Power Contracts: Provably Complete Power Leakage Models for Processors

Roderick Bloem*
Graz University of Technology
Graz, Austria

Barbara Gigerl
Graz University of Technology
Graz, Austria

Marc Gourjon
Hamburg University of Technology
Hamburg, Germany
NXP Semiconductors
Hamburg, Germany

Vedad Hadžić
Graz University of Technology
Graz, Austria

Stefan Mangard
Graz University of Technology
Graz, Austria

Robert Primas
Graz University of Technology
Graz, Austria

ABSTRACT

The protection of cryptographic software implementations against power-analysis attacks is critical for applications in embedded systems. A commonly used algorithmic countermeasure against these attacks is masking, a secret-sharing scheme that splits a sensitive computation into computations on multiple random shares. In practice, the security of masking schemes relies on several assumptions that are often violated by microarchitectural side-effects of CPUs. Many past works address this problem by studying these leakage effects and building corresponding leakage models that can then be integrated into a software verification workflow. However, these models have only been derived empirically, putting in question the otherwise rigorous security statements made with verification.

We solve this problem in two steps. First, we introduce a *contract* layer between the (CPU) hardware and the software that allows the specification of microarchitectural side-effects on masked software in an intuitive language. Second, we present a method for proving the correspondence between contracts and CPU netlists to ensure the completeness of the specified leakage models. Then, any further security proofs only need to happen between software and contract, which brings benefits such as reduced verification runtime, improved user experience, and the possibility of working with vendor-supplied contracts of CPUs whose design is not available on netlist-level due to IP restrictions. We apply our approach to the popular RISC-V IBEX core, provide a corresponding formally verified contract, and describe how this contract could be used to verify masked software implementations.

KEYWORDS

Power side-channel, Leakage model, Verification, Contract, Domain-specific language, Masking, Probing security

Author Email: first.last@iaik.tugraz.at and first.last@tuhh.de

*Authors listed in alphabetical order.

CCS '22, November 7–11, 2022, Los Angeles, CA, USA

1 INTRODUCTION

Physical side-channel attacks such as power or EM analysis allow attackers within proximity of a device to learn sensitive information like cryptographic keys [32, 47]. One of the most widely used algorithmic countermeasures for protecting a cryptographic implementation against these kinds of attacks is masking [14, 27, 31]. Masking is a secret-sharing technique that splits input and intermediate variables of cryptographic computations into $d \geq t+1$ random shares such that the observation of up to t shares does not reveal any information about their corresponding unmasked value. Masking schemes typically rely on certain assumptions, such as independent computations producing independent side-channel leakage. However, the structure of a CPU architecture can violate these assumptions and introduce additional leakage effects [15, 18, 25, 41, 45]. Such leakage is often referred to as *order-reducing leakage* because it induces a security loss and thus a gap between formal security assurance and practical resilience. The physical characteristics of CMOS gates are relatively well understood and give rise to extended leakage models, which allow constructing *hardware* implementations that reliably mitigate order-reducing leakage [17, 20, 35, 36]. Similarly, when designing masked *software* implementations of cryptographic algorithms, knowing the concrete power side-effects of different instruction types is indispensable. It allows developers to optimize the performance of masked implementations by simplifying the otherwise trial-and-error hardening process [9].

State of the Art. Many works address the problem of characterizing and understanding the leakage behavior of instructions. These can be divided into two categories: works that use empirical methods to determine side-channel leakage, and works that use formal verification approaches to verify side-channel resilience.

On the empirical side, the measurement of a CPU's power consumption combined with a subsequent analysis using statistical methods is a straightforward approach to determine whether cryptographic software is correctly masked. Any observed leakage effects can be reverse-engineered and taken into account in hardened versions of the respective masked software implementations [1, 23, 24, 38, 39, 45, 50]. The authors of ELMO [39] characterize leakage behavior by selecting "explanatory variables", e.g., operands of assembly instructions and determining whether the variables contribute to the measured leakage using statistical tests. Variables with a correlation to measured leakage form a "leakage model". Such a

model specifies the side-effects of a (sequence of) assembly instructions which may be exploited using for example Differential Power Analysis (DPA) [33, 40] to learn information about processed cryptographic keys. However, manual selection of variables, statistical methods and physical measurements bear potential for incomplete models which do not specify leakage of data which in practice could be exploited. Specifically, the restriction to sequences of three instructions in [39] may miss leakage effects spanning between two load instructions spaced by multiple instructions as observed in [9, 45]. Similar issues exist with other works that try to generalize power side-channel leakages that are often reported on a variety of devices in generic leakage models [9, 10, 39]. Any work building upon empirical or generalized models, e.g., ELMO [39], ROSITA [50] and TORNADO [11] only report, respectively protect against, vulnerabilities and leakages which are part of the model. Hence, all existing empirical approaches either require high practical effort or cannot guarantee completeness, thereby reducing the confidence in security assessments [5, 26, 41].

On the formal verification side, several works verify the security of masked software under specific masking-related security notions. MaskVerif performs algorithmic software masking verification using generic leakage models and supports the commonly acccepted t–NI and t–SNI security notions [7, 8]. scVerif improves upon this and verifies the absence of order-reducing leakage in user-provided leakage models for assembly instructions [9]. However, the security assurance still relies on the completeness of the used leakage models, as it is the case for pure empirical approaches.

Coco is a formal masking verification tool that avoids modeling the leakage by directly working with the processor netlist [25, 26, 30]. Their approach considers the leakage of every gate in an extended hardware leakage model and captures a wide range of microarchitectural side-effects. However, their method requires the processor's netlist, which may not be available.

Hardware-software contracts have previously been used to verify speculative and timing related side-channel resilience [29]. The contracts presented in this paper are more flexible and apply to arbitrary side-channel behavior, devices and software implementations. Their close correspondence to the respective ISAs facilitates understanding by users.

Our Contribution. We answer to the question of leakage model completeness and establish end-to-end (E2E) security for software executing on a processor. First, we introduce a *contract* between the hardware and the software that defines precise semantics and models side-channel behavior of assembly instructions. We then establish a technique to verify *compliance* of a processor with a contract. A processor is compliant when the leakage of each of its gates and the semantic of instruction is correctly specified in the contract. Put vice-versa, we prove the contract's model of instructions correct and its model of leakages complete. We pave the way for provable E2E security by defining software compliance for threshold probing security notions so that the approach of Barthe *et al.* [9] can be easily mapped to our slightly different language for contracts. We combine hardware compliance and software compliance to prove E2E security: *any* compliant software is secure w.r.t. *all* microarchitectural power side-channel leakage of *any* compliant CPU. Compared to related work, our approach comes with benefits such as more rigorous (practical) security statements, simplified software verification workflows, and the possibility of working with vendor-supplied contracts of CPUs whose design is not available on netlist-level due to IP restrictions. Our contributions also enable the construction of reliable hardened processors, as users can specify the desired leakage model in a contract and modify the CPU implementation to achieve compliance for fixed contracts. We emphasize that the intermediate contract layer enables for the first time portability of secure implementations across processors and improves the separation of secure hardware and software development in general. This separation also optionally allows the creation of vendor-supplied CPU contracts whose leakage specification is "high-level" on purpose to avoid potential IP-related issues yet modeling all real leakages in the actual hardware.

1. Contracts. We introduce an intuitive and industry-grade domain-specific language (DSL) called Genoa. Genoa allows specifying Instruction Set Architecture (ISA) semantics and device-specific leakages in contracts. Genoa extends the long-standing Sail language [4] to support leakage specifications. The RISC-V foundation recently picked Sail as the official tool to specify the reference RISC-V ISA and all standard extensions [42, 49]. Models for multiple architectures (e.g., ARM) exist, which can be freely adopted and compiled to software emulators [4]. We reuse existing models as the basis for contracts, augmenting them with leakage specifications and providing an interface for our verification tool. We show that whenever a program is secure with respect to a contract, its concrete execution on a compliant processor is also secure, i.e., no order-reducing leakage can occur. We emphasize that our contracts also support higher-order masking, branching, and secret-dependent memory accesses needed for masked table lookups [16].

2. Hardware Compliance. We present a method to automatically verify the compliance of a processor with a contract. Verification ensures that the leakage of every gate is captured by a leakage specification in the contract and that the contract specifies correct instruction semantics. Our methodology is based on the intuition that if the contract properly models the hardware, then any leakage arising in the hardware can be computed from leakage produced during an execution in the contract. It is hence up to the contract designer if they want to create exact specifications for every leakage of a CPU design, a high-level specification that does not contain any information about the CPUs microarchitecture (while still covering all real leakages), or any trade-off between the two. The verification encodes both hardware and contract execution, respectively leakage, as SMT formulas and checks for model gaps using the SMT solver Z3 [6, 19]. If the solver finds no cases where a hardware leakage is not modeled from the contract leakage, we have proven hardware compliance.

3. Case Study. We implement our methods in a tool and apply them to the popular RISC-V IBEX core [34], resulting in a verifiably complete contract for a wide range of instructions that are commonly used for cryptographic implementations. IBEX is a low-end processor, suitable for embedded or IoT applications that require cryptographic computations. Our analysis is focused on gate-level power leakage and we consider features such as speculative execution or data caches as out of scope since they are mostly deactivated while executing masked programs.

We showcase the applicability of contracts by incorporating them into scVerif and checking the software compliance of several masked programs. We then validate our results by checking the same programs with the independent verifiction tool Coco that directly uses the IBEX netlist. The contract, which is based on the official RISC-V reference models, is provided as part of the paper and its appendix.

2 SIDE-CHANNEL RESILIENCE

We introduce preliminaries for side-channel security. Hardware circuits and their power side-channel leakage are modeled in Section 2.1. In Section 2.2, we recall the masking countermeasure and the formal notions of provable side-channel resilience.

2.1 Hardware Model and Gate-level Leakage

Processors are digital hardware circuits which can be modeled using labeled directed graphs. For any given circuit (G, W, L), we say that G is the set of gates, $W \subseteq G \times G$ is the set of wires connecting the gates, and $L : G \to T$ is a labeling defining the type $\tau \in T$ of each gate $g \in G$. The types T depend on the technology that realizes the circuit. In addition to combinatorial gates we only require that the technology contains an input type τ_{in} and a register type τ_{reg}. Input gates only have outgoing wires, and register gates only have one incoming wire. Additionally, every cyclic path in the circuit contains at least one register gate. The state of a circuit is completely defined by the values of its inputs and registers, referred to as *locations*, denoted with $V^h = \left\{g \in G \mid L(g) \in \left\{\tau_{\text{in}}, \tau_{\text{reg}}\right\}\right\}$. Superscript h is for hardware, later we use c for contract. Hardware states are denoted with $\sigma^h \in \mathbb{B}^{|V^h|}$, with optional subscripts. Any location $v^h \in V^h$ just returns the appropriate bit of the state. Any gate g is a function of a state, *i.e.*, $g : \mathbb{B}^{|V^h|} \to \mathbb{B}$ where gate $g \in G \setminus V$ combines state bits according to its type τ.

The execution of a circuit happens in clock cycles. For a state σ^h_j, we denote the next state as σ^h_{j+1}. The registers of the next state have values reflecting the values of their inputs in the previous cycle, *i.e.*, $g(\sigma^h_{j+1}) := g'(\sigma^h_j)$ with $(g', g) \in W$, whereas the next state of circuit inputs is determined by the environment.

