], identifying key

insights to overcome the drawbacks of transactional HDLs. Table 1 shows that detailed simulation models, such as gem5 O3CPU [20], retain control over cycle and support general design. The key to this success is that they provide flexible, pure-software manipulation of temporal behavior. For example, the gem5 O3CPU model provides a `IEW` stage to handle the behavior of dispatching, issuing, executing, and writing back instructions. However, such temporal behavior across multiple clock cycles cannot be described in existing transactional HDLs. This motivates us to provide explicit temporal modeling support to hardware transactions. Accordingly, we propose *temporal hardware transactions*, which intuitively describe pipeline stages and their temporal relationships and support joint multi-cycle behavioral description of multiple stages, as shown in Figure 1c. The pseudo-description of multiple stages as one rule `IEC` shows that our abstraction eliminates the tedious human efforts of describing multiple individual rules that coordinate through manipulating shared instances, as did in Figure 1b.

Moreover, implementation overheads must be carefully considered for hardware design, which are beyond the concern of simulation models. We need synthesis techniques to translate between neighboring abstraction levels while preserving implementation efficiency. We implement the insights above in the `Cement2` framework, which addresses the fundamental tradeoff between design productivity and hardware quality. Table 1 shows that our approach solves the intra-cycle limitations of prior transactional HDLs and uses fewer hardware overheads on the CPU pipeline example, demonstrating *temporal hardware transactions* as a promising high-level abstraction for hardware and architecture design.

3 Frontend and Core Abstraction

In this section, we introduce the frontend language of our methodology, `CMT2-rs`, which uniformly supports both intra-cycle transactions and *temporal hardware transactions*.

3.1 `CMT2-rs`

`CMT2-rs` is a modern transactional HDL embedded in Rust. It implements Definition 2.1 in an *interface-constructor* pattern, whose basic syntax is shown in Figure 2a. An *interface*, declared by an `itfc_decl!` block, declares the *method* rules to be called. Specifically, the `param T` defines a data type parameter, and the `struct` definition specifies the ports with directions and data types. Figure 2a defines an interface named `Itfc`, which has two ports: input `a` and output `b`, both of which have the parametric data type `T`. It declares a method named `met`, which takes an argument through port `a` and returns a result through port `b`. A *constructor*, defined by a `#[module]` function, fills instances and rule implementations to construct a module of the given interface. Figure 2a defines a constructor named `constr`. It builds a module of the interface `Itfc` given a specific data type `t`. The `io!` statement specifies the data type parameter `T` as `t`, and the returned variable `io` can be used to access ports (e.g., `io.a`). The `instance!` statement instantiates a register by calling the `reg` constructor. The `method!` and `always!` statements define *method* and *always* rules, respectively, where the *guard* expression is enclosed by brackets and the fire *actions* are enclosed by braces. For example, the `constr` constructor defines two rules, the *method* `met` and the *always* `r0`, and the `prec!` statement

specifies the *precedence* relation between them, met≺r0. `CMT2-rs` combines the benefits of traditional hardware transactions with flexible and parameterized construction support of Rust embedding.

3.2 Temporal Hardware Transactions

We introduce the core features of *temporal hardware transactions*: *temporal relationships* and *multi-cycle rules*. They provide a temporal view of hardware behavior and enable inter-cycle analysis. Both features are implemented in `CMT2-rs`.

3.2.1 Temporal Relationships. We define *temporal relationships* among rules as two-fold: (1) *temporal guard* specifies the fire condition of a rule based on the execution history of predecessor rules; and (2) *temporal message passing* delivers data between temporally-related rules through *channels* to live across clock cycles.

Syntax. Figure 2b shows the syntax of the *temporal relationship* extension in `CMT2-rs`. Specifically, for the current rule r_cur, no matter *method* or *always*, it is guarded by a predecessor rule's execution history and communicates with the predecessor rule across cycles. It comprises four parts: (1) *predecessor declaration*: alias the predecessor rule (r_pre) to a local identifier (p); (2) *channel declaration*: specifies the output channels (o_ch); (3) *temporal guard*: specifies the predecessor's execution history as part of the current rule's guard through the `delay`, `dyndelay`, or `eagerdelay` operators; (4) *message passing*: calls the builtin `recv` and `send` methods of the predecessor rule's channel (p.ch) and the current rule's channel (o_ch) to deliver messages.

Semantics. The `delay` operator specifies an interval of the fixed number of clock cycles as one guard condition, named *latency-sensitive guard*. For example, `p.delay(k)` will hold exactly at clock cycle T+k if the predecessor rule p fires at clock cycle T. If the current rule cannot fire at clock cycle T+k due to other guard conditions' failure, the latency-sensitive guard will *expire* and will not hold in subsequent clock cycles. The `dyndelay` operator, on the other hand, specifies a variant interval of a minimum number of clock cycles as one guard condition, named *latency-insensitive guard*. That is, `p.dyndelay(k)`'s holding time will start from clock cycle T+k until a successful firing of the current rule, if the predecessor rule p fires at clock cycle T. The `eagerdelay` operator is a variant of `delay`, and is used only when the predecessor rule p is a *multi-cycle rule* (Section 3.2.2), and the delay countdown starts when p starts firing. Temporal relationships' semantics require that *temporal guards* must be coordinated with the *channel-based message passing*, named **guard-message atomicity**. Specifically, when a latency-sensitive guard expires, its carried message will be discarded from the channel at the same clock cycle; for a latency-insensitive guard, a message will remain in the channel until the guarded rule successfully fires. This temporal property helps avoid inter-cycle producer-consumer mismatch bugs: *the produced data does not last long enough for the consumer to consume or does not arrive when the consumer is ready to consume.* Figure 3 shows an example. Figure 3a uses intra-cycle hardware transactions: the producer rule prod writes the data io.x to the register r, and it calls shift2.enable to set the first stage of the 2-stage boolean shift registers shift2 high. Its time diagram is shown in Figure 3c. The consumer rule cons is fired two cycles after the producer rule prod. However, it cannot read the data x1, which has been overwritten by the data x2, causing a *data loss* bug.
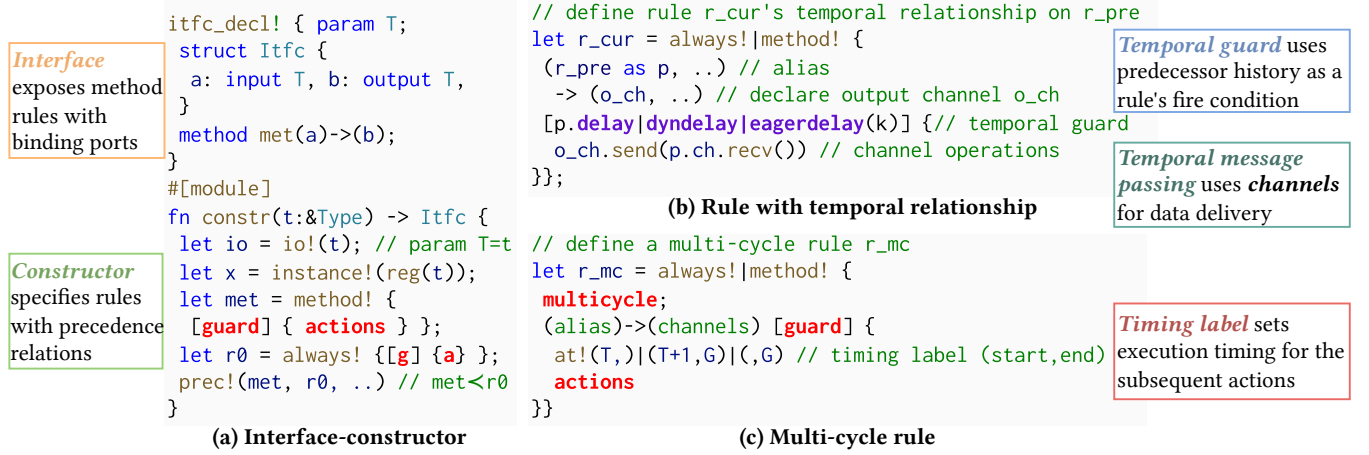
(a) Interface-constructor

(b) Rule with temporal relationship

(c) Multi-cycle rule

Figure 2: Syntax of `CMT2-rs` provides a unified description for intra-cycle and temporal hardware transactions.

