

PANTOMIME: Simulation-Based Leakage Proofs for Hardware Side-Channel Security

Robin Webbers
Vrije Universiteit Amsterdam
Amsterdam, Netherlands
r.m.a.webbers@vu.nl

Robert Schenck
Vrije Universiteit Amsterdam
Amsterdam, Netherlands
r@bert.lol

Alp Adnan Basar
Vrije Universiteit Amsterdam
Amsterdam, Netherlands
a.a.basar@vu.nl

Kristina Sojakova
Vrije Universiteit Amsterdam
Amsterdam, Netherlands
k.sojakova@gmail.com

Klaus v. Gleissenthall
Vrije Universiteit Amsterdam
Amsterdam, Netherlands
k.freiherrvongleissenthal@vu.nl

Abstract

Tools for verifying leakage descriptions of hardware aim to ensure that a given hardware design doesn't leak secrets via its microarchitecture when executing programs with appropriate countermeasures. However, current techniques either don't allow for leakage descriptions expressive enough to capture real-world software countermeasures like constant time programming, or they rely on expensive solvers and handwritten invariants, making them difficult to apply to larger designs—especially for hardware designers who are not experts in formal verification. In this paper, we present a new approach to leakage verification: *simulation-based leakage proofs*, where proofs are developed alongside the design.

Inspired by techniques in cryptography, simulation-based leakage proofs show that a leakage description correctly captures a hardware design by constructing a simulator—another hardware design that must faithfully replicate all attacker-observable behavior from explicitly leaked secrets. Simulation-based leakage proofs are easy to write and to debug: writing a simulator just means writing another hardware module, which is what hardware designers are already best at. They also capture all common software defenses and make proof checking orders of magnitude faster.

We implemented simulation-based leakage proofs in PANTOMIME, a tool that supports writing processors and their leakage proofs in Haskell; we report on using PANTOMIME to write and verify AIM-CORE, a 5-stage in-order processor. Unlike previous leakage proofs, which are conditional on the unproven functional correctness of the CPU, our proofs hold unconditionally.

Keywords

Hardware, Side-Channels, Leakage

1 Introduction

Context. Side-channel attacks via cache and timing can leak secrets used in private computations [22, 25, 30, 31, 45]. To prevent side-channels in software, programmers use countermeasures like branch balancing [1, 6, 44], and constant-time (or data-oblivious) programming [2, 9, 12, 43]. Whether these countermeasures are effective crucially depends on the underlying hardware. Unfortunately, modern hardware designs are mind-bendingly complex due to their highly parallel nature, their many microarchitectural optimizations like fast paths [3], pre-fetching [13], and speculative

execution [21, 28, 33], and their various micro-architectural buffers [39, 40]. This complexity makes it hard to manually audit even simple designs to check whether software with appropriate defenses will truly execute securely.

Problem. To increase our confidence that hardware designs are keeping up their end of the promise, a string of recent work aims to formally verify existing open-source hardware designs against descriptions of their intended leakage [4, 5, 16–18, 20, 36–38, 41]. While there has been tremendous progress in the area and these techniques have now managed to verify leakage descriptions for a number of open-source processors, their focus on verifying existing designs comes with a number of drawbacks. Their leakage descriptions are either not expressive enough to capture real-world defenses like the constant-time programming discipline [4, 5, 16–18], or they heavily rely on expensive solvers—either to find inductive invariants [37, 38, 41], or to exhaustively explore the design's state space [20, 36]—which limits their applicability to larger designs. As an alternative to fully automated proofs via solvers, one can instead ask the user to supply missing inductive invariants by hand. In fact, several existing methods already require users to supply certain hard-to-find invariants manually [18, 41]. While this helps with scaling, it places a heavy burden on the user: inductive invariants are difficult to come up with by hand and even harder to debug when they are wrong [27]. This is especially true for hardware designers, unless they also happen to be experts in formal verification.

Our Solution. In this paper, we propose simulation-based leakage proofs—a co-design approach to verifying hardware against explicit leakage descriptions, where proofs are developed alongside the code. Based on insights borrowed from cryptography [10, 24], a simulation-based leakage proof demonstrates the correctness of a leakage description by constructing a *simulator*—another hardware circuit which must faithfully replicate all attacker-observable behavior of the original design, while only being granted access to explicitly leaked secrets. As simulators are just programs, they can be written, debugged, and executed like any other hardware design. This makes simulation-based leakage proofs a good candidate for proofs that are written alongside the design by the hardware developers themselves. Indeed, the main proof effort lies in doing what hardware designers are already best at: writing hardware designs. Simulation-based leakage proofs are expressive and capture all common software defenses against cache and timing side-channels.

They are also fast to check, as we only need to show single-step equivalence between the simulator and the original design.

Simulation-Based Leakage Proofs. To show that a hardware circuit c is secure via a simulation-based leakage proof, we first have to provide a precise description of its intended leakage in the form of a circuit $leak$. For example, if c is an adder with a fast path on input 0 , the leakage description $leak$ reveals whether or not the input is 0 . Next, we model the attacker’s view of the computation as a circuit obs . For example, obs may reveal the time it takes to complete the computation by indicating whether an output was produced in a given clock cycle. To prove that the leakage description $leak$ faithfully captures everything an attacker can learn about potential secret information processed by c , we have to construct a *simulator* circuit sim such that the original circuit composed with the attacker observation function obs , written as $c \circ obs$, is *indistinguishable* from the composition of the leakage description $leak$ and the simulator sim , written as $leak \circ sim$. In our example, simulator sim reproduces the timing behavior of c by delaying inputs that don’t take the fast path. The existence of a simulator means that we can reconstruct all attacker-observable behavior from information that has been explicitly leaked. In turn, this means the attacker learns no more than what is specified via the explicit leakage description.

Functional Programs as Hardware. While simulation-based leakage proofs can be applied to hardware designs written in any language, we present a concrete instance of our proof method for circuits expressed as *functional programs*. Building on a long line of work connecting functional programs and hardware [8, 34], we represent the hardware’s single clock tick transition function as a functional program that takes a state and input to a state and output.

Equivalence Via State Projection. Next, we construct a proof system for proving equivalence between circuits. We build on the insight that functions describing hardware behavior *within* a clock tick are *loop-free* and therefore allow for easy equivalence proofs. Interestingly, our proofs generally require us to prove an equivalence between functions of *different types*. To enable these proofs, we propose a new *state projection rule*, which allows changing the type of a circuit’s state as long as this change doesn’t affect its observable behavior, thereby establishing a refinement relation between the two circuits [23]. Except for state projection functions—which translate between the states of different types and have to be provided by the user—equivalence proofs can be fully automated via an SMT solver.

Implementation and Evaluation. We implemented our method in PANTOMIME, a tool for writing processors and proofs in unrestricted Haskell (including support for higher-order functions, type classes, abstract data types, and monads). Hardware descriptions written with PANTOMIME can be extracted to hardware description languages using C λ aSH [7]. To prove equivalence between simulator and implementation, we developed a new symbolic execution engine based on Griset [26]. The engine proves function equivalence between expressions in GHC Core using the state projection rule and interprets C λ aSH data types to accurately model the behavior of the extracted hardware. We used PANTOMIME to write AIMCORE, a 5-stage RISC-V CPU that supports the full base integer




Figure 1: Circuits take state of type S , an input of type I , and return a pair of type (S, O) .

instruction set, and verify it against a leakage description. Unlike [36, 41], where leakage proofs are conditional on the (unproven) functional correctness of the CPU, our proof of AIMCORE is unconditional. Moreover, proof checking is orders of magnitude faster than previous methods, making PANTOMIME suitable for interactive use.

Contributions. In summary, we make the following contributions.

- **Simulation Based Proofs:** A co-design approach for verifying leakage descriptions of hardware circuits.
- **The State Projection Rule:** A proof rule which allows us to prove the equivalence of circuits with different types.
- **PANTOMIME:** An implementation of simulation-based leakage proofs in Haskell via a GHC plugin that symbolically executes GHC Core and interprets hardware data types in C λ aSH.
- **AIMCORE:** A 5-stage RISC-V processor and its leakage description, resulting in the first processor with an unconditional leakage proof against the constant-time discipline.

Availability. The full source code for PANTOMIME and AIMCORE is available in the supplementary material and will be open-sourced.

2 Overview

We illustrate our technique on a simple adder with a fast path (§ 2.1) and discuss how state (§ 2.2) and verify (§ 2.3) its leakage proof.

2.1 A Simple Adder and its Leakage Description

Circuits. We treat hardware designs as functional programs. A *circuit* is a function of type $S \rightarrow I \rightarrow (S, O)$, where S represents the circuit’s state (i.e., the current values of all its registers), I and O represent its inputs and outputs, and (S, O) represents a pair combining the circuit’s state with its output of type O —see Figure 1.

Adder. Our running example is an adder that computes the sum of two integer inputs a and b . The adder has a *fast path*: if its first input a is 0 (in which case the result is just its second input b), the adder produces an output right away in the same clock cycle. Otherwise, the adder needs some additional time to compute $a + b$ and outputs the result one cycle later. The timing difference between the paths may leak information about the inputs to an attacker. We show the adder’s code in Listing 1 using Haskell syntax. The adder’s internal state S is an optional integer of type `Maybe Int`, whose values are either `Nothing` or `Just v`, where v has type `Int`. When taking the slow path, the adder stores the intermediary result of the computation in its state. The adder’s input is a pair of optional integers and its output is a single optional integer. An input with value `Nothing` means no input is available; for example, when the client of the adder is waiting for the result of another computation.

```

1 add :: Maybe Int -> (Maybe Int, Maybe Int)
2   -> (Maybe Int, Maybe Int)
3 add _ (Just 0, Just b) = (Nothing, Just b)
4 add _ (Just a, Just b) = (Just (a+b), Nothing)
5 add s _ = (Nothing, s)

```

Listing 1: Running Example: A simple adder with a fast-path.

Computing the Output. When `add`'s first component is `Just 0`, we take the fast path (line 3): we update the state to `Nothing`—there's nothing to save—and output `Just b`. If it's non-zero, we take the slow path (line 4): we set the state to `Just (a + b)` to keep track of the sum that will be output in the next clock cycle and output `Nothing` to signal a pending result. If either of the inputs is `Nothing` (line 5), neither of the branches applies: we output the stored value, if any, and reset the state to `Nothing`.

From Circuits to Transducers. Circuits describe how hardware evolves from one clock tick to the next. `mealy`, shown below, transforms a circuit into a *transducer*—a function of type $[I] \rightarrow [O]$, which maps streams of inputs to streams of outputs. `mealy` takes a single-clock-tick circuit description and an initial state and returns the corresponding transducer: on input, it applies the circuit to update its state and produce an output, which is prepended onto the outputs produced by recursing on the remaining inputs.

```

1 mealy :: (S -> I -> (S, O)) -> S -> [I] -> [O]
2 mealy _ _ [] = []
3 mealy circuit s (i:is) =
4   let (s', o) = circuit s i in
5   o : mealy circuit s' is

```

Running the Adder. On `add`, initial state `Nothing`, and inputs

```

1 is = [(Just 1, Just 1), (Nothing, Nothing),
2       (Just 0, Just 1), (Nothing, Nothing)]

```

`mealy` produces the following output:

```

1 mealy add Nothing is
2 > [Nothing, Just 2, Just 1, Nothing]

```

Attacker View. Since we are interested in the *timing behavior* of the circuit, we model an attacker that can observe whether `add` produces an output in a given clock cycle. We model this with function `isJust`, shown below, defining our observation function as `obs = isJust`.


```

1 isJust (Just _) = True
2 isJust _       = False

```

Composing the circuit `add` with `obs` produces the combined circuit `add_obs`, which records the attacker-observable view of a computation. In particular, this circuit outputs `True` in a given clock cycle if and only if an output was produced by `add`. To construct `add_obs`, we must first define circuit composition.

Sequential Composition. Two circuits $c1 :: S1 \rightarrow I1 \rightarrow (S1, O1)$ and $c2 :: S2 \rightarrow I2 \rightarrow (S2, O2)$ may be composed to form a new circuit $c1 \circ c2$ that first executes $c1$ and then $c2$. Figure 2 illustrates circuit composition, and we define \circ below; in short, circuit composition uses circuit $c1$'s output as input to circuit $c2$, and combines their individual states.