We now proceed to define the power side-channel leakage that is exposed to an adversary. The root cause of power side-channels is that CMOS logic draws power, or emits electromagnetic radiation, mainly if a transistor switches its state. Thus, CMOS gates have a data-dependent power consumption. The leakage behavior of CMOS gates themselves is relatively well understood and can be modeled by a few simple leakage effects which allow an (idealized) *probing adversary* to observe the (intermediate) values of gates and wires without any loss due to measurement noise [5, 20, 31]. The seminal work of Ishai *et al.* [31] introduced *value leakage* which allows an idealized adversary to observe the value of any wire connected to a gate at the beginning or end of a cycle, *i.e.*, its stable signal. The value leakage λ_g exposed by gate g is its value $\lambda_g(\sigma^h_j) = g(\sigma^h_j)$ in the state σ^h_j. Besides value leakage, additional leakage effects are also observable in hardware. We define our extended probing model in close relation to the robust probing model of Faust *et al.* [20]. *Transition leakage* refers to the phenomeneon that the power consumption of CMOS gates depends on the charges (state) of the gate before computation. As such, transition leakage allows observing whether the value of a gate changed during a clock cycle but also whether the value changed from zero to one or vice-versa. Formally, our idealized adversary is able to observe the initial value and the resulting value of each gate. The observable gate leakage is then the concatenation of the old and new gate values, *i.e.*, $\lambda_g(\sigma^h_{j-1}, \sigma^h_j) = g(\sigma^h_{j-1}) || g(\sigma^h_j)$. This sufficiently captures any real-world transition leakage function computable from the old and new gate values. In addition to these main phenomena, there are *glitches* caused by propagation delay in the temporary logic states of combinatorial circuits within one clock cycle (and thus rather ephemeral) [37] and *couplings* caused by inductive coupling of adjacent wires [17]. These can be modeled by defining $\lambda_g(\sigma^h_{j-1}, \sigma^h_j)$ for non-register gates as the concatenation of all possible values the gate g could take on due to these effects. We use $\boldsymbol{\mathcal{L}}^h_{0,m}$ to denote the observable *gate-level* leakage throughout the execution starting in state σ^h_0 and ending in state σ^h_m. While the techniques described in the following apply to all effects, for the purpose of this paper, we focus on value leakage and transition leakage. Hence, we define $\boldsymbol{\mathcal{L}}^h_{0,m} = \{\lambda^h_g(\sigma^h_{j-1}, \sigma^h_j) \mid g \in G, 1 \le j < m\}$.

2.2 Provable Security and Simulatability

Applying masking to a cryptographic algorithm requires to replace the primitive operations (e.g., logical conjunction, exclusive or, addition) by masked computations, often called *gadgets*, which compute the same operation securely on shares [28, 31, 43, 48]. The challenge in the design and implementation of gadgets is to maintain the security of the secret-sharing: it must remain information theoretically impossible to learn the secrets or intermediate values by observing up to t leakages caused by the gadget. In sufficiently noisy environments this leads to an exponential gain of security in the order of t [13, 46]. However, many works do not take the full gate-level leakages into account, resulting in implementations that are exploitable at a lower-than-advertised security order. Especially for masked software, the resulting gap in the security assurance allows to break the implementation by observing, e.g., transition leakage of the processor executing the program [5, 35, 44]. This work aims to reduce the gap by enabling software security assessments to include the *complete* set of gate-level leakages.

We give an overview of the notation associated with masking, and formalize gadgets and their security. Masking heavily relies on random variables. We write the names of random variables in lowercase, e.g., x_i, and use lowercase boldface names for sets of variables, e.g., $\boldsymbol{x} = \{x_0, \ldots, x_n\}$. Each random variable x_i, respectively set $\boldsymbol{x}$, is associated with a probability distribution $\Pr[x_i]$, respectively $\Pr[\boldsymbol{x}]$. Each secret x_i is encoded (masked) using $d > t$ shares and we write $\overline{x_i} = \{x_i^0, \ldots, x_i^{d-1}\}$ for the shares which encode x_i, where x_i^j denotes for $0 \le i < n$ and $0 \le j < d$ the j^{th} share of the i^{th} secret. The superset of all shares is denoted by $\overline{\boldsymbol{x}} = \{\overline{\boldsymbol{x}}_0, \ldots, \overline{\boldsymbol{x}}_{n-1}\}$.

A gadget operates on input tuple $(\overline{\boldsymbol{x}}, \boldsymbol{r}, \boldsymbol{p})$ returning tuple $(\overline{\boldsymbol{y}}, \boldsymbol{o}, \boldsymbol{\mathcal{L}})$, each consisting of random variables. The input shares $\overline{\boldsymbol{x}}$ are a set of t-wise independent encodings $\overline{\boldsymbol{x_i}}$, each encoding a secret variable $\boldsymbol{x_i}$. $\boldsymbol{r}$ represents a set of independent and uniformly random variables, $\boldsymbol{p}$ are public inputs independent of secrets. The output of a gadget consist of output shares $\overline{\boldsymbol{y}}$, public outputs $\boldsymbol{o}$, and observable

leakage $\mathcal{L}$. Each individual output is a random variable y_i^j (respectively o_i) and computed as a function of the gadget's inputs, *i.e.*, $y_i^j = f_i^j(\overline{x}, \boldsymbol{r}, \boldsymbol{p})$. During its execution a gadget produces observable leakage $\mathcal{L} = \{\lambda_0(\overline{x}, \boldsymbol{r}, \boldsymbol{p}), \ldots, \lambda_m(\overline{x}, \boldsymbol{r}, \boldsymbol{p})\}$, which an attacker can observe, e.g., through power measurements. The attacker's goal is to learn information about the unshared secret inputs $\boldsymbol{x}$.

Threshold non-interference (t–NI) and strong threshold non-interference (t–SNI) are two prominent security notions for proving the security of gadgets against idealized adversaries [7, 8]. These have been extended in [9] into Stateful t–(S)NI to incorporate that physical execution involves state and public in- and outputs. Security of gadgets in these notions is shown by proving that the observations an attacker makes can be *simulated* without knowing the secret values, thereby proving that no information can be gained from t observations. In the following, we formalize what it means to simulate random variables, and restate t–NI and t–SNI.

Definition 1 (Simulation Procedure). *Let $\boldsymbol{c}$ and $\boldsymbol{h}$ be sets of possibly related random variables and $\boldsymbol{r}$ be a set of independent and uniformly distributed variables. The simulation procedure $S : Dom(\boldsymbol{c} \cup \boldsymbol{r}) \to Dom(\boldsymbol{h})$ (*simulator *for short) samples the random variables $\boldsymbol{r}$ to simulate the distribution of $\boldsymbol{h}$ from $\boldsymbol{c}$. We say that simulator S simulates $\boldsymbol{h}$ from $\boldsymbol{c}$ and $\boldsymbol{r}$ if $Pr[S(\boldsymbol{c}, \boldsymbol{r})] = Pr[\boldsymbol{h}]$.*

Importantly, the variables $\boldsymbol{c}$ and $\boldsymbol{h}$ are not necessarily independent, meaning $\Pr[\boldsymbol{h} \mid \boldsymbol{c}]$ could be different from $\Pr[\boldsymbol{h}]$, *i.e.*, their distributions are somehow related. This is central in the definitions of Stateful t–NI and t–SNI, however, we introduce a non-probabilistic way of *modeling* (instead of simulating) the outcome of a computation from a related but different value.

Definition 2 (Modeling Function). *Let $f_H : H \to V$ and $f_C : C \to U$ be deterministic functions. We say that a deterministic function $f_S : U \to V$ is a* modeling function *which models f_H from f_C under deterministic relation $\Psi : H \times C \to \mathbb{B}$ whenever*

$$\forall h \in H, c \in C : \Psi(h, c) \Rightarrow f_S \circ f_C(c) = f_H(h). \tag{1}$$

Definition 2 is strong: whenever modeling function f_S models f_H then it also simulates it, captured by Lemma 1.

Lemma 1 (Modeling Functions Simulate). *Let $\boldsymbol{h}$ and $\boldsymbol{c}$ be sets of possibly dependent random variables with deterministic function f_H computing a set of dependent random variables $\boldsymbol{v}$ and function f_C computing dependent random variables $\boldsymbol{u}$. If function f_S models f_H from f_C whenever $\Psi(\boldsymbol{h}, \boldsymbol{c})$ then it also simulates $\boldsymbol{v}$:*

$$Pr[f_S \circ f_C(\boldsymbol{c}) \mid \Psi(\boldsymbol{h}, \boldsymbol{c}) = \top] = Pr[\boldsymbol{h} \mid \Psi(\boldsymbol{h}, \boldsymbol{c}) = \top]. \tag{2}$$

Stateful t–(S)NI requires a probabilistic simulator to simulate observations on leakage or outputs shares independently of secrets and a function modeling public outputs from public inputs.

Definition 3 (Stateful t–(S)NI [7–9]). *Gadget $G(\overline{x}, \boldsymbol{r}, \boldsymbol{p}) = (\overline{\boldsymbol{y}}, \boldsymbol{o}, \mathcal{L})$ is Stateful t–(S)NI if for every set $\mathbf{e} \subseteq \mathcal{L} \cup \overline{\boldsymbol{y}}$, with $|\mathbf{e}| \leq t$, there exists a subset of input shares $\mathbf{s} \subseteq \overline{x}$, with $\forall i : |\mathbf{s} \cap \overline{x_i}| \leq t' \leq t$, a set of uniformly random variables $\boldsymbol{r}'$, a modeling function $F : Dom(\boldsymbol{p}) \to Dom(\boldsymbol{o})$* modeling *public outputs $\boldsymbol{o}$ from public inputs $\boldsymbol{p}$ and a simulator $S : Dom(\mathbf{s}, \boldsymbol{r}', \boldsymbol{p}) \to Dom(\mathbf{e})$* simulating *observations $\mathbf{e}$ from a subset of shares $\mathbf{s}$, random $\boldsymbol{r}'$ and public inputs $\boldsymbol{p}$ such that $Pr[S(\mathbf{s}, \boldsymbol{r}', \boldsymbol{p}), F(\boldsymbol{p})] = Pr[\mathbf{e}, \boldsymbol{o}]$ and $\boldsymbol{o} = F(\boldsymbol{p})$. For t–NI $t' = |\mathbf{e}|$ while the stricter t–SNI notion requires $t' = |\mathbf{e} \backslash \overline{\boldsymbol{y}}|$.*

The tool scVerif allows proving software gadgets secure under these notions for custom definitions of the observable leakage behavior $\mathcal{L}$ [9]. However, to prove security with respect to gate-level leakage, the provided model of $\mathcal{L}$ must capture absolutely all gate-level leakages of the CPU executing a gadget. In case a gate-level leakage is modeled incorrectly, the tool could assert security although the gadget can be broken with less than t observations at the gate-level. As analyzed by Balasch *et al.* [5], such leakage may halve the security order, *i.e.*, $t' = \frac{t}{2}$. Even worse, Gigerl *et al.* report protection losses that scale with the number of processor pipeline stages [26]. Our work mitigates such losses by verifying that all gate-level leakages are modeled.