(a) intra-cycle hardware transactions

(b) rules with temporal relationships

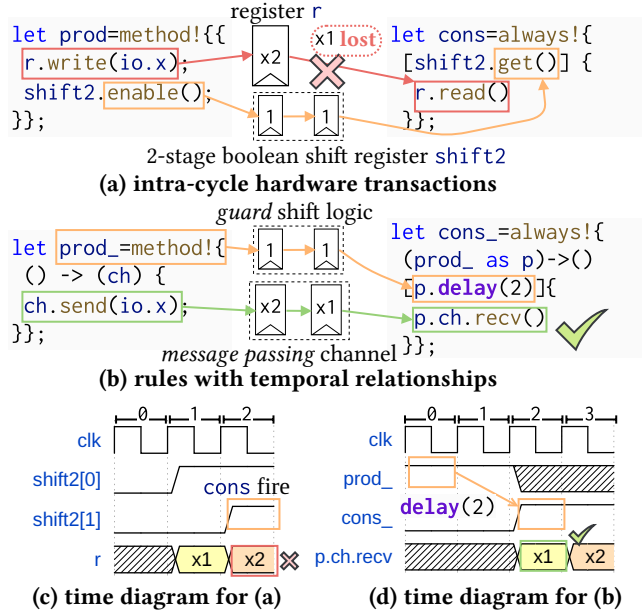**(c) time diagram for (a)**   **(d) time diagram for (b)**

Figure 3: Avoid producer-consumer mismatch.

Instead, Figure 3b adopts the temporal relationship: the producer rule prod_ sends io.x to its output channel ch, and the consumer rule cons_ uses prod_ as the predecessor rule p and is fired two cycles after prod_ with p.delay(2) as guard, as shown by the time diagram in Figure 3d. The *guard-message atomicity* guarantees that when cons_ is fired at cycle 2, it can receive the data x1 coordinated with the rule firing, avoiding data loss.
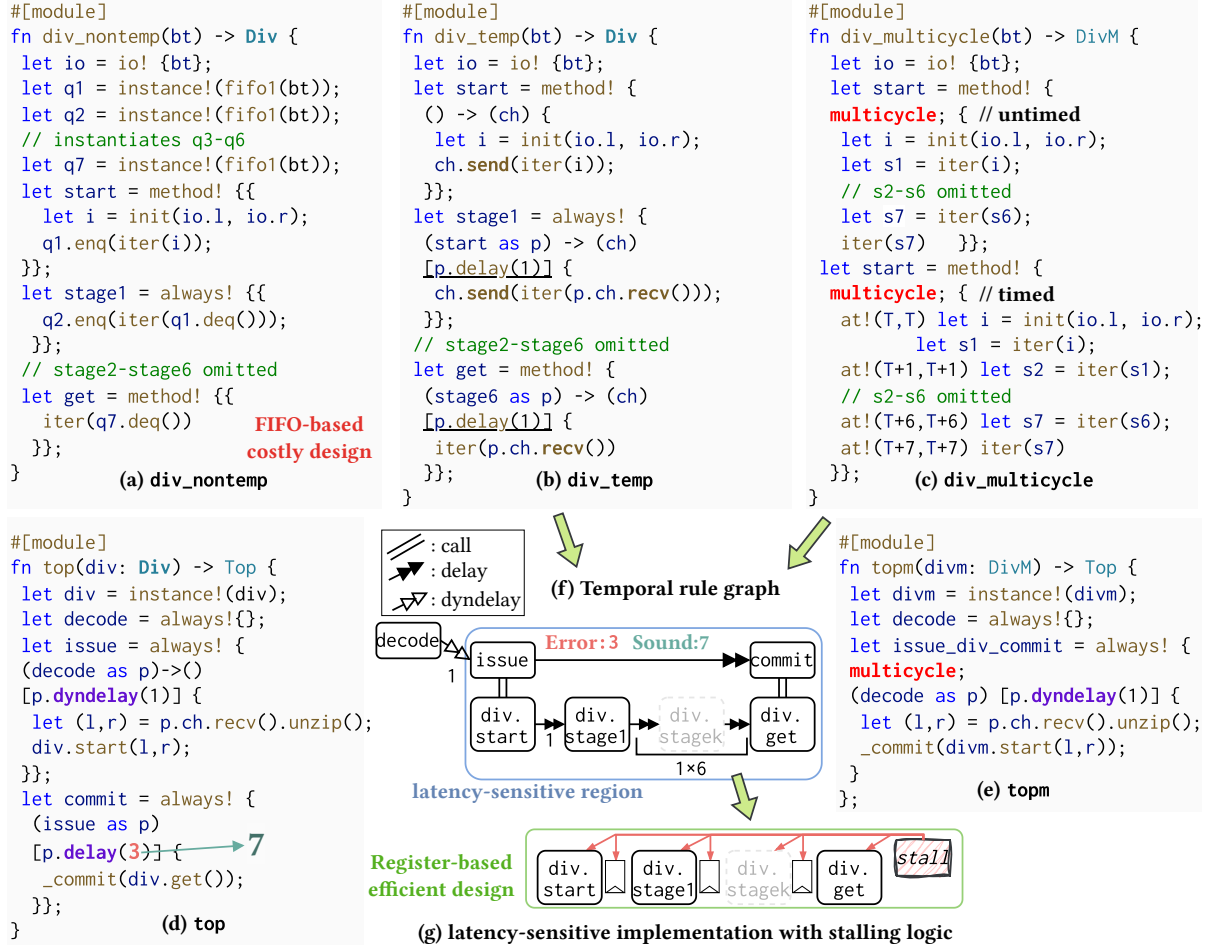
Hardware implementation. Both *temporal guard* and *temporal message passing* correspond to shifting logic: *temporal guard* shifts the firing history of predecessor rules, and *temporal message passing* shifts messages to keep pace with the guard. The delay is implemented as the efficient shift register for both guards and messages, while the dyndelay is implemented as a FIFO.

Figure 4 shows the description of an 8-bit restoring division pipeline in CMT2-rs. Function init and iter include actions for the computation initialization and iteration, respectively, which can be considered as two combinational logic blocks. The (a)div_nontemp constructor uses intra-cycle hardware transactions, describing the pipeline stages as rules (start, stage1, ..., get), each of which manipulates manually instantiated FIFO instances (q1, ..., q7) to coordinate one-by-one rule execution and deliver results. The (b)div_temp constructor adopts *temporal relationships* to describe the pipeline more intuitively with p.delay(1) temporal guards, and it uses channels to deliver results.

3.2.2 Multi-Cycle Rules. Although *temporal relationships* opens the door for temporal behavior modeling, it has two-fold drawbacks. First, it still requires designers to describe multiple rules and specify temporal relationships among them manually, which is verbose. As exemplified in Figure 4, the (b)div_temp constructor cannot reduce the number of lines of description code. Second, intra-cycle rules cannot be adjusted by the compiler, such as splitting a rule with a long critical path into multiple for frequency improvement. To tackle these problems, the *temporal hardware transaction* abstraction further introduces *multi-cycle rules*.

Syntax. Figure 2c shows the syntax of the *multi-cycle rule* extension in CMT2-rs. A *multi-cycle rule* definition is distinguished by the multicycle keyword. The firing actions of a multi-cycle rule can be specified exactly the same as the intra-cycle rules. We provide the *timing label* mechanism to specify the firing time (start time, finish time, or both) of the subsequent actions by the at! statements. Every timing label is associated with a *timing variable* (e.g., T and G) and an optional constant offset (e.g., 1 in T+1).