Figure 2: For circuits $c1$, and $c2$, we write $c1 \circ c2$ for their sequential composition.

```

1 leak :: (Maybe Int, Maybe Int) -> (Maybe Bool, Bool)
2 leak (Just a, i2) = (Just (a == 0), isJust i2)
3 leak (Nothing, i2) = (Nothing, isJust i2)

```

Listing 2: Leakage function `leak` takes two optional integers, like `add`, and returns leaked values for each component.

```

1 c1 o c2 :: (S1, S2) -> I1 -> ((S1, S2), O2)
2 c1 o c2 (s1, s2) i1 =
3   let (s1', o1) = c1 s1 i1 in
4   let (s2', o2) = c2 s2 o1 in
5   ((s1', s2'), o2)

```

Lifting. We cannot directly apply circuit composition to `obs` since it's not a circuit: it doesn't have state. Instead, for any function $f :: I \rightarrow O$, we write `lift f` to turn f into a circuit by augmenting it with an empty state: `lift f :: () -> I -> ((), O)`.

Combining the Circuits. We now have the machinery to produce `add_obs`, capturing the attacker-observable view of a computation:

```

1 add_obs :: (Maybe Int, ())
2   -> (Maybe Int, Maybe Int)
3   -> ((Maybe Int, ()), Bool)
4 add_obs = add o (lift obs)

```

The `obs` observation function describes an attacker that can observe the transducer corresponding to circuit `add_obs`. `obs` models a *timing attacker*: for a given initial state and sequence of inputs, the attacker can see whether an output has been produced in each clock cycle. For example, running `add_obs`'s transducer using the initial state `(Nothing, ())` on the inputs `is` produces the following output list:

```

1 mealy add_obs (Nothing, ()) is
2 > [False, True, True, False]

```

Modeling Leakage. Next, we want to capture which information about the circuit's inputs is leaked to the attacker via an explicit leakage description. The adder leaks information about its operands via timing: if `a` is 0, the fast-path computation produces an output in the same clock cycle; otherwise the output is delayed by one cycle. By observing the presence (or absence) of an output, the attacker can therefore determine whether `a` is 0. We capture this leakage via function `leak`, shown in Listing 2. The function `leak` takes the same inputs as `add`—a pair of optional integers—and returns a pair of type $(\text{Maybe Bool}, \text{Bool})$ that describes the component-wise input leakage. The first component of the leakage says whether the first input was a proper integer `a`, in which case we also learn whether `a == 0`. The second component only says whether the second input was a proper integer, but we learn nothing about its value. Applying `lift leak` on inputs `is` using `mealy` produces the following outputs:

```

1 sim :: Bool -> (Maybe Bool, Bool) -> (Bool, Bool)
2 sim _ (Just True, True) = (False, True)
3 sim _ (Just False, True) = (True, False)
4 sim s _ = (False, s)

```

Listing 3: Simulator `sim` which replicates the attacker-observable behavior of `add` using the outputs of `leak`.

```

1 mealy (lift leak) () is
2 > [(Just False, True), (Nothing, False),
3   (Just True, True), (Nothing, False)]

```

While the leakage and observation are pure, stateless functions here, they can be arbitrary stateful computations in general.

2.2 Proving Correctness via a Simulator Circuit

Our leakage description must contain enough information to reconstruct the full observable behavior of the adder for an attacker observing timing. To prove this, we build a *simulator*—a circuit that computes the observable behavior of the adder (as defined by `obs`) solely from the leakage description `leak`. The existence of a simulator guarantees that a timing attacker can learn no information beyond what is leaked explicitly via `leak` when observing the circuit `add`. Indeed, if `leak` were missing any relevant information, the simulator could not faithfully reproduce `add`’s observable behavior.

Simulator. In our proof system, the simulator is a proof artifact produced by the user. The simulator circuit `sim` is shown in Listing 3 and closely follows circuit `add`. However, `sim`’s internal state is of type `Bool` instead of `Maybe Int`. This is sufficient; the simulator only needs to keep track of whether there is a pending computation from a previous cycle—it doesn’t need to know the value. The simulator accepts an input of type `(Maybe Bool, Bool)`, matching `leak`’s output, and produces a `Bool`, matching `add_obs`’s output.

If the leakage is `(Just True, True)` (line 2), `sim` simulates the fast path: its state is updated to `False`—there is no new pending computation—and it outputs `True` to indicate that the current cycle yielded an output. If the leakage is `(Just False, True)` (line 3), `sim` simulates the slow path: its state is set to `True` to reflect that there is a pending computation and it outputs `False` to indicate that the current cycle yielded no new output. Otherwise (line 4), it outputs the saved state—which is `True` if there’s a pending computation and `False` otherwise—and sets the new state to `False`.

Combining the Circuits. Composing `leak` with `sim` produces the combined circuit that forwards the result of `leak` to the simulator:

```

1 leak_sim :: ((), Bool)
2   -> (Maybe Int, Maybe Int)
3   -> ((), Bool, Bool)
4 leak_sim = (lift leak) ∘ sim

```

Correctness of the Leakage Description. Circuits `add_obs` and `leak_sim` both take a pair of optional integers as an input and produce a Boolean as an output. The transducers `mealy add_obs (Nothing, ())` and `mealy leak_sim ((), Nothing)` thus have the same type: `[(Maybe Int, Maybe Int)] -> [Bool]`. Indeed, if the simulator is constructed correctly, the two transducers should coincide as functions. To prove that `leak` correctly captures circuit `add`’s




Figure 3: Proof obligation for the correctness of `leak` with respect to `add`. $S = (\text{Maybe Int}, ())$, $S' = ((), \text{Bool})$, $I = (\text{Maybe Int}, \text{Maybe Int})$, and $O = \text{Bool}$. \equiv is input/output equivalence.

leakage with respect to `obs`, we need to prove that the transducers of circuits `add_obs` and circuit `leak_sim` behave the same.

Different Types. However, when we view `add_obs` and `leak_sim` at the level of circuits, they are *not* equal: indeed, we cannot even compare them as they have different types. `add_obs`’s state has type `(Maybe Int, ())`, whereas `leak_sim`’s state has type `((), Bool)`. Fortunately, equivalence at the circuit level is not our ultimate goal; we only need the underlying *transducers* to be equal.

The State Projection Rule. Based on this observation, we introduce the *state projection rule* to reason about transducer equivalence of circuits with different state types: if information can be removed from the circuit’s state—thereby changing its type—without affecting its input/output behavior, then the transducer’s behavior is also unaffected. Indeed, if we look back at the definition of `add` in Listing 1, we can see that the integer value `v` stored in the state of form `Just v` is *irrelevant* in circuit `add_obs`. Lines (3) – (4) do not depend on the state at all. Line (5) does output `v`, but we discard the result in `add_obs` due to the output projection with `obs`.

2.3 Proof Checking via State Projection

Projection Function. We define the projection function `proj` using function `isJust` from before. This function discards `add_obs`’s unit state and applies `isJust` to discard the actual computation result, only keeping the information whether there is a value or not, and finally adds `leak_sim`’s unit state.

```

1 proj :: (Maybe Int, ()) -> ((), Bool)
2 proj (s, _) = ((), isJust s)

```

Applying the State Projection Rule. The state projection rule asks us to show that applying the projection to the *output state* of `add_obs` yields the same result as applying it to the *input state* of `leak_sim`, thereby showing that `leak_sim` can compute the same result as `add_obs` *without* the information we removed. We illustrate this proof obligation in Figure 3. As the resulting functions have the same type, we can now prove their equivalence using standard methods, e.g., an SMT solver.

Summary. To sum up, we’ve shown that the simulator reconstructs the observable behavior of `add`, confirming that leakage `leak` correctly captures `add`’s leakage. To apply the state projection rule, a PANTOMIME user only has to supply a state projection function, which removes the irrelevant part of the circuit state. In our implementation, we prove function equivalence using PANTOMIME’s symbolic execution engine, described in § 5.

3 Simulation-Based Leakage Proofs

3.1 Language

We formalize our proof method in a core language (Figure 4) based on the simply-typed lambda calculus (STLC), extended with records (labeled product types) and variants (sum types). Since STLC does not allow recursion, it is a good match for describing the single-step behavior of hardware, which must be loop-free. Our type system is a version of the standard STLC as described, e.g., in [32], slightly adapted to our purposes. Importantly, our language satisfies the *preservation* property that says the type of an expression is invariant across evaluation.

The language features the usual STLC staples: constants, variables, lambda abstractions, and applications. Constants are composed of integers and a number of binary and unary operators (functions), which are used in the translation of expressions in the language to a hardware description language (see § 3.4). Labels in variant and record types are denoted by l and are assumed to be unique within the same record or variant type. We omit type annotations on function binders as they are not important in our setting. We write \vec{k}_1 to represent a vector of expressions k_1, k_2, \dots, k_n . Records are constructed with $\{\vec{l}_1 = e_i\}$, wherein each term e_i is assigned the corresponding label l_i . For a record e , $e.l$ accesses its field with label l . Term $l\ e$ constructs an inhabitant of a variant type with label l .

Syntactic Sugar. We elaborate $\text{let } p = a \text{ in } e$ according to the choice of pattern p : $\text{let } x = a \text{ in } e$ is elaborated to $(\lambda x. e)\ a$, and $\text{let } \{l_1 = p_1, \dots, l_n = p_n\} = a \text{ in } e$ elaborates to $\text{let } p_1 = a.l_1 \text{ in } \dots \text{in let } p_n = a.l_n \text{ in } e$. We write $\lambda x_1 \dots x_n. e$ for $\lambda x_1. \dots \lambda x_n. e$, and $\lambda p. e$ for $\lambda x. \text{let } p = x \text{ in } e$. If e is a term of a record type $\{\vec{l}_1 : s_i\}$, we write $e\{l_i := e_i\}$ for the single-field update of the record e to e_i at field l_i , and $e\{l_{i_1} := e_{i_1}\} \dots \{l_{i_k} := e_{i_k}\}$ for the sequence of updates $e\{l_{i_1} := e_{i_1}\} \dots \{l_{i_k} := e_{i_k}\}$. Finally, infix applications aren't supported: when unambiguous, $e_1 \text{ bop } e_2$ should be understood as syntactic sugar for $\text{bop } e_1\ e_2$.

Evaluation. Figure 5 shows the evaluation rules for the core language. Evaluation is based on the reduction relation \rightsquigarrow , which is defined on closed, well-typed terms of the same type. The substitution operator $[x = e]$ performs a capture-avoiding substitution of free occurrences of the variable x with term e . Beta reduction and the step rules for evaluating applications, records, variants, and case expressions are standard. Rule [Record Reduce] reduces a record to the expression associated with the applied label. Rule [Case Reduce] evaluates a case expression $\text{case } l_k\ e \text{ of } \vec{l}_i\ x_i \rightarrow e_i$ over a variant by matching $l_k\ e$ with case $l_k\ x_k$, returning $e_k[x_k = e]$. Since the STLC is both confluent and strongly normalizing, the order in which we apply the evaluation rules does not matter and evaluation always terminates. We write $\llbracket e \rrbracket$ to represent the unique term obtained by exhaustive application of the evaluation steps to e . We write \rightsquigarrow_* for the reflexive-transitive closure of \rightsquigarrow .

Unit, Optional, and Tuple Types. We encode the unit type $()$ as the empty record type with no fields. For any type t , we encode Maybe t as the variant type $\langle \text{Just} : t, \text{Nothing} : () \rangle$. Conditionals $\text{if } e \text{ then } e_1 \text{ else } e_2$ stand for the corresponding case expressions over the variant type $\text{Bool} := \langle \text{true} : (), \text{false} : () \rangle$. We encode the tuple type (t_1, \dots, t_n) as the the record type $\{1 : t_1, \dots, n : t_n\}$.