3 HARDWARE-SOFTWARE CONTRACTS

A contract defines the instruction semantics and exposed side-channel information of a processor from the perspective of a software developer, *i.e.*, which data is leaked via power side-channels when an instruction is executed in conjunction with the semantic of the instruction. Contracts must specify correct instruction semantics to be able to express accurate data leakage. Besides the instruction perspective contracts allow to execute and thereby model entire programs. In practice, a contract is a user-supplied text file containing specifications of instructions written in Genoa.

In Sections 3.1 and 3.2, we describe how to build a contract which completely captures the leakage exposed by every single gate of a processor. In Section 3.3 we define how to verify the security of masked software against the model specified in a contract. We then turn towards the question of model completeness: In Section 3.4 we define *compliance*, a property which connects gate-level leakage of a processor to the leakage model specified in a contract. Finally, we prove E2E security by showing that if a processor's hardware complies with a contract, the contract models all gate-level leakages. Proving the security of a program against the leakage model specified in the contract implies that the same order of security is achieved when executed on real, compliant hardware. Section 4 introduces a way to check compliance of processors and contracts.

3.1 Expressing Contracts in Genoa

Genoa extends Sail by a dedicated `leak` statement to express that specific values are observable through a side-channel. For example, a statement of the form `leak(val1, val2)` indicates that the processor may leak any combination of the source operands, *i.e.*, any value that can be computed using these operands. Users are free to specify more fine-grained leakage using concrete functions, e.g., the Hamming-Distance. Barthe *et al.* applied this concept in [9] to a custom DSL but leave the error-prone and time-consuming task of modeling semantic and leakage to the user. As we will show, modeling leakage in a contract becomes as easy as adding few `leak` statements to one of the many existing Sail models for RISC-V, ARM, etc. (Genoa supports all Sail models), providing an interface to our tool and applying it to check for modeling gaps (more on this in Section 5). Parts of the IBEX contract are shown in Listings 1 to 6.

The Sail manual [2] and the work of Armstrong *et al.* [4] provide in-depth explanations of the syntax, we give a brief overview. In Listing 1 we define the architectural state of a processor, consisting of 32-bit registers which are declared as global variables.

Listing 1: Contract model of state defined in Genoa.

```
// adopted from RISCV Sail Model, see license in Listing 6
register PC : bits(32)
register nextPC : bits(32)
register x1 : bits(32) ...
// shadow registers
register rf_pA : bits(32) // register file read port A
register rf_pB : bits(32) // register file read port B
register mem_last_addr : bits(32) // address of last access
register mem_last_read : bits(32) // data from last instr.
```

Listing 2: Model of instruction-step χ defined in Genoa.

```
// adopted from RISCV Sail Model, see license in Listing 6
function step_ibex (op : bits(32)) -> bool = {
  nextPC = PC + 4;
  let instruction = encdec(op);
  let ret = execute(instruction);
  tick_pc();
  match ret { RETIRE_SUCCESS => return true,
              RETIRE_FAIL => return false}}
```

Additional *shadow registers* are introduced to model leakage which arises from microarchitectural state in hardware. For example, `rf_pA` is used to remember the value last read from the register file but is not used in the specification of instruction semantics. Its value is maintained in the model and later on leaked in `leak` statements to model leakage of instructions accessing the register file since such leakage involves the value of the register read last [45]. Every contract must specify a step function defining how a single instruction is executed. For IBEX, `step_ibex` shown in Listing 2 decodes the machine code instruction (`encdec`) provided as parameter `op` and returns whether the instruction executes (`execute`) successfully. Both `encdec` and `execute` are `scattered` into multiple clauses which describe the decoding, respectively execution, for a few instructions loosely belonging to a category. Each category is represented by a datatype `ast`, e.g., `RTYPE` for instructions operating on three ISA registers, represented by three indices for destination and two operands, as well as another datatype `rop` for different operations. Listing 3 shows the model of `RTYPE` instructions; `encdec` maps between instruction bits and `ast` representations using conditional pattern matching. In line 6 `rs1` represents the index bits of the first source register. The instruction semantic and leakage is specified in `execute`, `X(rs1)` returns the value of the register addressed by `rs1`. Leakage which is common across multiple instruction categories is exposed with a call to function `common_leakage` (we defer the descriptions to Section 5.2, Listing 4). The semantics of the different operations (add, signed less than, etc.) is defined in the `match` statement. Function `overwrite_leakage` specifies transition leakage emitted while writing the result to the destination register. In summary, Genoa allows designers to quickly construct and adjust contracts, while the human-readable specification supports the systematic development of side-channel protected software.

3.2 Contract Formalization

In the previous section we explained how to express contracts in the Genoa DSL. We now describe Genoa's profound formal semantics which is the basis for security verification and compliance checking.

Listing 3: Contract model of R-type instructions in Genoa.

```
// adopted from RISCV Sail Model, see license in Listing 6
type regidx = bits(5) // index of register 0b00001 = x1
enum rop = {RISCV_ADD, RISCV_SUB, RISCV_SLL, RISCV_SLT,
     ↪ RISCV_SLTU, RISCV_XOR, RISCV_SRL, RISCV_SRA,
     ↪ RISCV_OR, RISCV_AND}
union clause ast = RTYPE : (regidx, regidx, regidx, rop),
mapping clause encdec = RTYPE(rs2, rs1, rd, RISCV_ADD)
  <-> 0b0000000 @ rs2 @ rs1 @ 0b000 @ rd @ 0b0110011
  if (rs1[4] == bitzero) & (rs2[4] == bitzero) & (rd[4] ==
       ↪ bitzero)
mapping clause encdec = RTYPE(rs2, rs1, rd, RISCV_SLT)
  <-> 0b0000000 @ rs2 @ rs1 @ 0b010 @ rd @ 0b0110011
  if (rs1[4] == bitzero) & (rs2[4] == bitzero) & (rd[4] ==
       ↪ bitzero)
...
function clause execute (RTYPE(rs2, rs1, rd, op)) = {
  let rs1_val = X(rs1);
  let rs2_val = X(rs2);
  common_leakage(rs1_val, rs2_val);
  let result : bits(32) = match op {
    RISCV_ADD => rs1_val + rs2_val,
    RISCV_SLT => EXTZ(bool_to_bits(rs1_val <_s rs2_val)),
    ... };
  overwrite_leakage(rd, result);
  X(rd) = result;
  return RETIRE_SUCCESS}
```

The small-steps semantics of Genoa are defined as a reduction

$$(\delta, s, \mathcal{L}) \mapsto (\delta', s', \mathcal{L}').$$

δ is the context containing the definition of functions and the values of local and global variables, s is a sequence of statements and $\mathcal{L}$ is the leakage exposed during execution. After the execution of one Genoa statement (not to be confused with an instruction) δ' is the resulting context, s' are the remaining statements and $\mathcal{L}' \supseteq \mathcal{L}$ is the resulting leakage. Leakage cannot be erased. All statements except `leak` do not add leakage and their transformation rules stay the same as in Sail [3]. A `leak` statement appends its operands $v_1, \ldots, v_n$ to the execution leakage ($\cdot;\cdot$ is a sequence of statements):

$$(\delta, \texttt{leak}\,(v_1, \ldots, v_n)\,; s, \mathcal{L}) \mapsto (\delta, s, \mathcal{L} \cup \{v_1(\sigma^c) || \ldots || v_n(\sigma^c)\}).$$

A `leak` statement may expose multiple values, which allows abstracting away from particular assumptions such as Hamming-Distance leakage, as processors are allowed to leak any combination of the values exposed by a `leak`. While Genoa does not feature a construct to sample random values, sampling can be mimicked by reading from a dedicated state region containing randomness.

The behavior of a program is defined by user-supplied executions semantics which are specified in the contract. The contract specification written in Genoa thus defines the context δ for small-step execution and, as for hardware, the contract state $\sigma^c \in \mathbb{B}^{|V^c|}$ denotes the values of variables $v^c \in V^c$, further on referred to as locations. Based on these definitions, we can now define the semantics for the execution of an entire instruction, denoted by the step function χ, starting in state σ_i^c and returning the state σ_{i+1}^c and a set of side-channel leakages $\mathcal{L}_i^c$ of executing the i^{th} instruction:

$$\chi(\sigma_i^c) = (\sigma_{i+1}^c, \mathcal{L}_i^c).$$

The instruction to be executed is determined by the state σ_i^c itself, e.g., by the value of the program counter. The execution of an instruction corresponds to the many-steps evaluation of the instruction-steps function χ using the small-steps semantics described before. χ is part of the contract (for IBEX `step_ibex`) and supplied by the user; to simplify our tool state σ_i^c is implicitly passed while the instruction to be executed is passed explicitly. A step can either fail or succeed, indicated by a Boolean flag. The criteria for failing the execution is governed by user-defined assumptions. For IBEX these prohibit illegal instructions, accesses of non-existent registers or unaligned memory accesses. In the following, we depict execution of an entire instruction in the contract with

$$\sigma_i^c \xrightarrow{\mathcal{L}_i^c} \sigma_{i+1}^c.$$

Finally, we define the execution of programs in a contract. Contract $[\![\cdot]\!]^c$ models the execution of program P starting in initial state σ_0^c and resulting in state σ_n^c while producing the accumulated observable side-channel information $\mathcal{L}_{0,n}^c = \bigcup_{i=0}^{n-1} \mathcal{L}_i^c$, *i.e.*,

$$[\![\mathsf{P}]\!]^c \left(\sigma_0^c\right) = \left(\sigma_n^c, \mathcal{L}_{0,n}^c\right).$$

3.3 Software Security

In this section we link security of abstract gadgets to the execution in a contract or on hardware, *i.e.*, we define security w.r.t. the leakages specified in a contract or gate-level leakage of a processor.

Gadgets have as inputs and outputs either shares, random or public values, which are linked to the definition of Stateful t–(S)NI (Definition 3). However, when the implementation of a gadget is executed within a contract or hardware then the gadget's inputs are located in the state σ with an implementation-specific placement, e.g., shares could be in registers or memory. We introduce policies π to translate between the structured in- and outputs of a gadget and the states in a contract, respectively hardware. Input policy $\pi_{\text{in}} : (\overline{\boldsymbol{x}}, \boldsymbol{r}, \boldsymbol{p}) \leftrightarrow \sigma_0$ constructs a state given values of variables for input shares, random and public variables but also the converse; extracting the values of gadget inputs given a state. In practice, such a policy is an annotation which defines where shares, random and public (initial) values are located w.r.t. locations of the state. Similar, output policy $\pi_{\text{out}} : (\overline{\boldsymbol{y}}, \boldsymbol{o}) \leftrightarrow \sigma_n$ maps between the values of public outputs and output shares of a gadget and the state σ_n resulting from an execution. Since Stateful t–(S)NI is defined for random variables let $\boldsymbol{\sigma}$ denote the random variable for states. Using policies we can link Stateful t–(S)NI security of gadgets to the execution of their concrete implementation P within big-steps semantics $[\![\cdot]\!]$ representing either the contract or hardware:

Definition 4 (t–(S)NI of $[\![\mathsf{P}]\!]$ under π_{in}, π_{out}). *An implementation* P *of gadget* G $(\overline{\boldsymbol{x}}, \boldsymbol{r}, \boldsymbol{p})$ *is* t–(S)NI *secure w.r.t. semantic* $[\![\cdot]\!]$ *and placement policies* π_{in}, π_{out} *if the gadget* G′ $(\overline{\boldsymbol{x}}, \boldsymbol{r}, \boldsymbol{p}) = (\overline{\boldsymbol{y}}, \boldsymbol{o}, \mathcal{L}_{0,n})$ *is* t–(S)NI *according to Definition 3, where the inputs of the gadget correspond to the starting states* $\boldsymbol{\sigma}_0 = \pi_{in}(\overline{\boldsymbol{x}}, \boldsymbol{r}, \boldsymbol{p})$ *while leakages and gadget outputs correspond to the random variables* $\pi_{out}(\overline{\boldsymbol{y}}, \boldsymbol{o}) = \boldsymbol{\sigma}_n$ *resulting from execution* $[\![\mathsf{P}]\!](\boldsymbol{\sigma}_0) = (\boldsymbol{\sigma}_n, \boldsymbol{\mathcal{L}}_{0,n})$.