Semantics. When one multi-cycle rule is fired, its firing actions will be executed exactly *once* in the subsequent clock cycles, named **multi-cycle atomicity**. That is, when one action cannot execute in the current cycle due to *dependency violations*, guard failures, or *precedence violations*, the action will be retried in the subsequent cycles until it successfully executes. Action execution in multi-cycle rules must observe: (a) *data dependency*, one action cannot be fired until all the input values are valid; (b) *physical-timing dependency*,

```
#[module]
fn div_nontemp(bt) -> Div {
 let io = io! {bt};
 let q1 = instance!(fifo1(bt));
 let q2 = instance!(fifo1(bt));
 // instantiates q3-q6
 let q7 = instance!(fifo1(bt));
 let start = method! {{
   let i = init(io.l, io.r);
   q1.enq(iter(i));
 }};
 let stage1 = always! {{
   q2.enq(iter(q1.deq()));
  }};
 // stage2-stage6 omitted
 let get = method! {{
   iter(q7.deq())          FIFO-based
  }};                      costly design
}
        (a) div_nontemp
```

```
#[module]
fn div_temp(bt) -> Div {
 let io = io! {bt};
 let start = method! {
  () -> (ch) {
   let i = init(io.l, io.r);
   ch.send(iter(i));
  }};
 let stage1 = always! {
  (start as p) -> (ch)
  [p.delay(1)] {
   ch.send(iter(p.ch.recv()));
  }};
 // stage2-stage6 omitted
 let get = method! {
  (stage6 as p) -> (ch)
  [p.delay(1)] {
   iter(p.ch.recv())
  }};                (b) div_temp
}
```

```
#[module]
fn div_multicycle(bt) -> DivM {
  let io = io! {bt};
  let start = method! {
  multicycle; { // untimed
   let i = init(io.l, io.r);
   let s1 = iter(i);
   // s2-s6 omitted
   let s7 = iter(s6);
   iter(s7)   }};
 let start = method! {
  multicycle; { // timed
   at!(T,T) let i = init(io.l, io.r);
           let s1 = iter(i);
   at!(T+1,T+1) let s2 = iter(s1);
   // s2-s6 omitted
   at!(T+6,T+6) let s7 = iter(s6);
   at!(T+7,T+7) iter(s7)
  }};        (c) div_multicycle
}
```

```
#[module]
fn top(div: Div) -> Top {
  let div = instance!(div);
  let decode = always!{};
  let issue = always! {
  (decode as p)->()
  [p.dyndelay(1)] {
   let (l,r) = p.ch.recv().unzip();
   div.start(l,r);
  }};
  let commit = always! {
   (issue as p)
   [p.delay(3)] {
    _commit(div.get());
  }};
                (d) top
}
```


(f) Temporal rule graph

latency-sensitive region

Register-based efficient design

(g) latency-sensitive implementation with stalling logic

```
#[module]
fn topm(divm: DivM) -> Top {
  let divm = instance!(divm);
  let decode = always!{};
  let issue_div_commit = always! {
  multicycle;
  (decode as p) [p.dyndelay(1)] {
   let (l,r) = p.ch.recv().unzip();
   _commit(divm.start(l,r));
  }
};                 (e) topm
```

Figure 4: Implementations of a 8-bit restoring division pipeline with *temporal hardware transactions*.

one action path of total delay exceeding the target clock period must be fired in different cycles. In a timing label T+k, the timing variable T represents a certain clock cycle, and the constant k represents a latency-sensitive offset. T+k indicates the timing point k cycles after cycle T. For an action whose timing is specified by at!(T+1,G), the action will start firing one cycle after T and finish firing at the clock cycle G. Different timing variables (T and G) indicate the unresolved latency, modeling the latency-insensitive temporal behavior.

A multi-cycle rule can be either *timed*, with a determined and legal *schedule*, or *untimed*, with unspecified action execution timing. Timed multi-cycle rules give designers precise cycle-level control over action execution, while untimed multi-cycle rules relieve designers from scheduling, leaving the tasks to the *temporal scheduling* algorithm in Figure 4. Untimed multi-cycle rules are *reusable*: they can have different *schedules* for different configurations (e.g., target technology and frequency). In Figure 4, the (c)div_multicycle constructor describes the division pipeline by either an untimed or a timed multi-cycle rule. For the timed multi-cycle rule, all timing labels have the same timing variable T plus constant offsets, indicating *latency-sensitive* behavior.

3.2.3   Inter-cycle analysis and optimizations. Temporal hardware transactions boost productivity with intuitive syntax, and improve design robustness with semantic guarantees, including *guard-message atomicity* and *multi-cycle atomicity*. They also enable inter-cycle analysis and optimizations for compiler-enforced correctness and efficiency, beyond the scope of HDLs due to temporal unawareness.

Temporal rule graph. A temporal hardware transaction module with temporal relationships and *timed* multi-cycle rules can be abstracted as a unified *temporal rule graph* representation for analysis, checking, and optimization, where vertices represent *intra-cycle* rules and three types of edges represent different relationships among them: *call*, *delay*, and *dyndelay*. For example, Figure 4d presents the *temporal rule graph* for both the module (d)top instantiating (b)div_temp and the module (e)topm instantiating (c)div_multicycle. A *latency-sensitive region* is defined as a *connected component* of the *temporal rule graph* with only *call* and *delay* edges remaining. In Figure 4d, issue, commit, and the divider rules form a latency-sensitive region.

Timing inference and rule coordination checking. The compiler conducts *timing inference* for each latency-sensitive region

in a bottom-up manner. In Figure 4, `commit` can be inferred to fire 7 clock cycles after `issue`: `commit` fires together with the callee `div.get`, which fires 7 cycles after `issue`. The compiler leverages the inferred timing information for *rule coordination checking*. In Figure 4d, `commit` has the *temporal guard* `issue.delay(3)`, which mismatches the inferred delay 7. This mis-coordination causes the `commit` rule never to fire since the temporal guard holds when `div.get` cannot be called. This causes *deadlock*s in practice. For example, permanent commit failure in a CPU pipeline makes the scoreboard full and stalls the pipeline forever. The *rule coordination checking* reports such mis-coordination as a *temporal error* and provides a fix suggestion with the inferred timing.

Temporal relationship pruning. We define *redundant* temporal relationships as the ones implied by either *guarded atomicity* of its target rule or other temporal relationships ending at the same target rule. The compiler conducts *temporal relationship pruning* together with the bottom-up *timing inference*. When the timing of a rule is inferred, the pruning process removes all the temporal relationships targeting the rule if the timing is implied by its callees, or only keeps the one of the shortest delay. From the perspective of the *temporal rule graph*, the pruning transforms each *latency-sensitive region* into a *tree*, keeping the temporal behavior unchanged with implementation simplified. It is worth noting that the *temporal relationship pruning* only removes *temporal guards*, with all *temporal message passing channels* remaining.

False static prevention. A *false-static* pattern denotes that a rule has one *delay* guard accompanied by other guards after *temporal relationship pruning*. When the *delay* guard holds but any other guard fails, the rule cannot fire, and the latency-sensitive guard will expire immediately with the message discarded, causing **data loss** bugs. The compiler detects and reports *warning*s on the *false-static* pattern, recommending *dyndelay* operators for robustness.

3.2.4   Hybrid latency-sensitive/-insensitive temporal behavior support. Both *temporal relationships* and *multi-cycle rules* can describe hybrid latency-sensitive/-insensitive temporal behavior. Any hybrid temporal behavior can be represented as *latency-sensitive regions* connected by *dyndelay* edges in the *temporal rule graph*. Figure 4f presents one example, where decode can be treated as a latency-sensitive region of a single rule, and it is connected, with a *dyndelay* edge, to the latency-insensitive region including the `issue` rule. The compiler automatically implements the hybrid behavior by adding minimal *stall* logic, wrapping each latency-sensitive region with only one *stall* controller, which controls the state updating of the whole region, including both the rules and the temporal relationship implementations, as shown in Figure 4g. The *stall* controller watches all the *dyndelay* edges connected to the region. When any of them is blocked, either being full for sending or empty for receiving, the *stalling* logic will stall the whole region, preventing wrong behavior such as data loss or using invalid data.

4   Compiler

Cement2's compiler is built around the Cement2 Transaction Intermediate Representation (CTIR) and conducts multi-phase synthesis to generate rich backends. CTIR provides a unified representation for both *temporal* and intra-cycle hardware transaction abstraction. Figure 5a shows the CTIR representation of the untimed multi-cycle
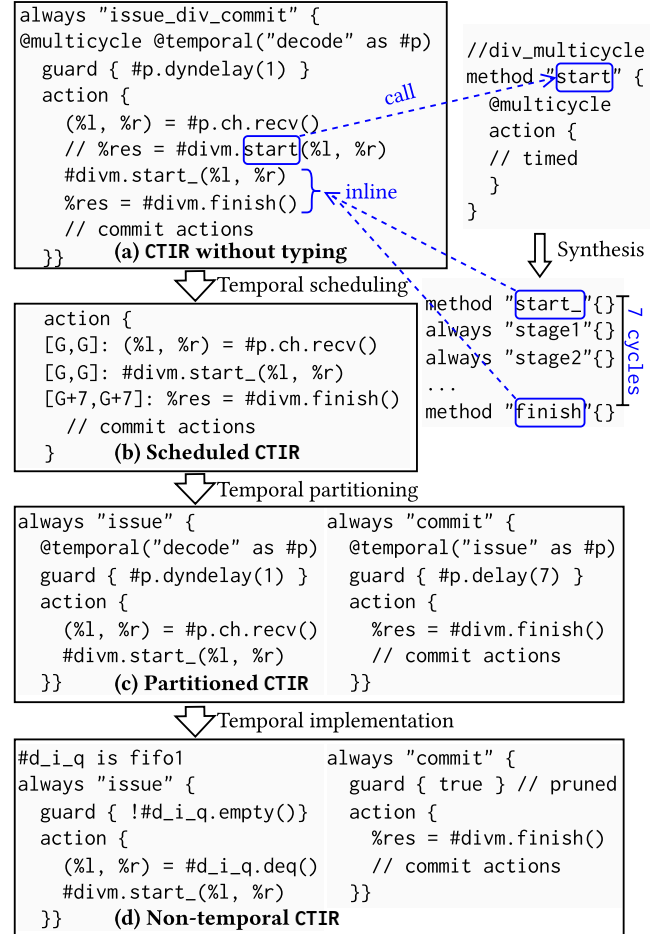