Base Types	$b ::= \text{Int, Word8, Word16, } \dots$	
Operators	$\text{bop} ::= +, -, \dots$	<i>binary operators</i>
	$\text{uop} ::= \sim, !, \dots$	<i>unary operators</i>
Constants	$c ::= 0, 1, \dots$	<i>numbers</i>
	$\text{bop} \mid \text{uop}$	<i>operators</i>
Types	$t ::= b$	<i>base types</i>
	$\mid \{\vec{l}_1 : t_i\}$	<i>record types</i>
	$\mid \langle \vec{l}_1 : t_i \rangle$	<i>variant types</i>
	$\mid t_1 \rightarrow t_2$	<i>function types</i>
Terms	$e ::= c$	<i>constants</i>
	$\mid x$	<i>variables</i>
	$\mid \lambda x. e$	<i>abstraction</i>
	$\mid e_1\ e_2$	<i>application</i>
	$\mid \{\vec{l}_1 = e_i\}$	<i>records</i>
	$\mid e.l$	<i>record fields</i>
	$\mid l\ e$	<i>variants</i>
	$\mid \text{case } e \text{ of } \vec{l}_i\ x_i \rightarrow e_i$	<i>pattern match</i>
Extractable	$s ::= b$	<i>base types</i>
Types	$\mid \{\vec{l}_1 : s_i\}$	<i>record types</i>
	$\mid \langle \vec{l}_1 : s_i \rangle$	<i>variant types</i>
Patterns	$p ::= x$	<i>variables</i>
	$\mid \{\vec{l}_1 = p_i\}$	<i>records</i>

Figure 4: Syntax of types and terms in our core language.

$\frac{e \rightsquigarrow e'}{\lambda x. e \rightsquigarrow \lambda x. e'} [\lambda]$
$\frac{e_1 \rightsquigarrow e'_1 \quad e_2 \rightsquigarrow e'_2}{e_1\ e_2 \rightsquigarrow e'_1\ e'_2} [\text{App}]$
$\frac{}{(\lambda x. e_1)\ e_2 \rightsquigarrow e_1[x = e_2]} [\text{Beta Reduce}]$
$\frac{e_i \rightsquigarrow e'_i}{l_i = e_i \rightsquigarrow l_i = e'_i} [\text{Record Field}]$
$\frac{e \rightsquigarrow e'}{e.l \rightsquigarrow e'.l} [\text{Record Label}]$
$\frac{e \rightsquigarrow e'}{l\ e \rightsquigarrow l\ e'} [\text{Variant}]$
$\frac{e \rightsquigarrow e'}{\text{case } e \text{ of } \vec{l}_i\ x_i \rightarrow e_i \rightsquigarrow \text{case } e' \text{ of } \vec{l}_i\ x_i \rightarrow e_i} [\text{Case}]$
$\frac{}{\{\vec{l}_1 = e_i\}.l_j \rightsquigarrow e_j} [\text{Record Reduce}]$
$\frac{}{\text{case } l_k\ e \text{ of } \vec{l}_i\ x_i \rightarrow e'_i \rightsquigarrow e'_k[x_k = e]} [\text{Case Reduce}]$

Figure 5: Evaluation rules for our core language.

Example 3.1. Consider the term

$\text{ex1} : \text{Int} \rightarrow \text{Maybe Word32} \rightarrow (\text{Word32}, \text{Word32})$
 $\text{ex1} = \lambda s \ i. (s + 1, \text{case } i \text{ of Just } a \rightarrow s + a, \text{Nothing} \rightarrow s).$

Then we have

$$\begin{aligned} \llbracket \text{ex1 } \emptyset \text{ Nothing} \rrbracket &= (1, \emptyset), \\ \llbracket \text{ex1 } \emptyset (\text{Just } 1) \rrbracket &= (1, 1). \end{aligned}$$

3.2 Circuits, Transducers, and Circuit Reduction

Definition 3.2 (circuit). Let S, I, O be types in the core language. A circuit is a term of type $S \rightarrow I \rightarrow (S, O)$.

For circuits to be hardware-extractable, S, I, O must be extractable types (*i.e.*, they can't contain function types). In practice, this means we can use higher-order functions freely whenever we don't need to synthesize the circuit into hardware—for example, for the leakage, observation or simulator circuits.

Example 3.3. ex1 is an extractable circuit that increments its state each clock tick. It outputs the sum of its input and current state when the input is non-empty, and outputs its state, otherwise.

Transducers. Circuits describe the behavior of a hardware design in a single clock tick. The *Mealy* function M transforms a circuit into a *transducer* that applies one input per clock tick (from a list of inputs) to the circuit and produces a corresponding list of outputs. Transducers therefore represent the actual input/output behavior of sequential hardware.

Definition 3.4 (Mealy). Let $c : S \rightarrow I \rightarrow (S, O)$ be a circuit. The Mealy function M is defined as follows, where $(s', o) = \llbracket c \ s \ i \rrbracket$, $:$ is list concatenation, and $[]$ is the empty list.

$$\begin{aligned} M : (S \rightarrow I \rightarrow (S, O), S, [I]) &\rightarrow [O] \\ M(c, s, is) &= \begin{cases} [] & \text{if } is = [], \\ o : M(c, s', is') & \text{if } is = i : is' \end{cases} \end{aligned}$$

M takes a circuit c and initial state s along with a list of inputs is . It then applies c to the first input of is to compute a new state s' and output o , which it prepends to the list of outputs and then recursively computes the remaining outputs using the new state s' and remaining inputs is' .

Example 3.5. Revisiting circuit ex1 , let $s = \emptyset$ be an initial state and $is = [\text{Just } 1, \text{Nothing}, \text{Just } 2, \text{Just } 3]$ a list of inputs. Applying the Mealy function to ex1 with these inputs yields

$$\begin{aligned} M(\text{ex1}, s, is) &= M(\text{ex1}, \emptyset, [\text{Just } 1, \text{Nothing}, \text{Just } 2, \text{Just } 3]) \\ &= 1 : M(\text{ex1}, 1, [\text{Nothing}, \text{Just } 2, \text{Just } 3]) \\ &= [1, 1, 4, 6]. \end{aligned}$$

Observational Equivalence. As discussed in §2, our proof method requires proving equivalence of functions. In particular, we are interested in *observable* function equivalence (rather than syntactic equivalence). For example, the functions $\lambda i. i + \emptyset$ and $\lambda i. \emptyset + i$ should be equivalent.

Definition 3.6 (Observational Equivalence). Two closed terms e, e' of the same type t are *observationally equivalent*, written $e \equiv e'$ when one of the following holds:

- (1) t is a base type and $\llbracket e \rrbracket = \llbracket e' \rrbracket$.

- (2) $t = t_1 \rightarrow t_2$ and if for any two closed terms $e_1, e'_1 : t_1$ with $e_1 \equiv e'_1$, then $e \ e_1 \equiv e' \ e'_1$.
- (3) $t = \{\overrightarrow{1_i : t_i}\}$, $e \rightsquigarrow_* \{\overrightarrow{1_i = e_i}\}$, $e' \rightsquigarrow_* \{\overrightarrow{1_i = e'_i}\}$, and for all i , $e_i \equiv e'_i$.
- (4) $t = \langle \overrightarrow{1_i : t_i} \rangle$, $e \rightsquigarrow_* 1_k \ e_1$, $e' \rightsquigarrow_* 1_k \ e'_1$, and $e_1 \equiv e'_1$.

We also lift observational equivalence to lists of terms: two lists of equal length are observationally equivalent if their elements are component-wise observationally equivalent.

Example 3.7. The three functions $\lambda i. i + \emptyset$, $\lambda i. \emptyset + i$, and $\lambda i. i$ are observationally equivalent because they return the same result on any input. Similarly, $\lambda b. b \equiv \lambda b. \text{if } b \text{ then true else false}$ as each function returns the same result on both true and false.

Definition 3.8 (State Projection Rule). Given circuits $c_1 : S_1 \rightarrow I \rightarrow (S_1, O)$ and $c_2 : S_2 \rightarrow I \rightarrow (S_2, O)$, and a function $p : S_1 \rightarrow S_2$, define

$$\begin{aligned} c_1 p &: S_1 \rightarrow I \rightarrow (S_2, O) \\ c_1 p &= \lambda s_1 \ i. \text{let } (s'_1, o) = c_1 \ s_1 \ i \text{ in } (p \ s'_1, o), \\ p c_2 &: S_1 \rightarrow I \rightarrow (S_2, O) \\ p c_2 &= \lambda s_1 \ i. c_2 \ (p \ s_1) \ i. \end{aligned}$$

We say that c_1 *reduces* to c_2 , written $c_1 \hookrightarrow c_2$, if $c_1 p \equiv p c_2$, and call p the *state projection function* witnessing this reduction.

Example 3.9. In the adder example from §2, we showed that $\text{add_obs} \hookrightarrow \text{leak_sim}$, *i.e.*, that add_obs reduces to leak_sim using state projection function proj , by the observational equivalence between $\text{c1p } s \ i = \text{let } (s', o) = \text{add_obs } s \ i \text{ in } (\text{proj } s', o)$, and $\text{pc2 } s \ i = \text{leak_sim } (\text{proj } s) \ i$. We wanted to show that the transducers of add_obs and leak_sim are equal, which confirms that lift leak is a suitable leakage for add under obs . In fact, circuit reduction implies transducer equivalence:

THEOREM 3.10 (CIRCUIT REDUCTION IMPLIES MEALY EQUIVALENCE). *Given two circuits $c_1 : S_1 \rightarrow I \rightarrow (S_1, O)$ and $c_2 : S_2 \rightarrow I \rightarrow (S_2, O)$, if $c_1 \hookrightarrow c_2$ with state projection function $p : S_1 \rightarrow S_2$, then for all initial states $s_1 : S_1$ there exists an initial state $s_2 : S_2$ such that for all input sequences is of type $[I]$, we have $M(c_1, s_1, is) \equiv M(c_2, s_2, is)$.*

We prove Theorem 3.10 in §A.

3.3 Leakage Description

Sequential Composition. To define leakage descriptions, we first define sequential circuit composition.

Definition 3.11 (Sequential Composition). Let $c_1 : S_1 \rightarrow I_1 \rightarrow (S_1, O_1)$ and $c_2 : S_2 \rightarrow O_1 \rightarrow (S_2, O_2)$ be circuits. Their *sequential composition*, denoted $c_1 \circ c_2$, is defined as

$$\begin{aligned} c_1 \circ c_2 &: (S_1, S_2) \rightarrow I_1 \rightarrow ((S_1, S_2), O_2) \\ c_1 \circ c_2 &\triangleq \lambda (s_1, s_2) \ i_1. \end{aligned}$$

$$\begin{aligned} &\text{let } (s'_1, o_1) = c_1 \ s_1 \ i_1 \text{ in} \\ &\text{let } (s'_2, o_2) = c_2 \ s_2 \ o_1 \text{ in } ((s'_1, s'_2), o_2). \end{aligned}$$

For transducers, sequential composition of circuits becomes function composition (which follows by induction on the list of inputs):

Basic Types

r	\in	Regs	registers
v	\in	Vars	variables
i	\in	$\text{Regs} \cup \text{Vars}$	identifiers
n	\in	$\{0, 1\}^*$	values

Syntax

e	$:=$	$n \mid i \mid \ominus e \mid e_1 \otimes e_2 \mid \{e_1, e_2\}$	expressions
		$\mid \text{if } e_1 \text{ th } e_2 \text{ el } e_3 \mid e_1[e_2 : e_3]$	
w	$:=$	$v = e$	wires
W	$:=$	$\{w_1, w_2, \dots\}$	wire sets
a	$:=$	$r \leftarrow v$	assignments
A	$:=$	$\{a_1, a_2, \dots\}$	assignments sets
I	$:=$	$\{v_1, v_2, \dots\}$	inputs
O	$:=$	$\{v_1, v_2, \dots\}$	outputs
C	$:=$	$I : W : A : O$	circuits

Figure 6: μ -Verilog syntax.