The actual verification of security for software implementations follows the same principles as outlined by Barthe *et al.* [9], which also describes representation of policies. However, the dependent type system of Genoa enables new approaches to verification of masked conversion functions and arithmetic masking in rings with prime moduli for security orders $t > 1$. We leave the development of software verification tools to dedicated future work.

$$\begin{array}{lccccccc} [\![\mathsf{P}]\!]^c\left(\sigma_0^c\right): & \sigma_0^c & \xrightarrow{\mathcal{L}_0^c} & \sigma_1^c & \xrightarrow{\mathcal{L}_1^c} & \cdots & \xrightarrow{\mathcal{L}_{n-1}^c} & \sigma_n^c \\ & \| \mathcal{M} & & & & & & \| \mathcal{M} \\ [\![\mathsf{P}]\!]^h\left(\sigma_0^h\right): & \sigma_0^h & \xrightarrow{\mathcal{L}_0^h} & \sigma_1^h & \xrightarrow{\mathcal{L}_1^h} & \cdots & \xrightarrow{\mathcal{L}_{m-1}^h} & \sigma_m^h \end{array}$$

Figure 1: Compliance for a full program execution

3.4 Hardware Compliance With a Contract

We now turn towards the question of model completeness and define *compliance with a contract*, a formal property expressing that the results and leakages from execution on a CPU are modeled by a contract according to Definition 2. This property is verified in Section 4 and ensures, as we prove in Section 3.5, that any program that is Stateful t–(S)NI secure in a contract mitigates all order-reducing leakage caused by the gate-level leakage of a processor.

A program P executed in initial hardware state σ_0^h leads to leakages $\mathcal{L}_{0,m}^h$ and final state σ_m^h when executed on processor $[\![\cdot]\!]^h$:

$$[\![\mathsf{P}]\!]^h \left(\sigma_0^h\right) = \left(\sigma_m^h, \mathcal{L}_{0,m}^h\right).$$

In contrast to contracts the execution proceeds in clock-cycles instead of instruction-steps, *i.e.*, one step in hardware corresponds to one clock cycle as defined in Section 2.1:

$$\sigma_j^h \xrightarrow{\mathcal{L}_j^h} \sigma_{j+1}^h.$$

Compliance expresses the property that all leakage and all outputs of hardware execution $[\![\cdot]\!]^h$ can be modeled (according to Definition 2) from execution in a contract $[\![\cdot]\!]^c$ as long as the starting states are similar, *i.e.*, execute the same program under equivalent inputs, depicted in Figure 1.

In Definition 5 we introduce a Boolean relation between hardware state σ^h and contract state σ^c expressing that the values contained at a specific location v^h in the hardware can be modeled from a location v^c in the contract, e.g., register x1 models its counterpart in hardware. Which contract locations model some hardware location is defined in the simulation mapping $\mathcal{M}$ provided by users alongside every contract and checked by our tool. The mapping specifies for all registers in the hardware (including finite state machines, decode stages, etc.) a location in the contract modeling the hardware location. To ease notation assume there are contract locations $v_0^c, v_1^c \in V^c$ which are constant zero, respectively one, and later on used to express constraints on hardware execution.

Definition 5 (Similar states: $\sigma^h \simeq_{\mathcal{M}} \sigma^c$). *Two states* σ^h *and* σ^c*, with respective locations* V^h *and* V^c *are* similar *under simulation mapping* $\mathcal{M} \subseteq V^h \times V^c$*, written* $\sigma^h \simeq_{\mathcal{M}} \sigma^c$*, if and only if*

$$\bigwedge_{(v^h, v^c) \in \mathcal{M}} v^h(\sigma^h) = v^c(\sigma^c) \tag{3}$$

Simulation mapping $\mathcal{M}$ *is said to be* complete *if for all hardware locations* V^h *a mapping is defined.*

We now give the definition of compliance, ensuring that semantic and leakage of execution of hardware is correctly modeled:

Definition 6 (Compliance: $[\![\cdot]\!]^h \vdash_{\mathcal{M}} [\![\cdot]\!]^c$). *A hardware implementation is compliant with a contract under simulation mapping $\mathcal{M}$ if for every program* P *and starting hardware and contract states σ_0^h and σ_0^c, the program executions*

$$[\![\mathsf{P}]\!]^c(\sigma_0^c) = \left(\sigma_n^c, \mathcal{L}_{0,n}^c\right) \quad \text{and} \quad [\![\mathsf{P}]\!]^h\left(\sigma_0^h\right) = \left(\sigma_m^h, \mathcal{L}_{0,m}^h\right)$$

fulfill the following conditions:

1. ***States remain similar:*** *Whenever σ_0^h and σ_0^c are similar under $\mathcal{M}$, so are resulting states σ_m^h and σ_n^c:*
$$\forall \sigma_0^h, \sigma_0^c : \sigma_0^h \simeq_{\mathcal{M}} \sigma_0^c \Rightarrow \sigma_m^h \simeq_{\mathcal{M}} \sigma_n^c.$$
2. ***Leaks are modeled:*** *For every leak $\lambda_g(\sigma_{j-1}^h, \sigma_j^h) \in \mathcal{L}_{0,m}^h$ observable in hardware, there exists a leak $\lambda(\sigma_i^c) \in \mathcal{L}_{0,n}^c$ in the contract and a function $f_\lambda : Dom(\lambda) \to Dom(\lambda_g)$ that models λ_g from λ under relation $\sigma_0^h \simeq_{\mathcal{M}} \sigma_0^c$ according to Definition 2:*
$$\forall \sigma_0^h, \sigma_0^c : \sigma_0^h \simeq_{\mathcal{M}} \sigma_0^c \Rightarrow f_\lambda \circ \lambda\left(\sigma_i^c\right) = \lambda_g\left(\sigma_{j-1}^h, \sigma_j^h\right).$$

The notion of similar states allows to express a key ingredient for the relational definition of compliance: if execution in a contract and hardware start in a similar state, then execution must end in similar states such that the hardware execution's results can be modeled according to the simulation mapping (Clause 1 of Definition 6). Further, the second part of compliance expresses that every gate-level leak observable during execution in hardware must be modeled by a single, fixed leak observable during execution in the contract (Clause 2 of Definition 6). Combined, this guarantees that software which is Stateful t–(S)NI secure when executed in the contract, is necessarily Stateful t–(S)NI when executed on compliant hardware.

3.5 End-to-end security

It remains to prove our E2E security claim: any implementation P of gadget G that is Stateful t–(S)NI w.r.t. the leakages of a contract must be Stateful t–(S)NI w.r.t. all gate-level leakage when executed on any compliant hardware and as such its security order cannot be decreased by leakage of the processor.

However, E2E security is claimed for the same software P implementing some gadget G and running on a processor and in the contract, i.e., both executions use the same structured inputs and outputs. Since the states in hardware and contract may have different structures we introduce a definition to ensure that the placement of inputs in hardware $\pi_{\text{in}}^h, \pi_{\text{out}}^h$ is similar to the ones $\pi_{\text{in}}^c, \pi_{\text{out}}^c$ for which t–(S)NI was shown in the contract. A hardware policy can be derived from a contract policy by substituting the locations which define where a value resides in the state according to the simulation mapping.

Definition 7 (Similar Policy ($\pi^h \triangleq_{\mathcal{M}} \pi^c$)). *Let contract policy $\pi^c : (\boldsymbol{d}_1, \ldots, \boldsymbol{d}_n) \leftrightarrow \sigma^c$ link sets of values $\boldsymbol{d}_1, \ldots, \boldsymbol{d}_n$ to contract state σ^c. Hardware policy $\pi^h : (\boldsymbol{d}_1, \ldots, \boldsymbol{d}_n) \leftrightarrow \sigma^h$ is similar to π^c, denoted $\pi^h \triangleq_{\mathcal{M}} \pi^c$ if any pair of contract and hardware states constructed from the same sets of values are similar under mapping $\mathcal{M}$:*

$$\forall \sigma^h = \pi^h(\boldsymbol{d}_1, \ldots, \boldsymbol{d}_n), \sigma^c = \pi^c(\boldsymbol{d}_1, \ldots, \boldsymbol{d}_n) \ : \ \sigma^h \simeq_{\mathcal{M}} \sigma^c.$$

Instead of proving the security reduction for t–(S)NI directly we prove a general *model reduction*: any observations made by an adversary interacting with hardware may be *modeled* with a contract the hardware complies with instead. We emphasize the difference: t–(S)NI requires the existence of a *simulation procedure* whereas compliance guarantees the existence of a (stronger) *modeling function* easing the subsequent security reduction.