Figure 5: CTIR and synthesis flow

rule `topm` in Figure 4e. It shows that CTIR strictly aligns with the language features of CMT2-rs, facilitating straightforward IR construction. For example, *timing labels* in multi-cycle rules appear in CTIR, as shown in Figure 5b, instructing the compiler to generate expected schedules and implementations.

The compiler translates CTIR of *temporal hardware transaction* features into a low-level description of the same functionality. The synthesizer works in a *bottom-up* manner: any callee multi-cycle rules should be synthesized before the current rule. When synthesizing a multi-cycle rule, the compiler will first *inline* all method calls to other multi-cycle rules. It will substitute the original method call with a call to the entry method among the synthesized intra-cycle rules. Besides, the compiler also inserts method calls to get the result values. For example, in Figure 5a, the method call to `#divm.start` is replaced by the partitioned `#divm.start_`, and a call to `#divm.finish` is added to retrieve the result `%res`. The synthesis process comprises the following phases:

Temporal scheduling. The *temporal scheduling* phase generates legal schedules for untimed multi-cycle rules. Since *temporal hardware transactions* support hybrid latency-sensitive/-insensitive temporal behavior as discussed in Section 3.2.4, we cannot adopt existing scheduling algorithms, such as SDC [17], since they only

assign operations into *static* timing positions. Instead, we introduce a custom ASAP (As-Soon-As-Possible) scheduling algorithm. The scheduler iterates the actions of the current rule and determines their earliest firing timing in order. For example, the scheduler groups the action `#p.ch.recv` and `#divm.start_` to be fired at clock cycle G, assuming the *physical-timing dependency* between them is satisfied. Next, the scheduler puts the action `#divm.finish` into the cycle G+7 according to the timing reported from `#divm`. The remaining commit actions consuming `%res` must be scheduled not earlier than G+7 to satisfy the *data dependency*. For hybrid latency-sensitive/-insensitive temporal behavior, our ASAP scheduler will automatically create new timing variables for actions whose firing time cannot be determined statically. For example, if `#divm` is a latency-insensitive division unit, the firing time of the action `#divm.finish` cannot be assigned a timing label in the G+k form. Instead, the scheduler will create a new timing variable (e.g., T) to indicate the firing time. The remaining commit actions will be scheduled at later cycles than T.

Our scheduler performs *retiming* on multi-cycle rules. Specifically, given a technology-specific propagation delay model and a clock period target, the scheduler considers the *physical-timing dependencies* for each action: it iterates the scheduled actions that start a data dependency path to the current action, and use the propogation delay model to estimate the path delay to determine the earliest clock cycle that the current action can be scheduled.

Temporal partitioning. The *temporal partitioning* pass partitions a timed multi-cycle rule into an equivalent set of intra-cycle rules with *temporal relationships* generated. It involves three steps: (1) building a new rule for every group of actions that have the same timing label; (2) creating *temporal guards* between new rules according to the schedule; (3) inserting *temporal message passing* channels and actions according to *data dependency*. For example, in Figure 5b, the partitioning pass creates two new rules `issue` and `commit` from the scheduled groups, and inserts a *latency-sensitive guard* of the 7-cycle delay between them, generating Figure 5c. This example does not create message passing channels since there is no data dependency across rule boundaries. For any dependency between two rules whose timing labels have different timing variables, the compiler will create a *latency-insensitive guard* to enforce the dependency and a channel for data delivery if required, which guarantees *multi-cycle atomicity*. This phase only creates necessary latency-insensitive logic, keeping the implementation efficient.

Temporal implementation. The *temporal implementation* pass translates *temporal guards* and *message passing* channels into non-temporal instances and actions. Before the implementation, the compiler builds the *temporal rule graph* representation and conducts *inter-cycle* analysis, checking, and optimizations, as described in Section 3.2.3. In Figure 5d, `d_i_q` is instantiated to implement the temporal guard and message passing between rule `issue` and its predecessor `decode`. The temporal relationship between `issue` and `commit` is optimized out by *temporal relationship pruning*.

The `Cement2` compiler supports rich backends. For RTL generation, it implements an efficient hardware transaction synthesis algorithm [1, 10]. Specifically, it translates CTIR into FIRRTL [27], and generates optimized SystemVerilog using firtool-1.108.0 [15] for Vivado synthesis and FPGA deployment. `CMT2-rs` supports transactional testbenches: test stimuli are programmed as *rules*,

Table 2: Temporal features in evaluation designs

Design	Temporal relationship	Multi-cycle rule	SLOC	Synth. interface
Soft processor	✓	untimed	571	LI
Rgba2gray, sobel, ..	✓	untimed	86	Hybrid
Polybench kernels	✓	untimed	771	LS
Systolic array	✓	timed	35	LS

Table 3: Evaluation results of RISC-V soft cores.

	Sodor [3]	HF [29]	CMT2-RV	CMT2-RV + Rgba2gray, sobel, ..
CPI	1.389	1.389	1.386	-
Frequency	367MHz	287MHz	377MHz	316MHz
LUT	1974	3055	1614	2729
FF	924	2829	779	1152

which peek and poke instances through method calls. `Cement2` supports the RTL simulator generator Verilator [50] and generates C++ harness code from testbenches to drive the RTL simulators.

5 Evaluation

We evaluate `Cement2` from four perspectives, including a general soft processor, custom instructions, linear algebra accelerators, and systolic array design. We discuss how temporal hardware transactions facilitate the case studies, as summarized in Table 2, and analyze design quality for FPGA implementation.

5.1 5-stage RISC-V Soft Processor

Implementation. We implement a RISC-V soft processor, denoted as `CMT2-RV`, in `CMT2-rs` according to the architecture of Sodor [3] (5-stage, fully bypassed). Our design adopts *temporal hardware transactions*, as illustrated in Figure 1c, especially using temporal relationships between pipeline stages and an *untimed* multi-cycle rule to describe backend behavior shown in Figure 1c. The multi-cycle rule is scheduled into latency-insensitive stages that support stalling for hazard resolution.

Baselines. We compare `CMT2-RV` with two baselines: Sodor in Chisel [7] and Sodor in HazardFlow [29]. Chisel is an embedded HDL for RTL design, and HazardFlow is an academic HDL featuring latency-insensitive pipeline description. Other approaches, including HLS [5, 30], Calyx [32, 41], and Filament [40], are not compared since they cannot describe the required architecture.

Methodology. To measure cycles per instruction (CPI), we use MachSuite [44] integer benchmarks and EEMBC CoreMark [2] benchmarks and run RTL simulation to collect the cycle counts. We use Verilator [50] v5.028 to simulate `CMT2-RV`. We synthesize and place-and-route all the cores with Vivado 2024.1, targeting an XCVU9P FPGA. We exclude memories for all the designs.