LEMMA 3.12. *Given two circuits $c_1 : S_1 \rightarrow I_1 \rightarrow (S_1, O_1)$ and $c_2 : S_2 \rightarrow O_1 \rightarrow (S_2, O_2)$ be circuits, states $s_1 : S_1$ and $s_2 : S_2$, and an input sequence is : $[I]$, we have*

$$M(c_1 \circ c_2, (s_1, s_2), is) = M(c_2, s_2, M(c_1, s_1, is)).$$

Leakage. We now have the machinery needed to define leakage:

Definition 3.13 (Leakage). Let $c : S \rightarrow I \rightarrow (S, O)$ be a circuit and $o : S_o \rightarrow O \rightarrow (S_o, O_o)$ an observation. A circuit $l : S_l \rightarrow I \rightarrow (S_l, O_l)$ is a *leakage description* for c under o if there exists a circuit $s : S_s \rightarrow O_l \rightarrow (S_s, O_o)$ such that $c \circ o \hookrightarrow l \circ s$.

Example 3.14. Revisiting the adder (§ 2), circuit `lift leak` is a leakage description for `add` under `lift obs` since there is a simulator `sim` such that `add` \circ (`lift obs`) \hookrightarrow (`lift leak`) \circ `sim`.

Next, we adapt contract equivalence [41]—a prior notion of leakage—to our setting and show it’s implied by our definition.

Definition 3.15 (Contract Equivalence). Let $c : S \rightarrow I \rightarrow (S, O)$ be a circuit and $o : S_o \rightarrow O \rightarrow (S_o, O_o)$ its observation. A circuit $l : S_l \rightarrow I \rightarrow (S_l, O_l)$ is called a *contract* for c with respect to the observation o if, for all states $s_l : S_l$ and all input sequences $is, is' : [I]$, whenever $M(l, s_l, is) = M(l, s_l, is')$, then for all states $s : S$ and $s_o : S_o$, we have $M(c \circ o, (s, s_o), is) = M(c \circ o, (s, s_o), is')$.

THEOREM 3.16 (LEAKAGE IMPLIES CONTRACT EQUIVALENCE). *Let $c : S \rightarrow I \rightarrow (S, O)$ be a circuit and $o : S_o \rightarrow O \rightarrow (S_o, O_o)$ be its observation. If circuit $l : S_l \rightarrow I \rightarrow (S_l, O_l)$ is a leakage for c with respect to observation o , then l is also a contract.*

Theorem 3.16 is proved in § A.

3.4 From Circuits to Hardware

If its types are extractable, we can compile a circuit into a hardware design that matches the behavior of its associated transducer. We show this process with μ -Verilog, a simplified version of Verilog.¹

μ -Verilog Syntax. Figure 6 shows μ -Verilog’s syntax, based on [41]. μ -Verilog circuits consist of a set of input variables I , a set of assignments A , and a set of output wires O . Input variables are assigned an

external input value at the start of each clock cycle. A wire assignment of the form $v = e$ immediately updates the value of variable v . An assignment of the form $r \leftarrow v$ updates the value of register r in the next clock cycle. We assume that assignments are disjoint—no two assignments update the same register or variable. Expressions are formed from bitvectors, register identifiers, unary and binary operations ($\ominus e$ and $e_1 \otimes e_2$), conditionals, the bit-selection operator $e_1[e_2 : e_3]$ (which selects a subset of bits from a value), and the wire concatenation operator $\{e_1, e_2\}$ (which forms a new wire consisting of both e_1 and e_2). Output wires specify the circuit’s outputs.

μ -Verilog Semantics. § B describes the semantics. In short, wire assignments take immediate effect, while register assignments occur concurrently, synchronized by the clock. Circuits first assign inputs to input wires, evaluate all wire expressions, and compute the outputs. Finally, they update registers for the next clock cycle. The circuit semantics $\llbracket C \rrbracket(\mu, i) = (\mu', o)$ captures this behavior: given a valuation μ that maps registers to values and a tuple of inputs i , it computes an output valuation μ' and a tuple of outputs o . Variable values don’t need to be recorded, as they aren’t preserved across cycles. The transducer semantics $M(C, \mu, is)$ is similarly defined.

Example 3.17. Consider the following μ -Verilog circuit

$$C_{\text{ex2}} = \{i\} : W : \{s \leftarrow s'\} : \{o\},$$

where $W = \{s' = s + 1, \quad o = \text{if } i[32 : 32] \text{ th } s + i[32 : 1] \text{ el } s\}$, $\text{size}(i) = 33$, and $\text{size}(s) = \text{size}(o) = 32$. C_{ex2} increments register s in each clock cycle. It then assigns to output register o : if i ’s most significant bit is 1, o is assigned $s + i[32 : 1]$ (i.e., the sum of s and i ’s remaining 32 bits), otherwise it’s assigned s .

From Circuits to μ -Verilog. This section sketches the translation from the core language (Figure 4) to Verilog; full translations rules appear in § C. A term $\lambda s \text{ i. } e$ of type $S \rightarrow I \rightarrow (S, O)$ (where types S, I , and O are extractable) is translated as follows, where v_i is an input variable, v_o an output variable, s is a state register, and v_e is the result of translating e .

$$\{v_i\} : \{v_o = \text{field}(v_e, 2), s' = \text{field}(v_e, 1), \dots\} : \{s \leftarrow s'\} : \{v_o\}$$

The field extraction function $\text{field}(v, l)$ accesses the field l of a record encoded into variable v by appropriately slicing the variable. For example, for the record type $\{a : \text{Word8}, b : \text{Word4}\}$, $\text{field}(v, a) = v[0 : 8]$.

The translation creates a variable $v_{e'}$ for each sub-term e' of e . It encodes terms of record types by creating a variable with enough bits to fit all components of the record and using the bit-slice operator to extract individual component (via field). To encode terms of variant types, the translation creates a variable with enough bits to encode all possible data-types for the variant, as well as a label indicating which choice the term represents. It encodes case expressions of the form $\text{case } e \text{ of } \overrightarrow{l_i x_i} \rightarrow e_i$ into conditional expressions, using bit-slicing to assign the relevant parts of expression e to variable x_i for each case.

Example 3.18. Applying the translation to the circuit `ex1` from Example 3.1 yields the μ -Verilog circuit C_{ex2} from Example 3.17. The optional type `Maybe Word32` is encoded into a 33-bit wire by conjoining a single bit to encode the choice between labels `Nothing` and `Just`, with 32 bits encoding the value of type `Word32`.

¹Verilog is one of the main hardware description languages; the other, VHDL, is similar.

```

1 fe :: State -> (Word16, Maybe Word8) -> (State, ())
2 fe state (rawInstr, jmp) =
3   let curPC = pc state in
4   let instr = decode rawInstr in
5   case jmp of
6     Just newPC -> (state {exInstr = Add 0,
7                          exPC = curPC,
8                          pc = newPC}, ())
9     Nothing -> (state {exInstr = instr,
10                      exPC = curPC,
11                      pc = curPC + 1}, ())

```

Listing 4: Fetch Stage.

4 Case Study: Pipelined Processor

This section shows how to construct a simulation based-leakage proof for a simple three stage pipelined processor with a single register. While we use Haskell syntax for readability, the processor is fully expressible in our core language.

4.1 The Processor

Instruction Set. We encode instructions as the variant type below:

```

1 data Instr = Add Word8 | Clr
2             | Out | Jmp Word8 | Beq Word8

```

Add adds a `Word8` value to the register, **Clr** resets it 0, **Out** outputs its current value, **Jmp** jumps to a `Word8` address, and **Beq** branches if the register is 0, adding its `Word8` offset to the current program counter.

Pipeline Stages. Each instruction passes through three processor stages. *Fetch* decodes a raw `Word16` instruction into an `Instr`. *Execute* runs the instruction, producing a register update (the *writeback*), an optional *output* value, and the next program counter. *Writeback* applies the register update and returns any outputs.

Toplevel Circuit and State. We encode the processor as a circuit `proc` (Listing 7) that takes a raw `Word16` instruction, and returns a pair: an optional `Maybe Word32` output and the new `Word8` program counter. The processor's state is encoded with the record type

```

1 data State = State
2   {pc :: Word8, reg :: Word32,
3    exInstr :: Instr, exPC :: Word8,
4    wbOut :: (Maybe Word32, Maybe Word32)}

```

Here, `pc` is the current program counter and `reg` holds the register value. The other fields will be explained as we go.

Fetch Stage. The fetch stage (Listing 4) decodes the raw instruction and updates the program counter based on the previous instruction's execution. It stores the decoded instruction in the state's `exInstr` field and the current program counter (needed when executing a `Beq` instruction) in `exPC`; these become state inputs to execute stage in the next cycle. If a jump occurred in the previous cycle (i.e., `jmp` is a `Just`-value), the instruction is discarded and is replaced with a no-op (`Add 0`). Otherwise, the program counter is incremented by 1.

Execute Stage. The execute stage (Listing 5) executes the instruction fetched last cycle, storing the result—a register update and the computed output—as `wbOut`. In the next cycle, the writeback

```

1 ex :: State -> () -> (State, Maybe Word8)
2 ex state _ =
3   let curReg = reg state in
4   case exInstr state of
5     Add imm -> let newReg = curReg + imm in
6       (state {wbOut = (Just newReg, Nothing)},
7        Nothing)
9     Beq off ->
10      if curReg == 0
11      then let newPC = (exPC state) + off in
12        (state {wbOut = (Nothing, Nothing)},
13         Just newPC)
13      else (state {wbOut = (Nothing, Nothing)},
14            Nothing)
15      ...

```

Listing 5: Execute Stage.

```

1 wb :: State -> () -> (State, Maybe Word32)
2 wb state _ =
3   let (wb, out) = wbOut state in
4   case wb of
5     Just newReg -> (state {reg = newReg}, out)
6     Nothing -> (state, out)

```

Listing 6: Writeback Stage.

```

1 proc :: State -> Word16
2       -> (State, (Maybe Word32, Word8))
3 proc state rawInstr =
4   let (state', out) = wb state () in
5   let (state'', jmp) = ex state' () in
6   let (state''', _) = fe state'' (rawInstr, jmp) in
7   (state''', (out, pc state'''))

```

Listing 7: The Pipelined Processor.

stage will update the register and perform the processor's output according to the stored `wbOut` values. For an `Add imm` instruction, we add `imm` to the register value and store the update in `wbOut`. For `Out`, we store the register value in `wbOut`, to be output by the writeback stage. For `Beq off`, if the current register value is 0, we jump by outputting `Just newPC` on a wire connected to the fetch stage.

Writeback Stage. The writeback stage (Listing 6) updates the register value in the processor state and returns the output.

Pipelining. In one clock cycle, `proc` executes all three pipeline stages concurrently, processing three instructions in parallel: it *fetches* the instruction at the current program counter, *executes* the instruction fetched last cycle, and performs the *writeback* for the instruction fetched two cycles ago. The pipeline executes stages in reverse order—writeback, then execute, then fetch—to prevent overwriting stage state inputs before they are used.


```

1 leak :: LState -> Word16 -> (LState, LInstr)
2 leak state rawInstr =
3   let curReg = reg state in
4   let instr = decode rawInstr in
5   case exInstr state of
6     Add imm -> let newReg = curReg + imm in
7       (state {reg = newReg,
8              exInstr = instr}, LOther)
9   Beq off ->
10    if curReg == 0
11    then (state {exInstr = Add 0}, LBeq off)
12    else (state {exInstr = instr}, LOther)
13 ...

```

Listing 8: Leakage.

4.2 Observation, Leakage, and Simulator

Observation. The observation function discards the processor output, only keeping the program counter:

```

1 obs :: (Maybe Word32, Word8) -> Word8
2 obs (_, pc) = pc

```

This observation encodes an attacker that can observe the control flow of the program and, therefore, its timing.

Leakage. Since the attacker only sees the program counter, we only need to simulate the control flow. To do so, for `Jmp` we need to know the target address, and for `Beq` the offset. For all other instructions, the program counter increases by 1. This suggests the following leakage (i.e., input for the simulator):

```

1 data LInstr = LJump Word8 | LBeq Word8 | LOther

```

The leakage circuit (Listing 8) must produce a leakage of type `LInstr` from raw `Word16` instructions. To determine when to branch, `leak` must track the register value. Like the processor, it decodes and stores instructions for subsequent execution—this suggests the following state for `leak`:

```

1 data LState =
2   LState {reg :: Word32, exInstr :: Instr}

```

For `Add imm`, we add `imm` to the register value, update the register, and decode the next instruction, storing it in `exInstr` for the next cycle. For `Beq off`, if the register value is 0, we jump: we discard the raw instruction, schedule a no-op (`Add 0`) for the next cycle, and—to allow the simulator to compute the address of the next instruction to fetch—record that a jump has occurred by leaking the offset as `LBeq off`.