Theorem 2 (Model Reduction). *Let* P *be a program, and the gadgets $\mathrm{G}^c(\overline{\boldsymbol{x}}, \boldsymbol{r}, \boldsymbol{p}) = \left(\overline{\boldsymbol{y}}^c, \boldsymbol{o}^c, \mathcal{L}_{0,n}^c\right)$ and $\mathrm{G}^h(\overline{\boldsymbol{x}}, \boldsymbol{r}, \boldsymbol{p}) = \left(\overline{\boldsymbol{y}}^h, \boldsymbol{o}^h, \mathcal{L}_{0,m}^h\right)$ correspond to the program executions $[\![\mathsf{P}]\!]^c(\boldsymbol{\sigma}_0^c) = (\boldsymbol{\sigma}_n^c, \mathcal{L}_{0,n}^c)$, respectively $[\![\mathsf{P}]\!]^h(\boldsymbol{\sigma}_0^h) = (\boldsymbol{\sigma}_m^h, \mathcal{L}_{0,m}^h)$, under policies π_{in}^h and π_{out}^h, respectively π_{in}^c and π_{out}^c, with $\boldsymbol{\sigma}_0^c = \pi_{in}^c(\overline{\boldsymbol{x}}, \boldsymbol{r}, \boldsymbol{p})$, $\boldsymbol{\sigma}_0^h = \pi_{in}^h(\overline{\boldsymbol{x}}, \boldsymbol{r}, \boldsymbol{p})$, $\boldsymbol{\sigma}_n^c = \pi_{out}^c(\overline{\boldsymbol{y}}^c, \boldsymbol{o}^c)$, and $\boldsymbol{\sigma}_m^h = \pi_{out}^h\left(\overline{\boldsymbol{y}}^h, \boldsymbol{o}^h\right)$. Furthermore, let $[\![\cdot]\!]^h \vdash_{\mathcal{M}} [\![\cdot]\!]^c$, $\pi_{in}^h \triangleq_{\mathcal{M}} \pi_{in}^c$ and $\pi_{out}^h \triangleq_{\mathcal{M}} \pi_{out}^c$ under complete mapping $\mathcal{M}$. For every set of observations in hardware on $\overline{\boldsymbol{y}}^h$ or $\boldsymbol{o}^h$ there is an equally sized set of observations in the contract on $\overline{\boldsymbol{y}}^c$ or $\boldsymbol{o}^c$ which allows to model the observations under the identity function:*

$$\forall \mathbf{e}_{\overline{\boldsymbol{y}}}^h \subseteq \overline{\boldsymbol{y}}^h \ \exists \ \mathbf{e}_{\overline{\boldsymbol{y}}}^c \subseteq \overline{\boldsymbol{y}}^c \ : \ \mathbf{e}_{\overline{\boldsymbol{y}}}^h = \mathbf{e}_{\overline{\boldsymbol{y}}}^c, \tag{4}$$

$$\forall \mathbf{e}_{\boldsymbol{o}}^h \subseteq \boldsymbol{o}^h \ \exists \ \mathbf{e}_{\boldsymbol{o}}^c \subseteq \boldsymbol{o}^c \ : \ \mathbf{e}_{\boldsymbol{o}}^h = \mathbf{e}_{\boldsymbol{o}}^c. \tag{5}$$

In addition, for every set of observations in hardware on $\mathcal{L}_{0,m}^h$, a modeling function $T^{\mathcal{L}}$ and a (potentially smaller) set of observations in the contract on $\mathcal{L}_{0,n}^c$ allow to model the observations in hardware:

$$\forall \mathbf{e}_{\mathcal{L}}^h \subseteq \mathcal{L}_{0,m}^h \ \exists \, \mathbf{e}_{\mathcal{L}}^c \subseteq \mathcal{L}_{0,n}^c : \left|\mathbf{e}_{\mathcal{L}}^c\right| \leq \left|\mathbf{e}_{\mathcal{L}}^h\right| \wedge \mathbf{e}_{\mathcal{L}}^h = T^{\mathcal{L}}\left(\mathbf{e}_{\mathcal{L}}^c\right). \tag{6}$$

Proof. The gadgets G^c and G^h operate on equally distributed inputs and the policies for hardware are similar, thus for every initial state σ_0^h there must be a starting state σ_0^c similar under mapping $\mathcal{M}$, *i.e.*, $\sigma_0^h \simeq_{\mathcal{M}} \sigma_0^c$. Since hardware is compliant with the contract, the resulting states are similar as well, *i.e.*, $\sigma_m^h \simeq_{\mathcal{M}} \sigma_n^c$ and since every observation in $\mathbf{e}_{\boldsymbol{o}}^h$, respectively $\mathbf{e}_{\overline{\boldsymbol{y}}}^h$, is an observation on the value of a location in σ_m^h it follows directly that there exists a single location in the contract $\mathbf{e}_{\boldsymbol{o}}^c$, respectively $\mathbf{e}_{\overline{\boldsymbol{y}}}^h$, according to the mapping $\mathcal{M}$ which models the observation, fulfilling (4) and (5). From Lemma 1 and Clause 2 of Definition 6 it follows that every observation $\lambda_g(\boldsymbol{\sigma}_{j-1}^h, \boldsymbol{\sigma}_j^h) \in \mathbf{e}_{\mathcal{L}}^h$ can be modeled from some contract leak $\lambda(\boldsymbol{\sigma}_i^c) \in \mathcal{L}_{0,n}^c$ using f_λ as modeling function. Grouping the necessary $\lambda(\boldsymbol{\sigma}_i^c)$ as the set of random variables $\mathbf{e}_{\mathcal{L}}^c$, results in $\left|\mathbf{e}_{\mathcal{L}}^c\right| \leq \left|\mathbf{e}_{\mathcal{L}}^h\right|$, and defining $T^{\mathcal{L}}$ as the set of respective f_λ implies (6), completing the proof. □

From Theorem 3, we derive simulatability of mixed observations in Corollary 3. Furthermore, the reduction from Stateful t–(S)NI in hardware to Stateful t–(S)NI in contract stated in Corollary 4 is a direct consequence of Theorem 3 and Corollary 3.

Corollary 3 (Mixed observations). *Let the setting be as in Theorem 2. Every set of mixed observations on leakage and shared outputs $\mathbf{e}_{\mathcal{L},\overline{\boldsymbol{y}}}^h \subseteq \mathcal{L}^h \cup \overline{\boldsymbol{y}}^h$, can be modeled from some an equally sized set $\mathbf{e}_{\mathcal{L},\overline{\boldsymbol{y}}}^c \subseteq \mathcal{L}^c \cup \overline{\boldsymbol{y}}^c$ by some modeling function $T^{\mathcal{L},\overline{\boldsymbol{y}}}$.*

Corollary 4 (End-to-End Security). *Let the setting be as in Theorem 2. If gadget* G^c *is* t*–(S)NI then gadget* G^h *is also* t*–(S)NI since there exist simulators* $T^{\mathcal{L},\overline{y}} \circ S$ *and* F *which simulate the outputs of* G^h *according to Definition 3.*

This proof is valid for higher-order masking, *i.e.*, $t \geq 1$, as *each* of the t hardware observations in $\mathbf{e}^h$ can be simulated from *one* observation in $\mathbf{e}^c$ in the contract. The presented model reduction can be of help in proving the preservation of other security notions like PINI [12], threshold implementations [43] or probing security [31].

4 VERIFYING HARDWARE COMPLIANCE

Whereas Section 3 introduces *contracts* and what it means for hardware to be *compliant*, this section presents a method to actually check hardware compliance for a given processor, i.e., that the leakage of each of gate in the CPU netlist (i.e. the synthesized CPU design) and the semantic of instruction is correctly specified in the corresponding contract. The method is broken down into verification steps. Each step checks if the processor satisfies some part of Definition 6. First, we check whether similar hardware and contract states stay similar after executing an instruction according to Clause 1. Then, we check that each hardware leak can be modeled from a single leak emitted in the contract, according to Clause 2.

4.1 Verification Concept

In this section, we first suggest that it is possible to prove a processor compliant without looking at full program executions. We argue that looking at all possible single instruction executions is sufficient to form an inductive argument of compliance. Next, we give an overview of the individual verification steps needed to verify that a processor is compliant with a given contract. Finally, we present a method for indirectly proving the existence of modeling functions. This method is the backbone of the verification procedure and relies on encoding constraints into SMT formulas and checking their satisfiability with an SMT solver.

Single instructions. Checking compliance for all programs and pairs of processors and contracts using SMT solvers is computationally intractable. Instead we prove compliance inductively by showing that Definition 6 holds for all possible executions of a single instruction. Consequently, we require compliant processors to fulfill the outlined properties at the start and end of each instruction, as shown in Figure 2. Starting with similar states σ_j^h and σ_i^c, the hardware executes k clock cycles and the contract executes one instruction step. The executions produce leaks $\boldsymbol{\mathcal{L}}_{j,j+k}^h$, respectively $\boldsymbol{\mathcal{L}}_i^c$, and state σ_{j+k}^h, respectively σ_{i+1}^c, for which we need to show:

(a) *States remain similar:* The states σ_{j+k}^h and σ_{i+1}^c are similar under $\mathcal{M}$ (marked red), *i.e.*, every location in σ_{j+k}^h is equal to the corresponding location in σ_{i+1}^c.
(b) *Leaks are modeled:* Every leak $\lambda_g(\sigma_{j+l-1}^h, \sigma_{j+l}^h) \in \boldsymbol{\mathcal{L}}_{j,j+k}^h$ produced by the processor (marked red) must be modeled by a single leak $\lambda(\sigma_i^c) \in \boldsymbol{\mathcal{L}}_i^c$ emitted in the contract execution.

These conditions are inductive and much stricter than the corresponding clauses in Definition 6, *i.e.*, if the execution of a single instruction in an arbitrary valid starting state maintains the compliance properties, the processor complies with the contract for all possible program executions.

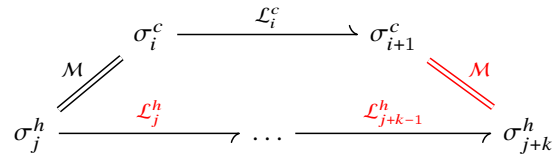

Figure 2: Compliance for single instruction execution

As seen in Figure 2, the hardware might require multiple clock cycles to execute an instruction while the contract always takes only one step. For the purposes of this paper, we define the starting point of an instruction as the moment it becomes in-flight, *i.e.*, it reaches the decode stage, and its end point when it retires, *i.e.*, the writeback completes. Therefore, we look at every possible instruction duration k on that particular processor. Concurrent execution of instructions in the pipeline complicates this approach. For simple pipelines this is not an issue because the fetch stage does not operate with security-critical data, and the writeback stage can be made the synchronisation point for the induction, instead of its full retirement. In more complex pipelines, the methods described in this paper require checking leakage produced by hardware components directly influenced by the instruction bits.

Verifying that states remain similar. As the very first step in the verification procedure, we show that the hardware and contract states are similar throughout the whole execution. That is, we show that if the relation $\sigma_j^h \simeq_{\mathcal{M}} \sigma_i^c$ holds, the relation $\sigma_{j+k}^h \simeq_{\mathcal{M}} \sigma_{i+1}^c$ must hold after the execution of a k-cycle instruction *i.e.*, $\sigma_{j+k}^h = \chi^k(\sigma_j^h)$ and $\sigma_{i+i}^c = \chi(\sigma_i^c)$, no matter what the starting states were. This is essentially a full-fledged functional equivalence proof between the hardware and the contract. If this check succeeds, we have shown that the processor satisfies Clause 1 of Definition 6 because $\simeq_{\mathcal{M}}$ is conserved over the execution of an instruction. Section 4.3 formalizes the verification step and gives a verification method.

Finding modeling functions for gates. Before verifying that leaks are modeled, we require an intermediate verification step that provides information about the old values of gates. This constrains the old values of each gate g, therefore implicitly constraining the possible values of the corresponding leak λ_g. Otherwise, the old value could directly leak secrets, trivially breaking leakage modeling. For every gate $g \in G$ in the hardware, we show that g can be modeled by some function $f_g : \mathbb{B}^n \to \mathbb{B}$ that only uses a (small) subset of contract state bits $\theta_g : \mathbb{B}^{|V^c|} \to \mathbb{B}^n$, *i.e.*,

$$\exists f_g : \forall \sigma_j^h, \sigma_i^c : \sigma_j^h \simeq_{\mathcal{M}} \sigma_i^c \Rightarrow f_g \circ \theta_g \circ \chi\left(\sigma_i^c\right) = g \circ \chi^{k-1}\left(\sigma_j^h\right). \quad (7)$$

Ideally, we want to prove the existence of a modeling function that uses as little contract state information θ_g as possible. Section 4.4 gives exact definitions of the verification checks and the greedy minimization procedure for θ_g.