Results. RTL simulation shows that all five cores achieve almost the same CPI. It indicates that `Cement2`'s *temporal hardware transactions* have the expressiveness to describe the required processor features, including data forwarding and branch prediction. Table 3 shows that `CMT2-RV` has the frequency of 377MHz, higher than Sodor's 367MHz. The HazardFlow core only achieves 287MHz. For
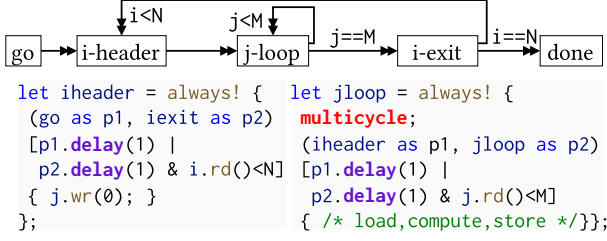
```
        i<N        j<M
  go → i-header → j-loop → i-exit → done
                   j==M      i==N
```

```
let iheader = always! {      let jloop = always! {
  (go as p1, iexit as p2)      multicycle;
  [p1.delay(1) |               (iheader as p1, jloop as p2)
   p2.delay(1) & i.rd()<N]    [p1.delay(1) |
  { j.wr(0); }                  p2.delay(1) & j.rd()<M]
};                          { /* load,compute,store */}};
```

Figure 6: Nested loops in `Cement2`

resource usage, HazardFlow uses 1.55× LUTs and 3.06× FFs compared to Sodor, which is a huge overhead. Instead, `CMT2-RV` uses 0.82× LUTs and 0.84× FFs compared to Sodor, indicating efficient resource usage. These results demonstrate that `Cement2`'s high-level abstraction does not sacrifice design quality for general hardware.

5.2   Custom CPU Instructions

With *temporal hardware transactions*, we can easily extend custom instructions to the `CMT2-RV` core in Section 5.1. By reserving an extension interface in `CMT2-RV`'s constructor, we can add new instructions by providing instruction encoding and rules for behavior description. We evaluate an image processing workload, *Edge Detection*, using the same methodology as Section 5.1 for evaluation. This case study aims to demonstrate that `Cement2` provides a convenient and efficient way to model and evaluate architecture decisions like adding instructions.

Implementation. We accelerate *Edge Detection* with custom instructions including `rgba2gray`, `sobel3x3`, `erode3x3`, and `dilate3x3`, all of which are described as untimed multi-cycle rules with the loop control described as temporal relationships according to software branching among basic blocks, as exemplified in Figure 6. The synthesized accelerators include hybrid latency-sensitive/-insensitive implementation as introduced in Section 3.2.4, since their behavior contains latency-insensitive memory access. It only takes 86 SLOC in total for the behavior description of these custom instructions, presenting software-level productivity.

Results. With the four custom instructions, the cycle count, reported by RTL simulation, is reduced by 75% compared to the original `Edge Detection` workload running without custom instructions. Table 3 shows the frequency and resource usage results of the extended `CMT2-RV` core. The synthesized frequency is 316MHz, which is 16% lower than the original `CMT2-RV` core. For resource usage, the extended `CMT2-RV` core uses 1.69× LUTs and 1.48× FFs compared to the original one, reporting the real hardware overheads. `Cement2` reports comprehensive performance results, including cycle count and frequency, and real resource usage, given the high-level description of the compound system, including both a processor and custom instructions. It saves human efforts to program at a tedious and error-prone low level to get the precise evaluation results.

5.3   Linear Algebra Kernels

We evaluate linear algebra kernels from PolyBench [34] to demonstrate `Cement2`'s capability for control-intensive designs.

Table 4: Performance and resource across Polybench kernels.

	SV	BSV[42]	Cement[52]	Vitis HLS[5]	Calyx[41]	Cement2
Cycle	1204	1.59×	0.92×	1.01×	3.83×	0.92×
Time	5.5µs	1.74×	1.14×	1.48×	2.73×	1.01×
LUT	431	1.87×	2.01×	1.33×	1.19×	0.87×
FF	169	3.06×	0.73×	2.23×	2.12×	0.81×

Implementation. We implement resource-efficient accelerators for PolyBench kernels, which contain nested loops to be described as hardware finite state machines (FSMs). `Cement2`'s abstraction facilitates such multi-cycle behavior description, as illustrated in Figure 6. The description concisely corresponds to the control flow graph (CFG) of the kernel, where transitions between basic blocks are described as temporal relationships, and complex basic blocks are described as multi-cycle rules, such as `jloop`. We implement 13 kernels from the benchmark suite, and the total SLOC is 771, including 121 rules (21 of them are multi-cycle rules, all untimed). All the kernels are latency-sensitive, since on-chip RAM access has a fixed latency. It only takes a Ph.D. student one day to implement. The same kernels in SystemVerilog [18] take 2610 SLOC in total, and it takes a Ph.D. student 1 week to implement and test.

Baselines. We compare `Cement2`-synthesized designs with five baselines: SystemVerilog for manual RTL design, BSV [42] for traditional rule-based design, Cement [52] and Dahlia-Calyx [39, 41] flow for automatic FSM generation from software-like description, and Vitis HLS 2024.1 [5] considered as the commercial state-of-the-art HLS tool. BSV generates FSMs by `Stmt` [9]. Static optimizations [32] are enabled in the Dahlia-Calyx flow. Filament [40] is not included since it cannot describe loop control, while dynamic HLS approaches [13, 30, 55] are also excluded since they present disadvantages against Vitis HLS on static kernels. All designs are configured to be sequential and use synchronous-read RAM for comparison fairness. We use Vivado 2024.1 for synthesis and place-and-route toward the XCVU9P FPGA with a target period of 7ns.

Results. Table 4 shows the geometric mean performance and resource utilization results across the PolyBench kernels. For performance, `Cement2` achieves better cycle counts (0.92×) and comparable execution time (1.01×) than the SystemVerilog baseline. The reason is that `Cement2` designs avoid redundant cycles compared to the manually crafted FSMs with acceptable frequency overhead. Although Cement and Vitis HLS achieve similar cycle counts as `Cement2`, `Cement2` designs achieve higher frequencies and faster execution time. The performance of BSV and Dahlia-Calyx designs is not competitive since their frontends introduce unnecessary idle cycles for loop control and computation. For resource usage, `Cement2` saves 13% LUTs and 19% registers than the SystemVerilog baseline. By observing Vivado-reported schematics, we found that `Cement2` compiler generates FSMs of one-hot encoded states natively from *temporal guards*, which is friendly for FPGA synthesis and can save FSM resources. Cement uses fewer registers than the SystemVerilog baseline but consumes the most LUTs due to extremely compact states and complicated transitions in the generated FSMs. Other approaches require more LUTs and registers than the SystemVerilog baseline. For example, BSV generates 24 rules for state transition, 42 rules for memory access, and 9 rules for FIFOs to implement the simple `atax` kernel. Most of the
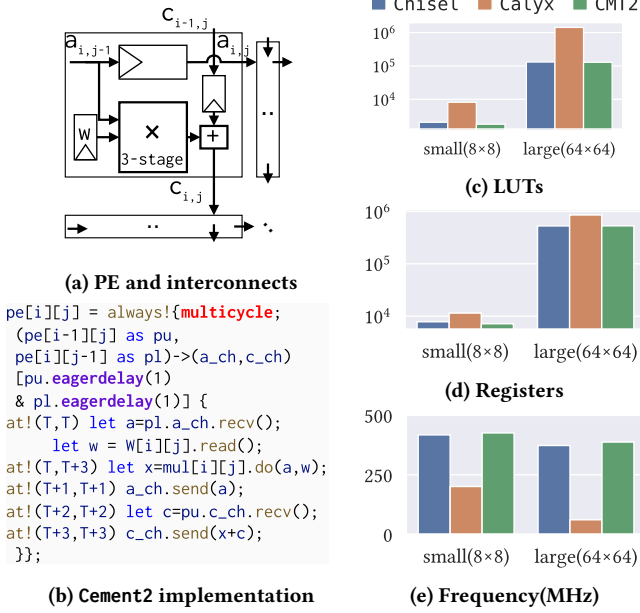
(a) PE and interconnects

```
pe[i][j] = always!{multicycle;
 (pe[i-1][j] as pu,
 pe[i][j-1] as pl)->(a_ch,c_ch)
 [pu.eagerdelay(1)
 & pl.eagerdelay(1)] {
at!(T,T) let a=pl.a_ch.recv();
     let w = W[i][j].read();
at!(T,T+3) let x=mul[i][j].do(a,w);
at!(T+1,T+1) a_ch.send(a);
at!(T+2,T+2) let c=pu.c_ch.recv();
at!(T+3,T+3) c_ch.send(x+c);
 }};
```

(b) Cement2 implementation

(c) LUTs

(d) Registers

(e) Frequency(MHz)

Figure 7: Systolic array implementation and evaluation.

verbose rules handle latency-insensitive actions. Cement2 avoids such overheads since its efficient synthesis flow does not introduce unnecessary latency-insensitive logic. The results demonstrate that Cement2 boosts productivity while achieving competitive performance against handcrafted RTL implementation.

5.4 Systolic Array Accelerators

We evaluate Cement2's applicability to high-performance architecture for FPGA acceleration.