Simulator. The simulator circuit must take `LInstrs` as inputs and produce the `Word8` program counter. The simulator state is the remainder of the original processor state (except for `wbOut`):

```

1 data SState = SState {pc :: Word8, exPC :: Word8}

```

Listing 9 shows the code. On `LJump addr` instructions, the next program counter is the specified address. For `LBeq off`, it's the original instruction address `exPC` plus the offset. Finally, on `LOther`, it's just the current counter incremented by 1.

State Projection. To show that the simulator matches the leakage, we need to prove that the circuit `proc` \circ (`lift obs`) reduces to

```

1 sim :: SState -> LInstr -> (SState, Word8)
2 sim state leakInstr =
3   let curPC = pc state in
4   case leakInstr of
5     LNonJump ->
6       (state {exPC = curPC, pc = curPC + 1},
7        curPC + 1)
8     LJump addr ->
9       (state {exPC = curPC, pc = addr}, addr)
10    LBeq off ->
11      let newPC = (exPC state) + off in
12      (state {exPC = curPC, pc = newPC}, newPC)

```

Listing 9: Simulator.

```

1 proj :: State -> (LState, SState)
2 proj (State pc reg exInstr exPC wbOut) =
3   let (wb, _) = wbOut in
4   case wb of
5     Just newReg ->
6       (LState newReg exInstr, SState pc exPC)
7     Nothing ->
8       (LState reg exInstr, SState pc exPC)

```

Listing 10: State Projection.

the circuit `leak` \circ `sim`. To do so, we construct the state projection function shown in Listing 10. The state projection function divides the processor state into leakage and simulator states. The exception is `wbOut`, which lacks an analogue in the leakage and simulator states. Instead, we discard the output part of `wbOut` and perform an in-flight register update according to the writeback part of `wbOut`.

5 Implementation

PANTOMIME. PANTOMIME is a full-path symbolic execution engine for GHC Core—the internal language of GHC (the Glasgow Haskell Compiler), which is based on System F_C.² Implemented as a GHC Core plugin, PANTOMIME verifies that a circuit's leakage specification is correct. To use it, the user provides the leakage specification as an annotation, like the one shown below.

```

1 {-# ANN circuit Pantomime
2   { observation = 'obs, leakage    = 'leak,
3     simulator   = 'sim, projection = 'proj } #-}
4 circuit :: s -> i -> (s, o)
5 circuit = ...

```

Given a specification, PANTOMIME generates constraints that entail its correctness. After constraint generation, PANTOMIME queries an SMT solver to check validity of the constraints. In case of a violation, PANTOMIME returns a pair of diverging observations along with their corresponding inputs. As PANTOMIME symbolically executes GHC Core, it supports all of System F_C. This allows hardware designs in PANTOMIME (and their corresponding leakage descriptions and simulators) to use the full range of functional programming techniques supported by Haskell—including monads, typeclasses

²System F_C [35] extends System F with non-syntactic type equality support.

and GADTs—provided they can be synthesized to hardware. PANTOMIME is written in approximately 7600 lines of Haskell code.

Grisette. We implemented PANTOMIME using the Grisette [26] symbolic execution library. Grisette provides an *Ordered Guards* representation, which allows for efficient merging of symbolic values. Additionally, Grisette folds constants and deduplicates expressions within terms where possible, further reducing constraint sizes.

In Grisette, a value of type `Maybe Word` is represented as

```
1 SymMaybeWord { tag = "tag" :: SymInt64
2               , word = "word" :: SymWord64 }
```

where `"tag"` and `"word"` are *symbolic constants*.³ Since `Maybe Word64` has only two variants (`Just` or `Nothing`), `SymMaybeWord`'s `tag` field needs to be restricted to be within range; Grisette allows us to encode this assumption directly in the expression:⁴

```
1 tag = If !(0 <= "tag" && "tag" < 2)
2       (Left Invalid) (Right "tag")
```

Any value outside of the range of variants is `Invalid` and should not be considered by the constraint solver. More generally, a symbolic expression is represented by a chain of if-then-else expressions.

Merging Expressions. Instead of exposing the primitive constructor `If` directly, Grisette wraps symbolic branching within the `mrGite` function. Directly using `If` hides opportunities for optimizations due to possible `Invalid` branches. We illustrate this in the following example, where we branch on conditional `condB` with both branches `x` and `y` containing `Invalid` in their body.

```
1 x = If condX (Left Invalid) (Right "x")
2 y = If condY (Left Invalid) (Right "y")
3
4 mrGite condB x y
5 > If (condX || condY)
6     (Left Invalid)
7     (Right (If condB "x" "y"))
```

Notice that merging opens up the possibility to simplify `condX || condY`. In comparison, this reduction is non-obvious when using the naive branch operation `If x y`.

Constraint Generation. The constraint generation algorithm parallels concrete evaluation in GHC Core. In fact, PANTOMIME doubles as an interpreter when given constant inputs as Grisette will perform constant folding on expressions. The symbolic evaluator takes an environment and a Core expression as inputs and converts it into a symbolic expression. For example, a Core variable expression triggers a lookup in the environment, and a Core lambda converts into a function that receives a symbolic argument.

PANTOMIME follows Haskell's evaluation semantics. It supports type and coercion reductions and tracks typecasts to ensure correct and well-typed reductions. Abstract data types are implemented as records containing a tag—that symbolically represents the data constructor—along with symbolic versions of the fields for every possible data constructor. PANTOMIME supports abstract data types generically, such that it can accept any new user definitions as is.

The symbolic valuation produced by PANTOMIME follows lazy semantics—expression constraints appear in the resulting expression only if the expression is used. This keeps constraint sizes small by including only what is strictly necessary.

ClaSH Support. We implemented AIMCORE, in ClaSH [7], a hardware description language embedded in Haskell. ClaSH exposes hardware primitives, such as the sized bit vector type `BitVec n` and operations on them (e.g., bit vector concatenation), which it can compile to synthesizable HDLs, like Verilog. In PANTOMIME, we implemented interpretations for these hardware primitives and operations, enabling verification of specifications for ClaSH circuits. For example, we encode `BitVec n` as the appropriately sized bit vector in the corresponding first-order theory. This requires PANTOMIME to perform type-level reductions to statically determine vector bit-sizes. We take a similar approach for primitive operations.

Function Equivalence. Given an annotation, PANTOMIME constructs two circuits whose equivalence determines the validity of the leakage specification, illustrated below.

```
1 real s i =
2   let (s', o) = circuit o obs $ s i in
3   (projection s', o)
4
5 synth s i = leakage o simulator $ (projection s) i
```

The first circuit, `real`, composes `circuit` with `observation`, and then applies `projection` to mask the output state. The second circuit, `synth` masks the input state with `projection` before passing it to `leakage` composed with `simulator`. The leakage specification is valid iff `real s i == synth s i` for all `s, i`—i.e., there is no pair of input state for which the circuits diverge.

6 Evaluation

We evaluate PANTOMIME by asking the following research questions.

RQ1: Can we use PANTOMIME to write/verify a RISC-V processor?

RQ2: What's the proof effort of verification with PANTOMIME?

RQ3: How does AIMCORE compare to other verified processors?

RQ4: How do PANTOMIME's leakage descriptions and guarantees compare to other methods?

RQ5: How long does PANTOMIME take to verify correctness?

RQ1: Implementing AIMCORE and its Proof. We answer RQ1 in the affirmative by reporting on implementing and verifying AIMCORE using PANTOMIME. AIMCORE is a 5-stage in-order pipelined processor that implements the RISC-V V2.1 RV32I Base Integer Instruction Set [42].⁵ AIMCORE consists of around 800 lines of Haskell and extracts to around 2400 lines of Verilog. Its design builds on the simpler processor in § 4, and consists of five pipeline stages. The core receives values from memory and the register file as inputs and produces memory and register file accesses as outputs. Programs are stored alongside values in memory. The core forwards values from the `wb` and `mem` stages to the `exe` stage (see `proc` below); pipeline stalls arise only from memory operations (load dependencies or a memory access blocking instruction fetch) and branches.

Processor in Haskell. Since the core is written in Haskell, it can make use of expressive monad abstractions and elide explicit passing of inputs, outputs, and state; the complete processor pipeline

³"tag" and "word" are *not* strings and are indeed of type `SymInt64` and `SymWord64`; the `OverloadedStrings` GHC extension enables us to represent these symbolic expressions as strings.

⁴`If` constructs *symbolic* if-then-else expressions.

⁵Excluding the `ecall` and `ebreak` instructions.

is specified simply as a sequence of its individual stages, where `runCPUM` extracts the function payload from the `CPUM` monad⁶ that the pipeline is written in.

```
1 proc :: State -> Input -> (State, Output)
2 proc = runCPUM (wb >> mem >> exe >> de >> fetch)
```

Since `proc` consists only of combinatorial logic, it can be directly compiled into an HDL using a tool like `ClaSH` and then synthesized.

Observation. We model an attacker that can observe the program counter, revealing control flow and timing. In a given cycle the processor may access a memory value instead of fetching an instruction; our attacker observation model reflects this by exposing the program counter only during instruction fetches:

```
1 obs :: Output -> Maybe Address
```

Leakage. As in § 4, we write a leakage circuit that takes as input the processor’s input (a memory read and two register file reads) and constructs the leakage:

```
1 leak :: LState -> Input -> (LState, Leak)
```

Since the register file and memory are external to the core, the core’s pipeline structure affects when inputs are requested: for instance, `de` must request register operands one cycle before `exe` needs them. Similarly, the result of a load request issued by `mem` appears in `wb` in the following cycle. Since jump and branch addresses (and, by extension, the program counter) can result from arbitrary computation, `leak` must faithfully replicate the processor’s architectural state, which requires matching the timing of requests to the register file and memory. Rather than distilling the core’s timing behavior into separate logic, we capture it implicitly by structurally matching `leak` to `proc`—like in our case-study. `leak` therefore has a 5-stage pipeline structure [11].

ISA Interpreter. To compute addresses, `leak` calls a black-box ISA interpreter that implements the ISA spec. As such, `leak` must only explicitly replicate the timing behavior of `proc`—all actual computation is delegated to the black-box interpreter, which is reusable across different leakage models. As a result, `leak` is easy to write for hardware designers: it looks just like `proc`, but with all computation abstracted away—retaining only the core’s input timing, stalling, and value-forwarding behavior.

Leakage Datatype. The actual leakage that `leak` computes is a tuple defined as:

```
1 type Leak = (LInstr, (Maybe RegId, Maybe RegId)
2             , Maybe Address)
```

It consists of three components: (1) the leakage instruction `LInstr`:

```
1 data LInstr = LJump | LLoad RegId | LStore | LOther
2             (2) (Maybe RegId, Maybe RegId), a pair of optional registers that
the core instruction corresponding to the given leakage instruction
may depend on, and (3) Maybe Address, the address of a jump
or branch. LJump has no Address payload because leak does not
yet know the jump target when it needs to issue the leakage
instruction. Instead, we leak jump addresses later, when they
become available, using the third component of Leak. This approach
also eliminates the need for a separate instruction for conditional
branches: leak outputs jump target Nothing for conditional
```

branches that end up not being taken. The simulator uses an instruction’s register dependencies to stall when the core stalls (e.g., for load dependencies). Instructions `LLoad` and `LStore` are needed to determine stalls. `LLoad` includes its destination register to determine load dependencies, and `LStore` has no payload. All other instructions act as a no-op, represented with `LOther`.

Simulator. The simulator `sim` takes as input the leakage `Leak` from `leak` and outputs the program counter:

```
1 sim :: SState -> Leak -> (SState, Maybe Address)
```

Like `leak`, `sim` focuses on timing: it must know when `leak` sends a jump address and when the core outputs the program counter (instead of the address for a memory access). So, `sim` too structurally mirrors `proc`: it’s a 5-stage pipeline. `sim`’s pipeline is simple—it models only the core’s stalling behavior (which also determines when `leak` sends a jump address).