Verifying that leaks are modeled. In this verification step, we check whether the hardware leakage is properly modeled from contract leakage for any possible instruction execution, starting in any pair of similar states $\sigma_j^h \simeq_{\mathcal{M}} \sigma_i^c$. If the check succeeds, the proper modeling throughout any program execution is implied by

composition of single instructions. Because we consider transition leakage, we constrain the possible values of gates at the end of the previous instruction. As mentioned before, we use the existence of a modeling function f_g according to (7) from the previous step.

The hardware leak $\lambda_g(\sigma^h_{j+l-1}, \sigma^h_{j+l}) = g(\sigma^h_{j+l-1}) || g(\sigma^h_{j+l})$ contains information about both the old, and the new values of gate g for any clock cycle l of the executed instruction. We analyze each hardware leak function λ_g separately by going through all leakage functions $\lambda : \mathbb{B}^{|V^c|} \to \mathbb{B}^m$ and checking if there is a function $f_\lambda : \mathbb{B}^m \to \mathbb{B}^2$ that models λ_g from λ, whenever states σ^h_j and σ^c_i are similar, the contract leak is emmitted, *i.e.*, $\lambda(\sigma^c_i) \in \mathcal{L}^c_i$, and f_g models the previous value of the gate g from θ_g. Written formally:

$$\begin{aligned} \forall \sigma^h_j, \sigma^c_i : \sigma^h_j \simeq_{\mathcal{M}} \sigma^c_i \wedge f_g \circ \theta_g \circ \chi\left(\sigma^c_i\right) = g \circ \chi^{k-1}\left(\sigma^h_j\right) \wedge \\ \lambda\left(\sigma^c_i\right) \in \mathcal{L}^c_i \Rightarrow\ f_\lambda \circ \lambda\left(\sigma^c_i\right) = \lambda_g\left(\sigma^h_{j+l-1}, \sigma^h_{j+l}\right). \end{aligned} \tag{8}$$

Additionally, we require that for any possible state σ^c_i at least one leak λ fulfills (8), guaranteeing that Clause 2 of the compliance definition is fulfilled. Section 4.5 gives a more detailed description.

Existence of modeling functions. Within the last two verification steps, we check that functions over the hardware state σ^h can be modeled from functions over the contract state σ^c whenever $\sigma^h \simeq_{\mathcal{M}} \sigma^c$. This involves proving the existence of modeling functions from Definition 2. However, automatically finding modeling functions is intractable in general [21]. We circumvent this issue by proving the existence of modeling functions without finding their definitions. Theorem 5 presents the condition we need to check.

Theorem 5 (Existence of Modeling Function). *There exists a modeling function $f : U \to V$ according to Def. 2 if and only if*

$$\begin{aligned} \forall h, h' \in H, c, c' \in C : \\ \Psi(h, c) \wedge \Psi\left(h', c'\right) \wedge f_C(c) = f_C\left(c'\right) \Rightarrow f_H(h) = f_H\left(h'\right). \end{aligned} \tag{9}$$

Proof. We prove the equality of the two statements by showing an implication in both directions. First, we prove that (9) follows from (1). From the functional congruence of f, we have:

$$\forall c, c' \in C : \left(f_C(c) = f_C\left(c'\right)\right) \Rightarrow \left(f \circ f_C(c) = f \circ f_C\left(c'\right)\right).$$

After instantiating the statement (1) separately for the primed and non-primed versions of $h \in H$ and $c \in C$, we get:

$$\begin{aligned} \forall h \in H, c \in C : \Psi(h, c) \Rightarrow f \circ f_C(c) = f_H(h), \\ \forall h' \in H, c' \in C : \Psi\left(h', c'\right) \Rightarrow f \circ f_C\left(c'\right) = f_H\left(h'\right). \end{aligned}$$

We see that if all three premises are fulfilled simultaneously, then also all consequences of the implication must be fulfilled simultaneously. Therefore, we consolidate the left- and right-hand sides. Afterwards, we simplify the right-hand side by substituting $f \circ f_C(c)$ with $f_H(h)$, and respectively $f \circ f_C(c')$ with $f_H(h')$, to get (9).

For the second direction of the proof, we assume (9) and construct f so that it fulfills (1) and is well defined for all $u \in U$. First, we define the subset $\widehat{U} \subseteq U$ of function inputs as

$$\widehat{U} := \{u \mid \exists h \in H, c \in C : u = f_C(c) \wedge \Psi(h, c)\}. \tag{10}$$

For inputs $u \in U \setminus \widehat{U}$, we define $f(u)$ as an arbitrary result $v \in V$. This partial definition trivially fulfills (1). For all other $u \in \widehat{U}$, we define $f(u) := f_H(h)$, for an arbitrary qualified h and c as in (10). We now argue that this portion of f is well defined, because $f_H(h)$ is fixed for u. Consider the case where we can pick two such pairs:

$$\begin{aligned} \exists h, h' \in H, c, c' \in C : u = f_C(c) \wedge \Psi(h, c) \wedge \\ u = f_C\left(c'\right) \wedge \Psi\left(h', c'\right). \end{aligned}$$

Because $f_C(c) = f_C(c') = u$, assumption (9) implies that $f_H(h)$ is unique since we always get $f_H(h') = f_H(h)$. □

The underlying principle behind Theorem 5 can be thought of as partial functional congruence. Plainly speaking, if equal *inputs* $f_C(c)$ and $f_C(c')$ always result in equal *outputs* $f_H(h)$ and $f_H(h')$, then there must also be a function mapping between them. Moreover, Theorem 5 can be translated into the quantifier-free SMT fragment and efficiently checked with modern SMT solvers.

4.2 Verification Prerequisites

Real program execution within a processor is subject to many internal assumptions and restrictions that need to be considered when checking the compliance properties. In particular, we define *normal operating conditions* for the execution of an instruction, as well as constraints related to the mapping $\mathcal{M}$ from Section 3.4.

Normal operating conditions. The hardware of a processor has many input ports and internal registers that are invisible to a software developer and are subject to hidden assumptions under *normal operating conditions*. In this section, we introduce predicates $\phi_\_$ to explicitly represent these assumptions. We use the predicate $\phi_{\text{dev}}(\sigma^h)$ to represent the usual assumptions a software developer might have, such as the processor not getting reset, triggering an interrupt, going into debug mode, or getting memory access errors. Similarly, there are several internal conditions for the processor to fetch, start the execution of, and retire an instruction. We formalize these conditions as $\phi^l_{\text{instr}}(\sigma^h), 0 \leq l < k$ for the l-th cycle in k-cycle instructions and apply them for the intermediate states $\sigma^h_{j+l} = \chi^l(\sigma^h_j)$. Sometimes, a contract is not able to execute an instruction because it violates some sanity conditions such as the instruction not being implemented or triggering a fault. We formalize the condition of the contract successfully retiring an instruction as $\phi_{\text{ret}}(\sigma^c)$. We aggregate these conditions into ϕ_{noc} as

$$\phi_{\text{noc}}\left(\sigma^h_j, \sigma^c_i\right) := \phi_{\text{ret}}\left(\sigma^c_i\right) \wedge \bigwedge_{l=0}^{k-1} \phi_{\text{dev}}\left(\sigma^h_{j+l}\right) \wedge \phi^l_{\text{instr}}\left(\sigma^h_{j+l}\right).$$

There are also some constraints that concern multiple executions of the hardware and contract. For such predicates we write $\phi^*_\_$ instead. We define $\phi^*_{\text{ports}}(\sigma^h, \sigma^{h'})$ as the constraint that certain processor input ports only contain public values. More concretely, for two executions of the hardware, these input ports are required to produce identical values. There are also similar execution-spanning conditions for the contract. For instance, the contract should forbid the program counter from becoming secret dependent. The predicate $\phi^*_{\text{ret}}(\sigma^c, \sigma^{c'})$ expresses these constraints, and is stricter than both $\phi_{\text{ret}}(\sigma^c)$ and $\phi_{\text{ret}}(\sigma^{c'})$ separately. Finally, we extend ϕ_{noc} to

ϕ^*_{noc} over several executions as

$$\phi^*_{\text{noc}}\left(\sigma^h_j, \sigma^{h'}_j, \sigma^c_i, \sigma^{c'}_i\right) := \phi^*_{\text{ret}}\left(\sigma^c_i, \sigma^{c'}_i\right) \wedge \bigwedge_{l=0}^{k-1} \phi^*_{\text{ports}}\left(\sigma^h_{j+l}, \sigma^{h'}_{j+l}\right) \bigwedge_{l=0}^{k-1} \phi_{\text{dev}}\left(\sigma^h_{j+l}\right) \wedge \phi_{\text{dev}}\left(\sigma^{h'}_{j+l}\right) \wedge \phi^l_{\text{instr}}\left(\sigma^h_{j+l}\right) \wedge \phi^l_{\text{instr}}\left(\sigma^{h'}_{j+l}\right).$$

Breaking any of the conditions from ϕ^*_{noc} breaks the guarantees provided in this work. Because these assumptions are instrumental for correct execution, we make sure that the restrictions imposed on the hardware still permit the execution of all instructions defined in the contract. This sanity check confirms that software can still execute within both the hardware and the contract, allowing the implementation of a sensible software verifier.

Applying mappings. As introduced in Section 3.4, hardware and contract states can be similar under a mapping. For expressing that two states σ^h and σ^c are similar under mapping $\mathcal{M}$, *i.e.*, $\sigma^h \simeq_{\mathcal{M}} \sigma^c$, we use the predicate $\phi^{\mathcal{M}}_{\text{rel}}(\sigma^h, \sigma^c)$ as defined in (3). Conversely, we also require a predicate expressing that all registers, which are not in the mapping $\mathcal{M}$, are equivalent across hardware executions of the same program. We specify this property of two hardware states σ^h and $\sigma^{h'}$ for the mapping $\mathcal{M}$ and locations V^h as

$$\phi^{\mathcal{M}*}_{\text{pub}}\left(\sigma^h, \sigma^{h'}\right) := \bigwedge_{v^h \in V^h, \nexists (v^h, v^c) \in \mathcal{M}} v^h\left(\sigma^h\right) = v^h\left(\sigma^{h'}\right).$$

4.3 Verifying that States Remain Similar

As introduced in Section 4.1, we verify that the hardware and contract states resulting from the execution of a program are similar by showing similarity after every instruction. Our inductive argument assumes that the states σ^h_j and σ^c_i are similar at the start of an instruction, and proves that the states σ^h_{j+k} and σ^h_{i+1} are also similar after the k-cycle instruction terminates. This is a straightforward functional equivalence check under the assumption that ϕ_{noc} holds.