Implementation. We implement a weight-stationary systolic array, whose processing element (PE) and interconnects are illustrated in Figure 7a. Such spatial architecture exhibits temporal behavior across clock cycles: every PE receives data from neighboring PEs, computes or temporarily stores data, and then sends data to other PEs. Thereby, Cement2's temporal hardware transactions can be easily applied to describe the behavior of the systolic array. As shown in Figure 7b, we describe the behavior of a PE as a *timed* multi-cycle rule and use temporal relationships to describe interconnects. Every PE gets data from its up and left neighbor, named pu and pl, respectively. The *temporal guard* of the current PE is set by two eagerdelay operators. The timing labels control the received data to be delayed by one cycle before being sent to channels, implementing systolic data movement. Besides, the timing labels also specify the pipeline depth of the multiplier. This demonstrates temporal hardware transactions' precise control over behavior to align with the architectural design intent. The complete temporal description of the systolic array passes the compiler checking in Section 3.2.3 to guarantee rule coordination and no data loss, and is synthesized into a fully latency-sensitive implementation.

Baselines. We compare the Cement2-synthesized systolic array with a high-performance weight-stationary systolic array in Chisel [36] and an output-stationary design produced by the newest Calyx systolic array generator [32]. We run Vivado 2024.1 with a target period

of 2.5ns for the XCU250 FPGA part. We target 32-bit fixed-point GEMM, and synthesize multipliers into DSP slices.

Results. Figure 7e shows the resource and frequency results. Cement2 designs save geomean 7% LUTs and 4% registers than Chisel designs, and save 86% LUTs and 38% registers than Calyx designs. The minor differences between the resource usages of the Cement2 designs and those of the Chisel designs demonstrate that Cement2 does not introduce unnecessary overhead than handcrafted designs. The significant resource savings of Cement2 over Calyx are due to Calyx's static inference of the clock-cycle timing of all data transfers among PEs and its generation of centralized FSMs to control all their execution, which incurs large overheads and worse frequencies. Cement2 designs achieve slightly higher frequencies, 1.03×, than those of Chisel. These results demonstrate Cement2's hardware quality for high-performance architecture.

6 Related Work

Embedded HDLs. Embedded HDLs leverage software languages for metaprogramming [7, 16, 28, 33, 43] and support flexible parameterization and construction. This line of work remains at the structural register-transfer level without raising design abstraction to provide better language promises about behavioral correctness.

Transactional HDLs. Transactional HDLs [11, 14, 42] embody the *guarded atomic action* concurrency model for digital hardware design. Synthesis algorithms [10, 24, 25] generate RTL circuits from rule descriptions. Prior rule-based language extensions [22, 31] generate additional hardware units like arbiters and reservation stations, causing unavoidable overheads. Although Cement2 introduces high-level abstraction, the temporal features are synthesized into efficient low-level implementations without introducing any unnecessary hardware components.

Type systems and models for hardware. Filament [40] and Aetherling [19] encode latency-sensitive timing properties in their type systems, while Shakeflow [23] and HazardFlow [29] introduce latency-insensitive combinator interfaces. Their type systems detect hardware issues such as resource conflicts and combinational loops. However, they are purpose-built for only latency-sensitive or latency-insensitive. Similarly, PDL [58], Cement [52], Spade [47], TL-Verilog [26], and Esterel [8] introduce language constructs to facilitate specific designs, such as control-intensive or pipeline circuits. Cement2's language features provide a general description of temporal behavior for hybrid latency-sensitive/-insensitive scenarios. Assassyn [51] provides an asynchronous programming model for both architecture and hardware design. Cement2's abstraction reaches a higher level to describe multi-cycle behavior while providing rich temporal analysis and synthesis capabilities.

High-level synthesis and IRs. High-level synthesis (HLS) [5, 12, 21, 30, 56, 60] starts with software programs and generates hardware implementations. They highly rely on synthesis algorithms like scheduling, which generate either static [17, 61], dynamic [30, 53, 54], or hybrid [13, 55, 57] circuits for design tradeoffs. However, HLS's abstraction is too high to expose enough control over hardware details, causing expressiveness limitations [4] and unpredictable performance [39]. Hardware IRs [32, 38, 41, 57] provide a middleground where high-level features co-exist with structural hardware. Cement2's high-level abstraction provides productivity

even close to software description for temporal behavior, but does not sacrifice expressiveness and low-level control.

7   Conclusion and Future Work

We present Cement2, a new FPGA programming approach with the raised abstraction, *temporal hardware transactions*, to provide an expressive description of temporal behavior for boosted productivity. We conduct comprehensive case studies, including soft processors, custom instructions, linear algebra kernels, and a systolic array, to demonstrate the effectiveness. For future work, we will perform real-world tasks like design exploration of out-of-order processors and heterogeneous systems with Cement2's help.

References

[1] 2025. B-Lang-org/bsc. https://github.com/B-Lang-org/bsc original-date: 2020-01-31T22:01:29Z.

[2] 2025. eembc/coremark. https://github.com/eembc/coremark original-date: 2018-05-23T00:53:13Z.

[3] 2025. ucb-bar/riscv-sodor. https://github.com/ucb-bar/riscv-sodor original-date: 2013-07-17T22:10:42Z.

[4] Abhinav Agarwal, Man Cheuk Ng, and Arvind. 2010. A Comparative Evaluation of High-Level Hardware Synthesis Using Reed–Solomon Decoder. *IEEE Embedded Systems Letters* 2, 3 (Sept. 2010), 72–76. https://doi.org/10.1109/LES.2010.2055231

[5] AMD Inc. 2025. Vitis High-Level Synthesis User Guide (UG1399). https://docs.amd.com/r/2024.2-English/ug1399-vitis-hls/Introduction

[6] AMD Inc. 2025. Vitis Libraries. https://docs.amd.com/r/en-US/Vitis_Libraries/index.html

[7] Jonathan Bachrach, Huy Vo, Brian Richards, Yunsup Lee, Andrew Waterman, Rimas Avižienis, John Wawrzynek, and Krste Asanović. 2012. Chisel: constructing hardware in a Scala embedded language. In *Proceedings of the 49th Annual Design Automation Conference (DAC '12)*. Association for Computing Machinery, New York, NY, USA, 1216–1225. https://doi.org/10.1145/2228360.2228584

[8] Gérard Berry and Georges Gonthier. 1992. The ESTEREL synchronous programming language: design, semantics, implementation. *Sci. Comput. Program.* 19, 2 (1992), 87–152. https://doi.org/10.1016/0167-6423(92)90005-V

[9] Bluespec Inc. [n. d.]. Bluespec SystemVerilog Language Reference Guide. https://github.com/B-Lang-org/bsc/releases/download/2025.01.1/BSV_lang_ref_guide.pdf

[10] Bluespec Inc. 2025. System and method for scheduling TRS rules. https://patents.justia.com/patent/7647567

[11] Thomas Bourgeat, Clément Pit-Claudel, Adam Chlipala, and Arvind. 2020. The essence of Bluespec: a core language for rule-based hardware design. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2020)*. Association for Computing Machinery, New York, NY, USA, 243–257. https://doi.org/10.1145/3385412.3385965

[12] Andrew Canis, Jongsok Choi, Mark Aldham, Victor Zhang, Ahmed Kammoona, Jason H. Anderson, Stephen Brown, and Tomasz Czajkowski. 2011. LegUp: high-level synthesis for FPGA-based processor/accelerator systems. In *Proceedings of the 19th ACM/SIGDA international symposium on Field programmable gate arrays (FPGA '11)*. Association for Computing Machinery, New York, NY, USA, 33–36. https://doi.org/10.1145/1950413.1950423

[13] Jianyi Cheng, Lana Josipovic, George A. Constantinides, Paolo Ienne, and John Wickerson. 2020. Combining Dynamic & Static Scheduling in High-level Synthesis. In *Proceedings of the 2020 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA '20)*. Association for Computing Machinery, New York, NY, USA, 288–298. https://doi.org/10.1145/3373087.3375297

[14] Joonwon Choi, Muralidaran Vijayaraghavan, Benjamin Sherman, Adam Chlipala, and Arvind. 2017. Kami: a platform for high-level parametric hardware specification and its modular verification. *Proc. ACM Program. Lang.* 1, ICFP (2017), 24:1–24:30. https://doi.org/10.1145/3110268

[15] CIRCT community. 2025. Release firtool-1.108.0. https://github.com/llvm/circt/releases/tag/firtool-1.108.0

[16] John Clow, Georgios Tzimpragos, Deeksha Dangwal, Sammy Guo, Joseph McMahan, and Timothy Sherwood. 2017. A pythonic approach for rapid hardware prototyping and instrumentation. In *2017 27th International Conference on Field Programmable Logic and Applications (FPL)*. 1–7. https://doi.org/10.23919/FPL.2017.8056860 ISSN: 1946-1488.