RQ2: Proof Effort. Table 1 shows proof statistics for AIMCORE and processors verified in LeaVe [41]. PANTOMIME proofs involve writing a simulator and state projection function, totalling around 200 lines of Haskell for AIMCORE.⁷ LeaVe proofs are invariant-based, consisting of assertions about the processor’s state space and relations between variables the solver can’t automatically find.

While writing 200 lines of simulator code might superficially seem like more work than writing a modest number of invariants, recall that in PANTOMIME the simulator—and hence leakage proof—is a simplified version of the core and is therefore easy to write. Moreover, simulators are just programs: they can be interactively executed and benefit from a broad ecosystem of programming aids—from static type checking to property-based testing frameworks like QuickCheck [14] (which we used extensively during development). Debugging incorrect or missing invariants, however, is notoriously challenging and requires expertise beyond hardware design.

RQ3: How does AIMCORE compare to other verified processors? Table 1 shows a comparison between AIMCORE and the processors verified in LeaVe. AIMCORE is similar in complexity to other verified cores. While the other cores implement 2-, and 3-stage pipelines, AIMCORE implements a more complex 5-stage pipeline. The largest processor in LeaVe, Ibex, implements further instructions outside the integer base-set, e.g., instructions for manipulating CSR registers, multiplication and division, and support for compressed instructions. We note that these extensions (except for multiplication and division), have been disabled during verification. Ibex and AIMCORE are comparable in size when measured in lines of code.

RQ4: Leakage guarantees vs. other verified cores. Our leakage description for AIMCORE leaks the class of instruction, the register dependencies, and the address of branch targets. As such, our leakage description verifies that both the constant-time discipline, as well as branch balancing (as long as it respects register dependencies) are effective against timing side-channels on this architecture. Table 1 summarizes the leakage descriptions verified by LeaVe. Notably, the leakage descriptions computed by LeaVe are less precise than that of PANTOMIME. For instance, every processor in Table 1 leaks the program counter value along with the executed

⁶Which is just a `newtype` wrapper around the `RWS` monad from the `mtl` package.

⁷The leakage circuit (which isn’t part of the proof) is about 300 lines of Haskell and relies on an ISA interpreter that’s a further 70 lines.

Table 1: AIMCORE compared with the processors verified in LeaVe [41].

	AIMCORE (Our work)	Processors verified by LeaVe [41]		
		Sodor	DarkRISCv-3	Ibex-small
<i>Architecture</i>				
ISA	RV32I	RV32I	RV32E	RV32IMC
Pipeline stages	5	2	3	2
Code size (loc)	800 Haskell / 2400 Verilog	400 Chisel / 2000 Verilog	620 Verilog	2500 Verilog
Forwarding	✓	✗	✗	✓
<i>Security Properties & Proof</i>				
Leakage	Inst types, Reg deps, Jump Addr	PC, Inst, Br taken	PC, Inst	PC, Inst, Br taken, Jump Addr, Div Op
Proof effort	180 loc simulator; 30 loc projection	16 manual invariants	13 manual invariants	59 manual invariants
Verification Time	18.4 seconds	97.8 min	11.1 min	118.7 min
Unconditional proof	✓	✗	✗	✗

instruction; in our framework, this would result in a trivial simulator that directly outputs the leaked program counter. We conjecture that invariants restrict the class of leakages for which proofs are tractable. Invariant-based approaches also assume an immutable program, disallowing processors from modifying their own code (as done in e.g., boot loaders and JIT compilers); PANTOMIME does not have this restriction. PANTOMIME also enables us to group instructions that have the same leakage behavior using the `LOther` instruction. Finally, we note that LeaVe’s proofs are conditional on functional correctness of the processor pipeline, which itself was not formally proven. By contrast, our proofs are unconditional.

RQ5: Verification Time. Table 1 reports the verification time for AIMCORE using PANTOMIME, alongside results for processors verified by LeaVe. LeaVe doesn’t report on the hardware used for benchmarking; for AIMCORE, we benchmarked using a consumer-grade AMD Ryzen 7 9700X CPU. While the verification results are not directly comparable as they concern different designs, we can see that PANTOMIME’s verification time for AIMCORE is two orders of magnitude lower than LeaVe’s verification time for Ibex-small. We attribute this to the fact that LeaVe has to rely on expensive solvers to compute inductive invariants, while PANTOMIME only has to perform one-step equivalence checking between implementation and simulator. As a result, PANTOMIME allows users to check proofs interactively—as they write the core—simplifying development.

7 Related Work

Verifying Security Properties of RTL Designs. UPEC [18] detects transient execution vulnerabilities in RTL designs (or proves their absence), but is restricted to fixed properties, while PANTOMIME can express arbitrary leakage properties via (stateful) observation functions and leakages. ConjunCT [16] and UPEC-DIT [15] identify subsets of a processor’s instruction set that are data-oblivious in the sense that their operands do not affect the timing of the computation. H-Houdini [17] checks for the same property, but scales better by proposing a new invariant synthesis approach that exploits locality, e.g., due to pipelining. While these techniques scale well, they offer limited guarantees: since branch instructions trivially affect timing, these instructions are not data-oblivious. Hence, these techniques cannot guarantee safety of code involving branches, as is used, e.g., in constant-time code in cryptographic libraries. SecVerilog [47] extends Verilog with a type system to statically check timing-sensitive information flow. SpecVerilog [46] builds upon

SecVerilog’s type system to express information-flow safety under speculative execution. Unlike PANTOMIME, they are too restrictive to capture software defenses like the constant-time discipline.

Verifying Leakage Descriptions. Iodine [37] proves secret-independent timing in hardware, given usage assumptions expressed on inputs and internal wires. Xenon [38] synthesizes such usage assumptions semi-automatically. While their ability to represent arbitrary assumptions makes Iodine and Xenon expressive, they can only do so on internal wires, making translation to assumptions on software difficult. By contrast, PANTOMIME can express assumptions directly on the instruction stream of the program, as described in § 4 and 6. LeaVe [41] verifies leakage contracts (Definition 3.15) of RTL designs. However, it relies on expensive solvers, and requires hand-written invariants to carry out proofs. By contrast, PANTOMIME uses simulation-based proofs written in the same language as the processor, making them easier to write and fast to check. Contract Shadow Logic [36] follows a similar approach to LeaVe, but uses exhaustive state space exploration via model checking instead of inductive invariants to validate contracts, which limits its scalability. Both Contract Shadow Logic and LeaVe assume functional correctness of the design against an ISA-level specification to prove contract satisfaction. PANTOMIME makes no such assumption and therefore delivers unconditional proofs.

Synthesizing Leakage Descriptions. [29] synthesizes leakage contracts using a user-supplied lists of “contract atoms,” which represent potential leakage sources, along with a set of test cases. However, [29] does not prove the correctness of the leakage contracts it synthesizes. RTL2MμPATH [20] uses model checking to enumerate microarchitectural paths of an instruction through the processor and records which instructions may influence the path choice. It requires designers to annotate with $\mu FSMs$ —microarchitectural finite state machines that govern updates of the processor state and relies on exhaustive state space exploration which limits scalability. Their notion of leakage as dependencies between instructions makes it hard to translate assumptions back to software, limiting generality.

Simulation-based Proofs in Crypto. In designing our approach, we were inspired by simulation-based proofs in cryptography [24] (as used e.g., in Universal Composability [10]). However, beyond the common idea of using simulators, the two proof methods diverge

significantly. In cryptography, the aim is often to show that a (potentially active) adversary—with bounded computational resources—can learn nothing of interest about the secrets processed in a cryptographic protocol. By contrast, we use simulators to precisely characterize the side-channel leakage of a processors against a passive attacker—one that can observe side-channels (e.g., timing), but cannot actively influence the computation.

8 Conclusion

This paper introduces simulation-based leakage proofs, an approach to leakage verification that integrates directly into the hardware design process, enabling rapid, interactive verification with unconditional leakage guarantees. We realized this approach in the PANTOMIME verification tool and used it to verify the AIMCORE RISC-V CPU. Unlike previous work, where proofs depend on the functional correctness of a CPU against an ISA spec, our proofs are unconditional.