Proposition 1. *Let σ^h_j be a hardware state and σ^c_i the corresponding contract state under mapping $\mathcal{M}$. Furthermore, let $\sigma^h_{j+k} = \chi^k(\sigma^h_j)$ and $\sigma^c_{i+1} = \chi(\sigma^c_i)$ be the hardware and contract state after the execution of an instruction. Inductively,*

$$\nexists\, \sigma^h_j, \sigma^c_i : \phi_{noc}\left(\sigma^h_j, \sigma^c_i\right) \wedge \phi^{\mathcal{M}}_{rel}\left(\sigma^h_j, \sigma^c_i\right) \wedge \neg\phi^{\mathcal{M}}_{rel}\left(\sigma^h_{j+k}, \sigma^c_{i+1}\right) \quad (11)$$

implies the first hardware compliance condition (Clause 1) from Definition 6 under normal operating conditions.

Proposition 1 specifies how exactly this check is performed. We use an SMT solver to check for states that satisfy both ϕ_{noc} and $\phi^{\mathcal{M}}_{\text{rel}}$, but their successors break $\phi^{\mathcal{M}}_{\text{rel}}$. Any such case is a counterexample to the state similarity property of Definition 6. Otherwise the property is inductive, and we use it in all further checks.

We also check that $\phi^{\mathcal{M}*}_{\text{pub}}(\cdot)$ is inductive because all of the further verification targets require this as an assumption. We check the inductiveness by asking an SMT solver

$$\nexists \sigma^h_j, \sigma^{h'}_j, \sigma^c_i, \sigma^{c'}_i : \phi^*_{\text{noc}}\left(\sigma^h_j, \sigma^{h'}_j, \sigma^c_i, \sigma^{c'}_i\right) \wedge \phi^{\mathcal{M}}_{\text{rel}}\left(\sigma^h_j, \sigma^c_i\right) \wedge \phi^{\mathcal{M}}_{\text{rel}}\left(\sigma^{h'}_j, \sigma^{c'}_i\right) \wedge \phi^{\mathcal{M}*}_{\text{pub}}\left(\sigma^h_j, \sigma^{h'}_j\right) \wedge \neg\phi^{\mathcal{M}*}_{\text{pub}}\left(\sigma^h_{j+k}, \sigma^{h'}_{j+k}\right). \quad (12)$$

Algorithm 1: Greedy minimization of required state bits.

Input : gate g to be modeled
$\Theta \leftarrow V^c$; $\theta_g \leftarrow$ `concat(`Θ`)` ;
if *formula (13) is SAT* **then**
  `error(`*"gate g cannot be modeled"*`)`;
for $v^c \in V^c$ **do**
  $\theta_g \leftarrow$ `concat(`$\Theta \setminus \{v^c\}$`)`;
  if *formula (13) is UNSAT* **then** $\Theta \leftarrow \Theta \setminus \{v^c\}$;
return θ_g;

If the solver is not able to find a solution, $\phi^{\mathcal{M}*}_{\text{pub}}$ is inductive and we assume $\phi^{\mathcal{M}*}_{\text{pub}}$ in addition to $\phi^{\mathcal{M}}_{\text{rel}}$ whenever we check properties under normal operating conditions over multiple executions.

4.4 Finding Modeling Functions for Gates

In this section, we introduce a method that finds a small number of contract registers from which the value of a hardware gate is modeled. This intermediate step determines, and later on restricts, the values a gate can have at the end of the previous instruction. Corollary 6 instantiates Theorem 5 under normal operating conditions and presents a method for checking whether the value of a hardware g can be modeled by contract state bits θ_g.

Corollary 6. *Let σ^h_j be a hardware state and σ^c_i the corresponding contract state under mapping $\mathcal{M}$ that fulfill both (11) and (12). Furthermore, let $\sigma^h_{j+k-1} = \chi^{k-1}(\sigma^h_j)$ be the last hardware state before the instruction terminates, and $\sigma^c_{i+1} = \chi(\sigma^c_i)$ be the contract state after the instruction terminates. The contract function θ_g models gate g in cycle $k-1$ under normal operating conditions if and only if*

$$\begin{aligned} \nexists \sigma^h_j, \sigma^{h'}_j, \sigma^c_i, \sigma^{c'}_i : \phi^*_{noc}\left(\sigma^h_j, \sigma^{h'}_j, \sigma^c_i, \sigma^{c'}_i\right) \wedge \\ \phi^{\mathcal{M}*}_{pub}\left(\sigma^h_j, \sigma^{h'}_j\right) \wedge \phi^{\mathcal{M}}_{rel}\left(\sigma^h_j, \sigma^c_i\right) \wedge \phi^{\mathcal{M}}_{rel}\left(\sigma^{h'}_j, \sigma^{c'}_i\right) \wedge \\ \theta_g\left(\sigma^c_{i+1}\right) = \theta_g\left(\sigma^{c'}_{i+1}\right) \wedge g\left(\sigma^h_{j+k-1}\right) \neq g\left(\sigma^{h'}_{j+k-1}\right). \end{aligned} \quad (13)$$

Corollary 6 instantiates Theorem 5 under assumption of ϕ_{noc} and $\phi^{\mathcal{M}*}_{\text{pub}}$. Here, g is the function f_H to be modeled, θ_g is the function f_C whose results are used as inputs for the modeling function, and $\phi^{\mathcal{M}}_{\text{rel}}$ is the relation Ψ between the hardware and contract states.

However, not all functions θ_g are useful, so a function like $\theta_g(\sigma^c) = \sigma^c$ would not really restrict the initial values of g. Instead, we propose the greedy minimization procedure shown in Algorithm 1. Here, we first check whether the hardware gate g can be modeled from the complete contract state. If this fails the contract does not model the hardware properly and the verification fails. Algorithm 1 iterates over all locations v^c in the contract state and checks whether they are needed for modeling g. In case they are not, *i.e.*, formula (13) is unsatisfiable, they are removed from θ_g. At the end, we have a (locally) minimal θ_g, where removing any component breaks the modeling of gate g.

4.5 Verifying that Leaks are Modeled

Lastly, we verify that the hardware leakage produced during the execution of an instruction can be modeled by the contract leakage

emitted during the execution of the same instruction. The set of transition leaks produced in the hardware, starting in state σ^h_j and executing a k-cycle instruction is given by

$$\mathcal{L}^h_{j,j+k} = \left\{ \lambda_g \left(\sigma^h_{j+l-1}, \sigma^h_{j+l} \right) \mid g \in G, 0 \leq l < k \right\}.$$

As established in Section 4.1, we analyze every hardware leak $\lambda_g(\sigma^h_{j+q-1}, \sigma^h_{j+q})$ separately and show that there is a set of leak statements $\mathcal{L}^g \subseteq \mathcal{L}^c_i$ such that every λ_g can be modeled by $\lambda(\sigma^c_i) \in \mathcal{L}^g$ whenever the corresponding `leak` statement is reached in the contract, written as $\phi_{\text{emit}}(\sigma^c_i, \lambda)$. Here, we use the intermediate proof of $g(\sigma^h_{j-1})$ being modeled by $\theta_g(\sigma^c_i)$ from Corollary 6.

Proposition 2. *Let σ^h_{j-1} be the predecessor of hardware state $\sigma^h_j = \chi(\sigma^h_{j-1})$, and σ^h_j be similar to contract state σ^c_i under mapping $\mathcal{M}$, fulfilling both (11) and (12). Furthermore, let $\sigma^h_{j+l} = \chi^l(\sigma^h_j)$ with $0 \leq l < k$ be the hardware states reached throughout the execution of a k-cycle instruction. Let λ_g be the leakage function of a hardware gate g, and θ_g be a contract function such that (13) holds. Crucially, let $\mathcal{L}^g \subseteq \mathcal{L}^c_i$ be a set of contract leaks, such that for every $\lambda(\sigma^c_i) \in \mathcal{L}^g$:*

$$\begin{aligned} \nexists \sigma^h_{j-}, \sigma^{h'}_{j-1}, \sigma^c_i, \sigma^{c'}_i : \phi^*_{noc}\left(\sigma^h_j, \sigma^{h'}_j, \sigma^c_i, \sigma^{c'}_i\right) \wedge \phi_{emit}\left(\sigma^c_i, \lambda\right) \wedge \\ \phi^{\mathcal{M}*}_{pub}\left(\sigma^h_j, \sigma^{h'}_j\right) \wedge \phi^{\mathcal{M}}_{rel}\left(\sigma^h_j, \sigma^c_i\right) \wedge \phi^{\mathcal{M}}_{rel}\left(\sigma^{h'}_j, \sigma^{c'}_i\right) \wedge \\ \left(\theta_g\left(\sigma^c_i\right) = \theta_g\left(\sigma^{c'}_i\right) \Rightarrow g\left(\sigma^h_{j-1}\right) = g\left(\sigma^{h'}_{j-1}\right)\right) \wedge \\ \lambda(\sigma^c_i) = \lambda(\sigma^{c'}_i) \wedge \lambda_g\left(\sigma^h_{j+l-1}, \sigma^h_{j+l}\right) \neq \lambda_g\left(\sigma^{h'}_{j+l-1}, \sigma^{h'}_{j+l}\right). \end{aligned}$$

The leak function λ_g in cycle l of a k-cycle instruction is modeled by a single contract leak function λ under relation $\phi^{\mathcal{M}}_{rel}$ and normal operating conditions ϕ_{noc}, according to Definition 6 if

$$\forall \sigma^h_j, \sigma^c_i : \phi_{noc}\left(\sigma^h_j, \sigma^c_i\right) \wedge \phi^{\mathcal{M}}_{rel}\left(\sigma^h_j, \sigma^c_i\right) \Rightarrow \bigvee_{\lambda(\sigma^c_i) \in \mathcal{L}^g} \phi_{emit}\left(\sigma^c_i, \lambda\right).$$

Again, the method outlined in Proposition 2 uses an SMT solver to show that the hardware cannot leak more information than the contract. If the solver is able to find a pair of states σ^h_j, σ^c_i for which the check fails, it has found a counterexample and the hardware does not comply with the leakage specified in the contract.

4.6 Modeling and Implementation

In this section, we briefly discuss the implementation and modeling details enabling our verification method. In particular, we discuss how all the formulas given to the SMT solver are constructed.

Unfolding circuits into SMT. Our method relies on the synthesized processor netlist to build the SMT formulas shown in Section 4. We follow the procedure established in the model checking community: the hardware state σ^h_j is represented symbolically using propositional variables. Each gate g in the processor is a symbolic expression of the variables representing hardware locations V^h. The expressions are generated by topologically iterating through the circuit and building the representation of each gate g from its inputs and gate type. With regard to clock cycles, the registers of the very first state σ^h_j, repsectively σ^h_{j-1}, are variables. In successor states σ^h_{j+l}, the registers are determined by their write-backs from the previous cycle. In a sense, we unfold the processor circuit symbolically k times for our verification.

Genoa to SMT translation. The translation of a contract to a SMT formula is based on an existing Sail back-end which allows to generate SMT formulas for custom predicates. However, the back-end cannot handle `leak` statements. We perform two code-rewriting passes from Genoa to Genoa. The first adds global state for each value in a `leak` statement and replaces the `leak` by an assignment to the respective global state. This reduces the Genoa DSL to the Sail subset supported by the SMT back-end. The second pass duplicates the variables representing contract state σ^c and leakages into prime and non-primed variants and duplicates the instruction-step function χ by rewriting it to operate on either σ^c_i or $\sigma^{c'}_i$ and resulting in σ^c_{i+1}, respectively $\sigma^{c'}_{i+1}$. Finally, Genoa ensures that the initial and final states are preserved by the SMT back-end and asserts that the predicates ϕ_{ret}, ϕ^*_{ret} and ϕ_{emit} hold. Our tool receives the resulting SMT code as input.