[17] Jason Cong and Zhiru Zhang. 2006. An efficient and versatile scheduling algorithm based on SDC formulation. In *Proceedings of the 43rd annual Design Automation Conference (DAC '06)*. Association for Computing Machinery, New York, NY, USA, 433–438. https://doi.org/10.1145/1146909.1147025

[18] Design Automation Standards Committee. 2024. IEEE Standard for SystemVerilog–Unified Hardware Design, Specification, and Verification Language. *IEEE Std 1800-2023 (Revision of IEEE Std 1800-2017)* (Feb. 2024), 1–1354. https://doi.org/10.1109/IEEESTD.2024.10458102 Conference Name: IEEE Std 1800-2023 (Revision of IEEE Std 1800-2017).

[19] David Durst, Matthew Feldman, Dillon Huff, David Akeley, Ross Daly, Gilbert Louis Bernstein, Marco Patrignani, Kayvon Fatahalian, and Pat Hanrahan. 2020. Type-directed scheduling of streaming accelerators. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2020)*. Association for Computing Machinery, New York, NY, USA, 408–422. https://doi.org/10.1145/3385412.3385983

[20] gem5.org. [n. d.]. gem5: Out of order CPU model. https://www.gem5.org/documentation/general_docs/cpu_models/O3CPU

[21] Google Inc. 2025. XLS: Accelerated HW Synthesis. https://google.github.io/xls/

[22] David J. Greaves. 2019. Further sub-cycle and multi-cycle schedulling support for Bluespec Verilog. In *Proceedings of the 17th ACM-IEEE International Conference on Formal Methods and Models for System Design (MEMOCODE '19)*. Association for Computing Machinery, New York, NY, USA, 1–11. https://doi.org/10.1145/3359986.3361199

[23] Sungsoo Han, Minseong Jang, and Jeehoon Kang. 2023. ShakeFlow: Functional Hardware Description with Latency-Insensitive Interface Combinators. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS 2023)*. Association for Computing Machinery, New York, NY, USA, 702–717. https://doi.org/10.1145/3575693.3575701

[24] J.C. Hoe and Arvind. 2000. Synthesis of operation-centric hardware descriptions. In *IEEE/ACM International Conference on Computer Aided Design. ICCAD - 2000. IEEE/ACM Digest of Technical Papers (Cat. No.00CH37140)*. 511–518. https://doi.org/10.1109/ICCAD.2000.896524 ISSN: 1092-3152.

[25] J.C. Hoe and Arvind. 2004. Operation-centric hardware description and synthesis. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 23, 9 (Sept. 2004), 1277–1288. https://doi.org/10.1109/TCAD.2004.833614 Conference Name: IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems.

[26] Steven F. Hoover. 2017. Timing-Abstract Circuit Design in Transaction-Level Verilog. In *2017 IEEE International Conference on Computer Design (ICCD)*. 525–532. https://doi.org/10.1109/ICCD.2017.91 ISSN: 1063-6404.

[27] Adam Izraelevitz, Jack Koenig, Patrick Li, Richard Lin, Angie Wang, Albert Magyar, Donggyu Kim, Colin Schmidt, Chick Markley, Jim Lawson, and Jonathan Bachrach. 2017. Reusability is FIRRTL ground: Hardware construction languages, compiler frameworks, and transformations. In *2017 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. 209–216. https://doi.org/10.1109/ICCAD.2017.8203780 ISSN: 1558-2434.

[28] Jane Street. 2025. Hardcaml is an OCaml library for designing hardware. https://github.com/janestreet/hardcaml

[29] Minseong Jang, Jungin Rhee, Woojin Lee, Shuangshuang Zhao, and Jeehoon Kang. 2024. Modular Hardware Design of Pipelined Circuits with Hazards. *Proceedings of the ACM on Programming Languages* 8, PLDI (June 2024), 28–51. https://doi.org/10.1145/3656378

[30] Lana Josipović, Radhika Ghosal, and Paolo Ienne. 2018. Dynamically Scheduled High-level Synthesis. In *Proceedings of the 2018 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA '18)*. Association for Computing Machinery, New York, NY, USA, 127–136. https://doi.org/10.1145/3174243.3174264

[31] Michal Karczmarek and Arvind. 2008. Synthesis from multi-cycle atomic actions as a solution to the timing closure problem. In *Proceedings of the 2008 IEEE/ACM International Conference on Computer-Aided Design (ICCAD '08)*. IEEE Press, San Jose, California, 24–31.

[32] Caleb Kim, Pai Li, Anshuman Mohan, Andrew Butt, Adrian Sampson, and Rachit Nigam. 2023. Unifying Static and Dynamic Intermediate Languages for Accelerator Generators. https://doi.org/10.48550/arXiv.2312.16300 arXiv:2312.16300 [cs].

[33] Derek Lockhart, Gary Zibrat, and Christopher Batten. 2014. PyMTL: A Unified Framework for Vertically Integrated Computer Architecture Research. In *Proceedings of the 47th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO-47)*. IEEE Computer Society, USA, 280–292. https://doi.org/10.1109/MICRO.2014.50

[34] Louis-Noel Pouchet and Tomofumi Yuki. 2018. PolyBench/C 4.2. https://sourceforge.net/projects/polybench/

[35] Jason Lowe-Power, Abdul Mutaal Ahmad, Ayaz Akram, Mohammad Alian, Rico Amslinger, Matteo Andreozzi, Adrià Armejach, Nils Asmussen, Brad Beckmann, Srikant Bharadwaj, Gabe Black, Gedare Bloom, Bobby R. Bruce, Daniel Rodrigues Carvalho, Jeronimo Castrillon, Lizhong Chen, Nicolas Derumigny, Stephan Diestelhorst, Wendy Elsasser, Carlos Escuin, Marjan Fariborz, Amin Farmahini-Farahani, Pouya Fotouhi, Ryan Gambord, Jayneel Gandhi, Dibakar Gope, Thomas Grass, Anthony Gutierrez, Bagus Hanindhito, Andreas Hansson, Swapnil Haria, Austin Harris, Timothy Hayes, Adrian Herrera, Matthew Horsnell, Syed Ali Raza Jafri, Radhika Jagtap, Hanhwi Jang, Reiley Jeyapaul, Timothy M. Jones, Matthias

Jung, Subash Kannoth, Hamidreza Khaleghzadeh, Yuetsu Kodama, Tushar Krishna, Tommaso Marinelli, Christian Menard, Andrea Mondelli, Miquel Moreto, Tiago Mück, Omar Naji, Krishnendra Nathella, Hoa Nguyen, Nikos Nikoleris, Lena E. Olson, Marc Orr, Binh Pham, Pablo Prieto, Trivikram Reddy, Alec Roelke, Mahyar Samani, Andreas Sandberg, Javier Setoain, Boris Shingarov, Matthew D. Sinclair, Tuan Ta, Rahul Thakur, Giacomo Travaglini, Michael Upton, Nilay Vaish, Ilias Vougioukas, William Wang, Zhengrong Wang, Norbert Wehn, Christian Weis, David A. Wood, Hongil Yoon, and Éder F. Zulian. 2020. The gem5 Simulator: Version 20.0+. https://doi.org/10.48550/arXiv.2007.03152 arXiv:2007.03152 [cs].