References

- [1] Johan Agat. 2000. Transforming out timing leaks. In *Proceedings of the 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (Boston, MA, USA) (POPL '00). Association for Computing Machinery, New York, NY, USA, 40–53. doi:10.1145/325694.325702
- [2] José Baelar Almeida, Manuel Barbosa, Gilles Barthe, François Dupressoir, and Michael Emmi. 2016. Verifying {Constant-Time} Implementations. In *25th USENIX Security Symposium* (USENIX Security 16). 53–70.
- [3] Marc Andryscio, David Kohlbrenner, Keaton Mowery, Ranjit Jhala, Sorin Lerner, and Hovav Shacham. 2015. On subnormal floating point and abnormal timing. In *2015 IEEE Symposium on Security and Privacy*. IEEE, 623–639.
- [4] Armaiti Ardeshiricham, Wei Hu, and Ryan Kastner. 2017. Clepsydra: Modeling timing flows in hardware designs. In *2017 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. 147–154. doi:10.1109/ICCAD.2017.8203772
- [5] Anish Athalye, M. Frans Kaashoek, and Nikolai Zeldovich. 2022. Verifying Hardware Security Modules with Information-Preserving Refinement. In *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*. USENIX Association, Carlsbad, CA, 503–519. <https://www.usenix.org/conference/osdi22/presentation/athalye>
- [6] Konstantinos Athanasiou, Byron Cook, Michael Emmi, Colm MacCarthaigh, Daniel Schwartz-Narbonne, and Serdar Tasiran. 2018. SideTrail: Verifying Time-Balancing of Cryptosystems. In *Verified Software. Theories, Tools, and Experiments*, Ruzica Piskac and Philipp Rümmer (Eds.). Springer International Publishing, Cham, 215–228.
- [7] C.P.R. Baaij. 2015. *Digital circuit in ClASH: functional specifications and type-directed synthesis*. PhD Thesis - Research UT, graduation UT. University of Twente, Netherlands. doi:10.3990/1.9789036538039 eemcs-eprint-23939.
- [8] Per Bjesse, Koen Claessen, Mary Sheeran, and Satnam Singh. 1998. Lava: hardware design in Haskell. *SIGPLAN Not.* 34, 1 (Sept. 1998), 174–184. doi:10.1145/291251.289440
- [9] Pietro Borrello, Daniele Cono D’Elia, Leonardo Querzoni, and Cristiano Giuffrida. 2021. Constantine: Automatic Side-Channel Resistance Using Efficient Control and Data Flow Linearization. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security* (Virtual Event, Republic of Korea) (CCS ’21). Association for Computing Machinery, New York, NY, USA, 715–733. doi:10.1145/3460120.3484583
- [10] Ran Canetti. 2000. Universally Composable Security: A New Paradigm for Cryptographic Protocols. Cryptology ePrint Archive, Paper 2000/067. <https://eprint.iacr.org/2000/067>
- [11] Sunjay Cauligi, Craig Disselkoen, Klaus v. Gleissenthall, Dean Tullsen, Deian Stefan, Tamara Rezk, and Gilles Barthe. 2020. Constant-time foundations for the new spectre era. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation* (London, UK) (PLDI 2020). Association for Computing Machinery, New York, NY, USA, 913–926. doi:10.1145/3385412.3385970
- [12] Sunjay Cauligi, Gary Soeller, Fraser Brown, Brian Johannsmeyer, Yunlu Huang, Ranjit Jhala, and Deian Stefan. 2017. FaCT: A Flexible, Constant-Time Programming Language. In *2017 IEEE Cybersecurity Development (SecDev)*. 69–76. doi:10.1109/SecDev.2017.24
- [13] Boru Chen, Yingchen Wang, Pradyumna Shome, Christopher W. Fletcher, David Kohlbrenner, Riccardo Paccagnella, and Daniel Genkin. 2024. GoFetch: Breaking Constant-Time Cryptographic Implementations Using Data Memory-Dependent Prefetchers. In *USENIX Security*.
- [14] Koen Claessen and John Hughes. 2011. QuickCheck: a lightweight tool for random testing of Haskell programs. *SIGPLAN Not.* 46, 4 (May 2011), 53–64. doi:10.1145/1988042.1988046
- [15] Lucas Deutschmann, Johannes Müller, Mohammad Rahmani Fadiheh, Dominik Stoffel, and Wolfgang Kunz. 2024. A Scalable Formal Verification Methodology for Data-Oblivious Hardware. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 43, 9 (2024), 2551–2564. doi:10.1109/TCAD.2024.3374249
- [16] Sushant Dinesh, Madhusudan Parthasarathy, and Christopher W. Fletcher. 2024. Conjunct: Learning Inductive Invariants to Prove Unbounded Instruction Safety Against Microarchitectural Timing Attacks. In *2024 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, Los Alamitos, CA, USA, 3735–3753. doi:10.1109/SP54263.2024.00180
- [17] Sushant Dinesh, Yongye Zhu, and Christopher W. Fletcher. 2025. H-Houdini: Scalable Invariant Learning. In *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1* (Rotterdam, Netherlands) (ASPLOS ’25). Association for Computing Machinery, New York, NY, USA, 603–618. doi:10.1145/3669940.3707263
- [18] Mohammad Rahmani Fadiheh, Alex Wezel, Johannes Müller, Jorg Bormann, Sayak Ray, Jason M. Fung, Subhashish Mitra, Dominik Stoffel, and Wolfgang Kunz. 2023. An Exhaustive Approach to Detecting Transient Execution Side Channels in RTL Designs of Processors. *IEEE Trans. Comput.* 72, 01 (Jan. 2023), 222–235. doi:10.1109/TC.2022.3152666
- [19] Mike Gordon. 1995. The Semantic Challenge of Verilog HDL. In *Proceedings of the 10th Annual IEEE Symposium on Logic in Computer Science (LICS ’95)*. IEEE Computer Society, USA, 136.
- [20] Yao Hsiao, Nikos Nikolieris, Artem Khzyha, Dominic P. Mulligan, Gustavo Petri, Christopher W. Fletcher, and Caroline Trippel. 2024. RTL2MμPATH: Multi-μPATH Synthesis with Applications to Hardware Security Verification. In *2024 57th IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 507–524. doi:10.1109/MICRO61859.2024.00045
- [21] Paul Kocher, Jann Horn, Anders Fogh, Daniel Genkin, Daniel Gruss, Werner Haas, Mike Hamburg, Moritz Lipp, Stefan Mangard, Thomas Prescher, et al. 2019. Spectre attacks: Exploiting speculative execution. In *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1–19.
- [22] Paul C Kocher. 1996. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In *Annual International Cryptology Conference*. Springer, 104–113.
- [23] Leslie Lamport and Stephan Merz. 2022. Prophecy Made Simple. *ACM Trans. Program. Lang. Syst.* 44, 2, Article 6 (April 2022), 27 pages. doi:10.1145/3492545
- [24] Yehuda Lindell. 2016. How To Simulate It – A Tutorial on the Simulation Proof Technique. Cryptology ePrint Archive, Report 2016/046. <https://eprint.iacr.org/2016/046> <https://eprint.iacr.org/2016/046.pdf>
- [25] Fangfei Liu, Yuval Yarom, Qian Ge, Gernot Heiser, and Ruby B Lee. 2015. Last-Level Cache Side-Channel Attacks are Practical. In *2015 IEEE Symposium on Security and Privacy*. IEEE, 605–622.
- [26] Sirui Lu and Rastislav Bodík. 2023. Grisette: Symbolic Compilation as a Functional Programming Library. *Proc. ACM Program. Lang.* 7, POPL, Article 16 (Jan. 2023), 33 pages. doi:10.1145/3571209
- [27] Haojun Ma, Aman Goel, Jean-Baptiste Jeannin, Manos Kapritsos, Baris Kasikci, and Karem A. Sakallah. 2019. I4: incremental inference of inductive invariants for verification of distributed protocols. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles* (Huntsville, Ontario, Canada) (SOSP ’19). Association for Computing Machinery, New York, NY, USA, 370–384. doi:10.1145/3341301.3359651
- [28] Daniel Moghimi. 2023. Downfall: Exploiting Speculative Data Gathering. In *32th USENIX Security Symposium* (USENIX Security 2023).
- [29] Gideon Mohr, Marco Guarnieri, and Jan Reineke. 2024. Synthesizing Hardware-Software Leakage Contracts for RISC-V Open-Source Processors. In *2024 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. 1–6. doi:10.23919/DAT58400.2024.10546681
- [30] Dag Arne Osvik, Adi Shamir, and Eran Tromer. 2006. Cache attacks and countermeasures: the case of AES. In *Topics in Cryptology—CT-RSA 2006*. Springer, 1–20.
- [31] Colin Percival. 2005. *Cache missing for fun and profit*. Technical Report. BSDCan.
- [32] Benjamin C. Pierce. 2002. *Types and Programming Languages* (1st ed.). The MIT Press.
- [33] Hany Ragab, Enrico Barberis, Herbert Bos, and Cristiano Giuffrida. 2021. Rage Against the Machine Clear: A Systematic Analysis of Machine Clears and Their Implications for Transient Execution Attacks. In *30th USENIX Security Symposium* (USENIX Security 21). USENIX Association, 1451–1468. <https://www.usenix.org/conference/usenixsecurity21/presentation/ragab>
- [34] Mary Sheeran. 2005. Hardware Design and Functional Programming: a Perfect Match. *JUCS - Journal of Universal Computer Science* 11, 7 (2005), 1135–1158. arXiv:https://doi.org/10.3217/jucs-011-07-1135 doi:10.3217/jucs-011-07-1135

- [35] Martin Sulzmann, Manuel MT Chakravarty, Simon Peyton Jones, and Kevin Donnelly. 2007. System F with type equality coercions. In *Proceedings of the 2007 ACM SIGPLAN international workshop on Types in languages design and implementation*. 53–66.
- [36] Qinhan Tan, Yuheng Yang, Thomas Bourgeat, Sharad Malik, and Mengjia Yan. 2024. RTL Verification for Secure Speculation Using Contract Shadow Logic. arXiv:2407.12232 [cs.AR] <https://arxiv.org/abs/2407.12232>
- [37] Klaus v. Gleissenthall, Rami Gökhan Kıcı, Deian Stefan, and Ranjit Jhala. 2019. IO-DINE: Verifying Constant-Time Execution of Hardware. In *28th USENIX Security Symposium (USENIX Security 19)*. USENIX Association, Santa Clara, CA, 1411–1428. <https://www.usenix.org/conference/usenixsecurity19/presentation/von-gleissenthall>
- [38] Klaus v. Gleissenthall, Rami Gökhan Kıcı, Deian Stefan, and Ranjit Jhala. 2021. Solver-Aided Constant-Time Hardware Verification. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (Virtual Event, Republic of Korea) (CCS '21)*. Association for Computing Machinery, New York, NY, USA, 429–444. doi:10.1145/3460120.3484810
- [39] Jo Van Bulck, Daniel Moghimi, Michael Schwarz, Moritz Lipp, Marina Minkin, Daniel Genkin, Yarom Yuval, Berk Sunar, Daniel Gruss, and Frank Piessens. 2020. LVI: Hijacking Transient Execution through Microarchitectural Load Value Injection. In *41th IEEE Symposium on Security and Privacy (S&P'20)*.
- [40] Stephan van Schaik, Alyssa Milburn, Sebastian Österlund, Pietro Frigo, Giorgi Maisuradze, Kaveh Razavi, Herbert Bos, and Cristiano Giuffrida. 2019. RIDL: Rogue In-flight Data Load. In *S&P*. Paper=<https://mdsattacks.com/files/ridl.pdf>Slides=<https://mdsattacks.com/slides/slides.html>Web=<https://mdsattacks.com>Code=<https://github.com/vusec/ridl>Press=<http://mdsattacks.com> Intel Bounty Reward (Highest To Date), Pwnie Award Nomination for Most Innovative Research, CSAW Best Paper Award Runner-up, DCSR Paper Award.
- [41] Zilong Wang, Gideon Mohr, Klaus von Gleissenthall, Jan Reineke, and Marco Guarnieri. 2023. Specification and Verification of Side-channel Security for Open-source Processors via Leakage Contracts. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (Copenhagen, Denmark) (CCS '23)*. Association for Computing Machinery, New York, NY, USA, 2128–2142. doi:10.1145/3576915.3623192
- [42] Andrew Waterman and Krste Asanović. 2024. *The RISC-V Instruction Set Manual Volume I: Unprivileged ISA*. RISC-V Foundation. <https://lf-riscv.atlassian.net/wiki/spaces/HOME/pages/16154769/RISC-V+Technical+Specifications> Document Version 20240411.
- [43] Conrad Watt, John Renner, Natalie Popescu, Sunjay Cauligi, and Deian Stefan. 2019. CT-wasm: type-driven secure cryptography for the web ecosystem. *Proc. ACM Program. Lang.* 3, POPL, Article 77 (Jan. 2019), 29 pages. doi:10.1145/3290390
- [44] Hans Winderix, Marton Bogner, Lesly-Ann Daniel, and Frank Piessens. 2024. Libra: Architectural Support For Principled, Secure And Efficient Balanced Execution On High-End Processors. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security (Salt Lake City, UT, USA) (CCS '24)*. Association for Computing Machinery, New York, NY, USA, 19–33. doi:10.1145/3658644.3690319
- [45] Yuval Yarom and Katrina Falkner. 2014. FLUSH+RELOAD: A High Resolution, Low Noise, L3 Cache Side-Channel Attack. In *23rd USENIX Security Symposium*. 719–732.
- [46] Drew Zagieboylo, Charles Sherk, Andrew C. Myers, and G. Edward Suh. 2023. SpecVerilog: Adapting Information Flow Control for Secure Speculation. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (Copenhagen, Denmark) (CCS '23)*. Association for Computing Machinery, New York, NY, USA, 2068–2082. doi:10.1145/3576915.3623074
- [47] Danfeng Zhang, Yao Wang, G. Edward Suh, and Andrew C. Myers. 2015. A Hardware Design Language for Timing-Sensitive Information-Flow Security. *SIGARCH Comput. Archit. News* 43, 1 (March 2015), 503–516. doi:10.1145/2786763.2694372

A Proofs for Section 3

PROOF OF THEOREM 3.10. Let s be a state of type S and is be an input sequence of Is . We proceed by induction on is .

- **Base Case:** $is = []$. By definition.
- **Induction hypothesis:** $M(c_1, s, is') = M(c_2, p, s, is')$.

- **Inductive Step:** $is = i : is'$. We show the equivalence between $M(c_1, s, is)$ and $M(c_2, p, s, is)$ as follows:

$$\begin{aligned}
 & M(c_1, s, is) \\
 & \equiv^{(1)} o : M(c_1, s', is') \\
 & \equiv^{(2)} o : M(c_2, p, s', is') \\
 & \equiv^{(3)} o : M(c_2, \llbracket c_1 p \ s \ i \rrbracket .1, is') \\
 & \equiv^{(4)} o : M(c_2, \llbracket pc_2 \ s \ i \rrbracket .1, is') \\
 & \equiv^{(5)} o : M(c_2, (c_2(p \ s) \ i).1, is') \\
 & \equiv^{(6)} M(c_2, p, s, is).
 \end{aligned}$$

Step (1) follows by definition of M , where $(s', o) = \llbracket c_1 \ s \ i \rrbracket$; step (2) from our induction hypothesis; and step (3) from the definition of $c_1 p$ in the reduction proof, where 1 is the label for the first element of a tuple. Step (4) follows from the observational equivalence between $c_1 p$ and pc_2 , which we get from the assumption $c_1 \hookrightarrow c_2$. Step (5) follows from the definition of pc_2 ; and step (6) from the definition of M . \square

PROOF OF THEOREM 3.16. Let $\text{sim} : S_s \rightarrow O_l \rightarrow (S_s, O_o)$ be the simulator, with respect to which we have the mealy reduction $c \circ o \hookrightarrow \text{leak} \circ \text{sim}$. Given an initial state $(s, s_o) : (S, O)$ of the circuit $c \circ o$, from Theorem 3.10 we get an initial state $(s_l, s_s) : (S_l, S_s)$ of the circuit $\text{leak} \circ \text{sim}$ such that

$$M(c \circ o, (s, s_o), is) = M(\text{leak} \circ \text{sim}, (s_l, s_s), is) \quad (\star)$$

for all input sequences $is : [I]$. Fix input sequences $is, is' : [I]$. Then we have

$$\begin{aligned}
 & M(c \circ o, (s, s_o), is) \\
 & \equiv^{(1)} M(\text{leak} \circ \text{sim}, (s_l, s_s), is) \\
 & \equiv^{(2)} M(\text{sim}, s_s, M(\text{leak}, s_l, is)) \\
 & \equiv^{(3)} M(\text{sim}, s_s, M(\text{leak}, s_l, is')) \\
 & \equiv^{(4)} M(\text{leak} \circ \text{sim}, (s_l, s_s), is') \\
 & \equiv^{(5)} M(c \circ o, (s, s_o), is')
 \end{aligned}$$

Steps (1) and (5) follow from (\star) , steps (2) and (4) from Lemma 3.12, and step (3) from the assumption. \square

B Verilog Semantics

We now sketch the semantics of μ -Verilog. For a more thorough description of Verilog's semantics, see [19, 37, 41, 47]. The circuit state consists of a pair (ρ, μ) , where ρ maps variables to values, and μ maps registers to values. Values are fixed length bit-strings, where we write $\text{size}(v)$ to denote the number of bits in v . At the beginning of each clock cycle, we update ρ with the valuation of the inputs in the current cycle. That is, if the value for input i is value b in the current cycle, we set $\rho(i) = b$. Next, we execute wire assignments updating ρ accordingly, until we reach a fix-point, i.e., until re-executing any wires no longer changes ρ . For this, we make use of an evaluation function $\llbracket \cdot \rrbracket^{\rho, \mu}$, where $\llbracket e \rrbracket^{\rho, \mu}$ denotes the result of evaluating expression e under wire valuation ρ , and registers valuation μ . Finally, we pass the values of output wires on as outputs of the circuit. After assigning variables, we perform updates to registers stemming from assignments of the form $r \leftarrow v$.

For this, we construct μ' , the register valuation in the next clock cycle by setting $\mu' = \mu[r \leftarrow \rho(v)]$, for each $r \leftarrow v \in A$, where we write $f[x \leftarrow v]$ to denote the function that acts like f everywhere, except for x , where it returns v . Note that since we computed ρ using μ thereby using the values of registers in the last clock-cycle, effectively all register updates are performed at once. This is crucial, and reflects Verilog's synchronous semantics.

C Translation from STLC to Verilog

We now describe a translation from STLC circuits to μ -Verilog programs. We first change the names of the variables in the circuit to make sure each lambda abstraction and case expression variable has a unique name. We exhaustively apply β -reductions and η -expansions. β -reduction is defined in the evaluation rules of the core language. η -expansion expands function terms, that is, for each function term f that is not a lambda abstraction and does not occur in the function position of an application, we replace it with the abstraction $\lambda x. f x$, where x is fresh. Exhaustively applying these two rules ensures that there are no partially applied functions left, and all remaining applied functions are primitive functions that can be represented in μ -Verilog.

Representing Records and Variants. Our translation needs to encode record and variant types into variables. For records, it encodes the whole record into a variable of appropriate size, and then uses the bit-slicing operator to extract the parts. For variants, it encodes a representation of the label together with the actual data. We now define utility functions that help with this encoding.

Representation Size. The function $sizeof(t)$ computes the number of bits required to represent a type t in μ -Verilog. A base-type of size $Word_n$ requires n bits. For a record type $t = \{l_1 : t_1, \dots, l_n : t_n\}$, we require the sum of the sizes of the record types, i.e., $sizeof(t) = \sum_i sizeof(t_i)$. For a variant type $t = \langle l_1 : t_1, \dots, l_n : t_n \rangle$, the size is defined as $sizeof(t) = \max_i(sizeof(t_i)) + \lceil \log_2(n) \rceil$, to account for both the size of the largest field and the size of the label.

Field Extraction. Field extraction function $field(v, l)$ lets us access field l of a record type encoded into variable v . For a record of type $\{l_1 : t_1, \dots, l_n : t_n\}$, the expression $field(v, l_n)$ is defined as $v[\sum_{i=0}^{n-1} sizeof(t_i) : \sum_{i=0}^n sizeof(t_i)]$, meaning the field is accessed by slicing the appropriate bit range based on the total size of the preceding fields.

Label Encoding. The function $tolabel(l_i)$ converts the label l_i of a variant type $t = \langle l_1 : t_1, \dots, l_n : t_n \rangle$ to a μ -Verilog bit-vector of size $\lceil \log_2(n) \rceil$ representing the integer i , which is a tag for the variant case corresponding to l_i . For example, for the variant type $\langle Just : Word8, Nothing : () \rangle$, $tolabel(Just) = 0$ and $tolabel(Nothing) = 1$.

Label Extraction. Similar to $field$, the label extraction function $label(v, l)$ lets us access the label l of a record type encoded into variable v . For a record of type $\{l_1 : t_1, \dots, l_n : t_n\}$, the expression $label(v, l_m)$ is defined as $v[\sum_{i=0}^m sizeof(t_i) : \sum_{i=0}^m sizeof(t_i) + \lceil \log_2(n) \rceil]$.

Casting. The function $cast(v, l)$ casts a variable by assuming its size has label ls size where l is a label of a variant type $\langle l_1 : t_1, \dots, l_n : t_n \rangle$. For example, for the variant type $\langle First : Word8, Second : Word9 \rangle$, $cast(v, First) = v[0 : 8]$ and

$cast(v, Second) = v[0 : 9]$. We use the casting function while transforming a case-expression. The translated μ -Verilog circuit includes a wire for each e_i in $\text{case } e_c \text{ of } l_i x_i \rightarrow e_i$.

Mapping STLC Operators to μ -Verilog Operators. We assume there exist corresponding unary and binary μ -Verilog operators for each unary and binary STLC operator. For each unary STLC operator uop , we represent the corresponding μ -Verilog operator as \ominus_{uop} . Similarly, for each binary STLC operator bop , we represent the corresponding μ -Verilog operator as \otimes_{bop} . Note that in STLC, there are no built-in operators. Operators are represented as primitive functions.

Translation Function. We translate an STLC circuit $\lambda s \text{ i. } e$ into the μ -Verilog circuit

$$C = \{v_i\} : T(e, V_0) \cup \{o = field(v_e, 2), s' = field(v_e, 1)\} : \{s \leftarrow s'\} : \{o\}.$$

The initial valuation is $V_0 = \{s \mapsto v_s, i \mapsto v_i\}$. The translation function T , defined in Figure 7, recursively translates each STLC subexpression into a corresponding μ -Verilog wire. It takes a valuation V as an additional argument, which maps μ -Verilog variables and registers to μ -Verilog expressions. This valuation tracks bindings introduced by case expressions and bookkeeps assumptions made during translation.

For example, when translating the expression

$$e = \text{case } e_c \text{ of} \\ \quad \text{Just } x \rightarrow e_j \\ \quad \text{Nothing} \rightarrow e_n$$

we must translate e_j under the assumption that x refers to a Just-tagged value carried by e_c . If Just is of type $Word4$, we extend the valuation as $V' = V[x \mapsto v_{e_c}[0 : 4]]$ while evaluating $T(e_j, V')$.

Example C.1. We now demonstrate our STLC to μ -Verilog translation with a complete example. Consider the STLC circuit

$$e : Word8 \rightarrow \text{Maybe } Word8 \rightarrow (Word8, Word8) \\ e = \lambda s \text{ i. } (\lambda f. (fs, fs)) \\ (\text{case } i \text{ of} \\ \quad \text{Just } x \rightarrow (-) x \\ \quad \text{Nothing} \rightarrow (+) 1)$$

To translate e to a μ -Verilog circuit, we first exhaustively apply η -expansions and β reductions to get

$$e' : Word8 \rightarrow \text{Maybe } Word8 \rightarrow (Word8, Word8) \\ e' = \lambda s \text{ i.} \\ (\text{case } i \text{ of} \\ \quad \text{Just } x \rightarrow x - s \\ \quad \text{Nothing} \rightarrow 1 + s, \\ \text{case } i \text{ of} \\ \quad \text{Just } x \rightarrow x - s \\ \quad \text{Nothing} \rightarrow 1 + s)$$

$$\begin{aligned}
 T(x, V) &= \{v_x = V(x)\} \\
 T(c, V) &= \{v_c = c\} \\
 T(e_1 \text{ bop } e_2, V) &= \{v_e = v_{e_1} \otimes_{\text{bop}} v_{e_2}\} \cup T(e_1, V) \cup T(e_2, V) \\
 T(uop\ e_1, V) &= \{v_e = \ominus_{\text{uop}} v_{e_1}\} \cup T(e_1, V) \\
 T(e_1.l, V) &= \{v_e = \text{field}(v_{e_1}, l)\} \cup T(e_1, V) \\
 T(\overrightarrow{\{l_i = e_i\}}, V) &= \{v_e = \{v_{e_1}, \dots, v_{e_n}\}\} \cup \bigcup_i T(e_i, V) \\
 T(1\ e, V) &= \{v_e = \{\text{tolabel}(1), v_e\}\} \cup T(e, V) \\
 T(\text{case } e_c \text{ of } \overrightarrow{l_i x_i \rightarrow e_i}, V) &= T(e_c, V) \cup \bigcup_i T(e_i, V[x_i \mapsto \text{cast}(e_c, l_i)]) \\
 &\quad \cup \{v_e = (\text{if } \text{label}(v_{e_c}) == 0 \text{ th } v_{e_0} \text{ el} \\
 &\quad \quad \text{if } \text{label}(v_{e_c}) == 1 \text{ th } v_{e_1} \text{ el } \dots \\
 &\quad \quad \text{if } \text{label}(v_{e_c}) == n - 1 \text{ th } v_{e_{n-1}} \text{ el } v_{e_n})\}
 \end{aligned}$$

Figure 7: The $T(e, V)$ transformation from Circuits to μ -Verilog.

We then apply the wire calculation function on the body of the circuit expression e' as

$$\begin{aligned}
 T(e_{\text{body}}, V_0) &= T(e_{\text{case}}, V_0) \cup \{v_{e'} = \{v_{\text{case}}, v_{\text{case}}\}\} \\
 &= T(x - s, V_0[x \mapsto v_i[0 : 8]]) \cup T(1 + s, V_0) \\
 &\quad \cup \{v_{\text{case}} = \text{if } v_i[8 : 9] == 0 \text{ th } v_{x-s} \text{ el } v_{1+s}\} \\
 &= T(x, V_0[x \mapsto v_i[0 : 8]]) \cup \{v_{x-s} = v_x - v_s\} \\
 &\quad \cup \{v_{1-s} = v_1 - v_s\} \cup T(1, V_0[x \mapsto v_i[0 : 8]]) \\
 &\quad \cup \{v_{\text{case}} = \text{if } v_i[8 : 9] == 0 \text{ th } v_{x-s} \text{ el } v_{1+s}\} \\
 &= \{v_x = v_i[0 : 8]\} \cup \{v_{x-s} = v_x - v_s\} \\
 &\quad \cup \{v_{1-s} = v_1 - v_s\} \cup \{v_1 = 1\} \\
 &\quad \cup \{v_{\text{case}} = \text{if } v_i[8 : 9] == 0 \text{ th } v_{x-s} \text{ el } v_{1+s}\} \\
 &= \{v_x = v_i[0 : 8], v_{x-s} = v_x - v_s, v_{1-s} = v_1 - v_s, \\
 &\quad v_1 = 1, v_{\text{case}} = \text{if } v_i[8 : 9] == 0 \text{ th } v_{x-s} \text{ el } v_{1+s}\}.
 \end{aligned}$$

Which finally gives us the circuit

$$C = \{v_i\} : T(e, V_0) \cup \{o = \text{field}(v_{\text{body}}, 2), s' = \text{field}(v_{\text{body}}, 1)\} : \{s \leftarrow s'\} : \{o\}.$$