[36] Zizhang Luo, Liqiang Lu, Size Zheng, Jieming Yin, Jason Cong, Jianwei Yin, and Yun Liang. 2023. Rubick: A Synthesis Framework for Spatial Architectures via Dataflow Decomposition. In *2023 60th ACM/IEEE Design Automation Conference (DAC)*. 1–6. https://doi.org/10.1109/DAC56929.2023.10247743

[37] Jiacheng Ma, Gefei Zuo, Kevin Loughlin, Haoyang Zhang, Andrew Quinn, and Baris Kasikci. 2022. Debugging in the brave new world of reconfigurable hardware. In *Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 2022)*. Association for Computing Machinery, New York, NY, USA, 946–962. https://doi.org/10.1145/3503222.3507701

[38] Kingshuk Majumder and Uday Bondhugula. 2024. HIR: An MLIR-based Intermediate Representation for Hardware Accelerator Description. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 4 (ASPLOS '23)*. Association for Computing Machinery, New York, NY, USA, 189–201. https://doi.org/10.1145/3623278.3624767

[39] Rachit Nigam, Sachille Atapattu, Samuel Thomas, Zhijing Li, Theodore Bauer, Yuwei Ye, Apurva Koti, Adrian Sampson, and Zhiru Zhang. 2020. Predictable accelerator design with time-sensitive affine types. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2020)*. Association for Computing Machinery, New York, NY, USA, 393–407. https://doi.org/10.1145/3385412.3385974

[40] Rachit Nigam, Pedro Henrique Azevedo De Amorim, and Adrian Sampson. 2023. Modular Hardware Design with Timeline Types. *Proceedings of the ACM on Programming Languages* 7, PLDI (June 2023), 343–367. https://doi.org/10.1145/3591234

[41] Rachit Nigam, Samuel Thomas, Zhijing Li, and Adrian Sampson. 2021. A compiler infrastructure for accelerator generators. In *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*. Association for Computing Machinery, New York, NY, USA, 804–817. https://doi.org/10.1145/3445814.3446712

[42] R. Nikhil. 2004. Bluespec System Verilog: efficient, correct RTL from high level specifications. In *Proceedings. Second ACM and IEEE International Conference on Formal Methods and Models for Co-Design, 2004. MEMOCODE '04*. 69–70. https://doi.org/10.1109/MEMCOD.2004.1459818

[43] Charles Papon and Yindong Xiao. 2025. SpinalHDL. https://github.com/SpinalHDL/SpinalHDL original-date: 2015-01-25T11:42:00Z.

[44] Brandon Reagen, Robert Adolf, Yakun Sophia Shao, Gu-Yeon Wei, and David Brooks. 2014. MachSuite: Benchmarks for accelerator design and customized architectures. In *2014 IEEE International Symposium on Workload Characterization (IISWC)*. 110–119. https://doi.org/10.1109/IISWC.2014.6983050

[45] Daniel Sanchez and Christos Kozyrakis. 2013. ZSim: fast and accurate microarchitectural simulation of thousand-core systems. *SIGARCH Comput. Archit. News* 41, 3 (2013), 475–486. https://doi.org/10.1145/2508148.2485963

[46] Maico Cassel dos Santos, Tianyu Jia, Martin Cochet, Karthik Swaminathan, Joseph Zuckerman, Paolo Mantovani, Davide Giri, Jeff Jun Zhang, Erik Jens Loscalzo, Gabriele Tombesi, Kevin Tien, Nandhini Chandramoorthy, John-David Wellman, David Brooks, Gu-Yeon Wei, Kenneth Shepard, Luca P. Carloni, and Pradip Bose. 2022. A Scalable Methodology for Agile Chip Development with Open-Source Hardware Components. In *Proceedings of the 41st IEEE/ACM International Conference on Computer-Aided Design (ICCAD '22)*. Association for Computing Machinery, New York, NY, USA, 1–9. https://doi.org/10.1145/3508352.3561102

[47] Frans Skarman and Oscar Gustafsson. 2022. Spade: An HDL Inspired by Modern Software Languages. In *2022 32nd International Conference on Field-Programmable Logic and Applications (FPL)*. 454–455. https://doi.org/10.1109/FPL57034.2022.00075 ISSN: 1946-1488.

[48] Synopsys. 2025. VC SpyGlass Lint: RTL Design Analysis. https://www.synopsys.com/verification/static-and-formal-verification/vc-spyglass/vc-spyglass-lint.html

[49] Michael Bedford Taylor. 2018. Basejump STL: systemverilog needs a standard template library for hardware design. In *Proceedings of the 55th Annual Design Automation Conference (DAC '18)*. Association for Computing Machinery, New York, NY, USA, 1–6. https://doi.org/10.1145/3195970.3199848

[50] Veripool. 2025. Verilator. https://www.veripool.org/verilator/

[51] Jian Weng, Boyang Han, Derui Gao, Ruijie Gao, Wanning Zhang, An Zhong, Ceyu Xu, Jihao Xin, Yangzhixin Luo, Lisa Wu Wills, and Marco Canini. 2025. Assassyn: A Unified Abstraction for Architectural Simulation and Implementation. In *Proceedings of the 52nd Annual International Symposium on Computer Architecture (ISCA '25)*. Association for Computing Machinery, New York, NY, USA, 1464–1479. https://doi.org/10.1145/3695053.3731004

[52] Youwei Xiao, Zizhang Luo, Kexing Zhou, and Yun Liang. 2024. Cement: Streamlining FPGA Hardware Design with Cycle-Deterministic eHDL and Synthesis. In *Proceedings of the 2024 ACM/SIGDA International Symposium on Field Programmable Gate Arrays (FPGA '24)*. Association for Computing Machinery, New York, NY, USA, 211–222. https://doi.org/10.1145/3626202.3637561

[53] Jiahui Xu and Lana Josipovic. 2025. CRUSH: A Credit-Based Approach for Functional Unit Sharing in Dynamically Scheduled HLS. In *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1 (ASPLOS '25)*. Association for Computing Machinery, New York, NY, USA, 249–263. https://doi.org/10.1145/3669940.3707273

[54] Jiahui Xu and Lana Josipović. 2024. Suppressing Spurious Dynamism of Dataflow Circuits via Latency and Occupancy Balancing. In *Proceedings of the 2024 ACM/SIGDA International Symposium on Field Programmable Gate Arrays (FPGA '24)*. Association for Computing Machinery, New York, NY, USA, 188–198. https://doi.org/10.1145/3626202.3637570

[55] Jiahui Xu, Emmet Murphy, Jordi Cortadella, and Lana Josipovic. 2023. Eliminating Excessive Dynamism of Dataflow Circuits Using Model Checking. In *Proceedings of the 2023 ACM/SIGDA International Symposium on Field Programmable Gate Arrays (FPGA '23)*. Association for Computing Machinery, New York, NY, USA, 27–37. https://doi.org/10.1145/3543622.3573196

[56] Ruifan Xu, Jin Luo, and Yun Liang. 2024. Hermes: Enhancing Extensibility in High-Level Synthesis through Multi-Level IRs. In *Proceedings of the 2024 ACM/SIGDA International Symposium on Field Programmable Gate Arrays (FPGA '24)*. Association for Computing Machinery, New York, NY, USA, 186. https://doi.org/10.1145/3626202.3637606

[57] Ruifan Xu, Youwei Xiao, Jin Luo, and Yun Liang. 2022. HECTOR: A Multi-Level Intermediate Representation for Hardware Synthesis Methodologies. In *Proceedings of the 41st IEEE/ACM International Conference on Computer-Aided Design (ICCAD '22)*. Association for Computing Machinery, New York, NY, USA, 1–9. https://doi.org/10.1145/3508352.3549370

[58] Drew Zagieboylo, Charles Sherk, Gookwon Edward Suh, and Andrew C. Myers. 2022. PDL: a high-level hardware design language for pipelined processors. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI 2022)*. Association for Computing Machinery, New York, NY, USA, 719–732. https://doi.org/10.1145/3519939.3523455

[59] Sizhuo Zhang, Andrew Wright, Thomas Bourgeat, and Arvind Arvind. 2018. Composable Building Blocks to Open up Processor Design. In *2018 51st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 68–81. https://doi.org/10.1109/MICRO.2018.00015

[60] Zhiru Zhang, Yiping Fan, Wei Jiang, Guoling Han, Changqi Yang, and Jason Cong. 2008. AutoPilot: A Platform-Based ESL Synthesis System. In *High-Level Synthesis: From Algorithm to Digital Circuit*, Philippe Coussy and Adam Morawiec (Eds.). Springer Netherlands, Dordrecht, 99–112. https://doi.org/10.1007/978-1-4020-8588-8_6

[61] Zhiru Zhang and Bin Liu. 2013. SDC-based modulo scheduling for pipeline synthesis. In *2013 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. 211–218. https://doi.org/10.1109/ICCAD.2013.6691121 ISSN: 1558-2434.

[62] Gefei Zuo, Jiacheng Ma, Andrew Quinn, and Baris Kasikci. 2023. Vidi: Record Replay for Reconfigurable Hardware. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3 (ASPLOS 2023)*. Association for Computing Machinery, New York, NY, USA, 806–820. https://doi.org/10.1145/3582016.3582040