Gluing it all together. Configuration files play a central role in the generation of formulas. In particular, our verification procedure expects an input where all of the hardware locations are declared, and either mapped onto contract registers with $\phi^{\mathcal{M}}_{\text{rel}}$, subjected to developer assumptions ϕ_{dev}, port restrictions ϕ^*_{ports}, or instruction execution constraints ϕ^I_{instr}. Everything specified in the configuration is heavily sanity-checked, making sure that execution works properly, public signals $\phi^{\mathcal{M}*}_{\text{pub}}$ remain public, and every hardware location is declared. Similarly, intermediate results such as θ_g are cached in configuration files and checked upon use. The IBEX configuration is provided in Listing 7.

5 VERIFICATION PROCESS

We apply the verification method presented to the IBEX processor and detail the process and results. IBEX is an open source RISC-V processor that supports the **I**nteger, **E**mbedded, **M**ultiplication, **C**ompressed and **B**it manipulation ISA extensions [34]. For the purpose of our paper, we mainly target the **E** extension, although adding support for the others is possible. The IBEX pipeline consists of two stages, Instruction Fetch (IF) and Instruction Decode/Execute (ID/EX). Computations take place in the ID/EX stage, which consists of a decoder, a controller, and the register file, which forward the data into the arithmetic-logic unit (ALU) and the load-store unit (LSU). The outputs of the ALU and LSU are routed to the write-back logic (WBL), that decides which data is written into the register file. In the same pipeline stage, and hence in the same clock cycle, the result is written back into the register file.

The verification requires two manual and four automated steps:

(1) Configuration of the processor by defining constraints
(2) Definition of a mapping between hardware and contract
(3) Automated sanity-check to ensure instructions defined in the contract can still execute in the processor under constraints
(4) Automated check of similarity for resulting states (Section 4.3)
(5) Automated check for gate modeling functions (Section 4.4)
(6) Automated check for leakage modeling functions (Section 4.5)

In case any of the steps fail, the verification framework produces a detailed counterexample explaining the verification failure. The

developer must then adjust the configuration, the annotation, the contract, or even the processor in order to fix the problem and restart verification. Therefore, development and verification form a refinement loop producing improved contracts.

5.1 IBEX Configuration

We align the contract and the hardware by restricting the state of the processor throughout the execution of an instruction. We constrain the values of all registers with regard to the current instruction length and analyzed cycle. For the verification, we look at instructions when they reach the ID stage. At this point, signal `instr_rdata_id` carries the instruction bits and must be set equivalent to the argument of `step_ibex` in the contract.

Additionally, we need to make sure that instructions are only retired in the last cycle $k-1$ of a k-cycle instruction by constraining `instr_id_done` to be $\top$ in the last cycle and $\bot$ otherwise. Similarly, we enforce that the next instruction is fetched exactly in cycle $k-1$ by constraining `fetch_valid` and `id_in_ready`. We assert that there are no outstanding errors caused by the previous instruction by constraining registers `lsu_err_q`, `pmp_err_q`, `branch_set_raw`, and `data_err_i` to be $\bot$. To make sure that the state machines in the LSU and control unit start off in a valid state when the instruction starts executing, we add several further constraints. Finally, we also assert that there is no reset through `rst_ni` and no interrupt signals `irq_*`, `debug_req_i` are triggered to match the developers expected behavior.

One of the main challenges in modeling the processor environment is the memory interface. Whenever the processor requests data by setting `data_req_o` to $\top$, the next cycle provides a grant with `data_rvalid_i` set to $\top$ and the corresponding read data being available at `data_rdata_i`. Here, we additionally require memory to only provide acknowledgement through `data_rvalid_i` if there was a request, and not provide any data on the input `data_rdata_i` otherwise. This is due to an oversight in IBEX, which causes the `data_rvalid_i` signal to overrule all other signals in the processor and ultimately issue an erroneous write-back.

5.2 Complete Power Contract for IBEX

We have proven that IBEX is compliant with the contract. In this section, we discuss the observed behavior and compare the findings to existing models for other architectures.

Most instructions have a common leakage pattern modeled in `common_leakage` in Listing 4. The IBEX processor combines the previous outputs of the register file (modeled in leakage states `rf_pA`, `rf_pB`) with the current outputs `rs1_val` and `rs2_val`, as well as the address and value of the last memory access `mem_last_addr` and `mem_last_read`. This leak statement models all transition leakage and value leakage produced in the ALU and the WBL. None of the operands in the `leak` statement can be removed without breaking compliance since distinct parts of IBEX cause these combinations. The WBL causes additional combinations: ALU or branch instructions after a memory load cause a transition between their results.

This common leakage covers leak effects previously discussed in related works. It models transition leakage produced in the ALU and WB stage, whose source are the two read ports of the register file. Transitions in the first and second operand of instructions are

Listing 4: Common leakage occurring in every instruction.

```
1  // see license in Listing 6
2  function common_leakage(rs1_val, rs2_val) = {
3    leak(rs1_val, rs2_val, rf_pA, rf_pB,
4         mem_last_addr, mem_last_read);
5    rf_pA = rs1_val; rf_pB = rs2_val; /* update read ports */
6    mem_last_read = 0x00000000; /* clear data memory port */ }
```

Listing 5: Specialized leakage occurring during loads.

```
7  // see license in Listing 6
8  function load_leakage(rs1_val : xlenbits, rs2_val : xlenbits
       ↪ , addr: xlenbits, req_data: xlenbits) = {
9    leak(rf_pA, rf_pB, rs1_val, rs2_val);
10   leak(rf_pA, rf_pB, mem_last_addr, mem_last_read);
11   leak(addr, req_data, mem_last_addr);
12   rf_pA = rs1_val; rf_pB = rs2_val;
13   mem_last_read = req_data; mem_last_addr = addr; }
```

well-known [38, 39]. Interestingly, prior empirical analysis of the ARM M0 [39], a processor in the same performance and size class as IBEX, did not report interactions between the data loaded in the previous instruction and the current or past ALU operands.

Furthermore, the leakage is even caused by instructions which have no register operands like `LUI` (load unsigned immediate). The root cause of this effect is that the register file always decodes specific instruction bits as register addresses and forwards their contents to the ALU. In the case of `LUI`, these bits are actually part of the immediate value. This effect was observed by Gigerl *et al.* [25] but we characterize and describe the behavior more accurately. Prior analyses of the ARM M0 did not report similar effects [39], and instead reports that instructions with an immediate field behave as if they have only one operand which could be to the microarchitecture of the CPU or gaps in the empirical modeling procedure.

The leakage of load instructions modeled by `load_leakage` (Listing 5) differs from all other instructions. Here, the leak statement in `common_leakage` can be broken down into smaller leak statements. First off, because the ALU is always active, the current and previous values of the register file outputs are combined in line 9. Line 10 specifies leakage inherited from the prior instruction's WBL through transition leakage. In contrast to prior work [25, 45], consecutive load instructions do not cause transition leakage between the loaded data because the contract disallows misaligned memory accesses. Similarly, IBEX does not expose transition leakage in subsequent memory writes, unlike several ARM architectures [9, 38]. This is likely because IBEX does not have additional registers in the memory path like other processors. However, IBEX does produce transition leakage between memory access addresses of loads and stores separated by an arbitrary amount of other instructions, as shown in line 11. IBEX causes this leakage because it always stores the last address for error-handling purposes.

We emphasize that the contract is provably complete for branching instructions; the behavior of these central instructions was so far not characterized.

5.3 Discussion

Performance. Verifying that a CPU design complies with a contract is computationally intensive, but can be well parallelized. We

ran the full verification on an Intel Xeon E5-4669 CPU with 88 logical cores running at 2.20GHz. The most time-consuming verification task is the search for gate modeling functions, which takes about 30.6 hours. This step is done once and then cached, and any changes to leak statements in the contract do not require it to be run again. Verifying the leakage modeling requires another 4.9 hours.

Adaptability and scalability. While we demonstrate our approach on the RISC-V IBEX core we emphasize that it is neither limited to RISC-V processors, nor the IBEX core. Verifying contract compliance for similar architectures and processors requires adapting the tool to their pipeline and properly configuring the verification procedure. We believe that the effort of including other architectures or processors mainly comes from adapting a contract to their CPU pipeline and properly configuring the verification. In any case, this effort is only required once per CPU netlist, and can be made either by the CPU designers themselves, or by any other person in case the CPU netlist is not IP restricted. Anybody with access to the contract can then verify masked software against it, without the need of having access to the CPU netlist itself.

Hardware constraints and shortcomings. We currently limit the scope of our contracts, and thus also the scope of masking verification, to the CPU core itself. Clearly, there may also be some other components that could cause power side-channel leaks during the execution of masked software. For example, RAM or data caches are another location where unintentional combinations of shares could occur. Within our framework, the leakage of such components would need to be verified separately and modeled within the CPU contract. Additionally, there are also cases when no "good" contract can be written for a CPU. One such example would be CPUs where the register file output is computed with a multiplexer tree, and would lead to a `leak` statement containing the whole register file.

So far, we analyze an instruction starting with the decode stage, which assumes that the fetch stage does not expose leakage depending on the fetched instruction. From what we have seen in IBEX and other processors, there is no leakage in the fetch stage that depends on the bits of the fetched instruction. For a similar reason, speculation or (secret-dependent) branch prediction is not a primary concern for our current analysis since these are usually not present in embedded devices.

Masking verification. Our tool currently focuses on value leakage and transition leakage, while our theoretical framework supports arbitrary gate-level leakage. Extending our verification tool to include further effects such as glitches makes an interesting future research question, and could be achieved by extending the encoding of leakage from Section 4.5. Our verification methodology and the contracts themselves support bitsliced and n-sliced masking [11], which are among the most popular implementation techniques for masked software. Our analysis, as well as prior work by others, observes joint leakage of bits stored in the same 32-bit register [22, 25, 39], rendering exotic concepts like share-slicing inherently insecure.

6 APPLICATION AND VALIDATION

We implement higher-order masked gadgets in software and verify their security against the IBEX contract. We demonstrate the benefit of contracts and validate our methodology by repeating the verification with an independent tool that directly verifies programs against the processor netlist and all of its side-effects. Finally, we assess the precision of contracts by confirming that the abstract leakage specification does not demand needless protection.

Table 1: Verifying software implementations of 2^{nd}-order probing secure gadgets using the contract or the netlist of IBEX results in the same confirmation of security (✔) at reduced verification time and validates our approach.

Gadget	t	# Instr.	# Clear.	Verification time	
				Contract	Netlist
`AND`	2	62	10	< 1 s ✔	284.63 s ✔
`Refresh`	2	19	0	< 1 s ✔	32.85 s ✔
`XOR`	2	16	1	< 1 s ✔	50.79 s ✔
`NOT`	2	5	0	< 1 s ✔	63.32 s ✔

Validation of the methodology and tool. We port multiple 2^{nd} order masked gadgets presented by Barthe *et al.* [9] to RISC-V and check their security (software compliance) using scVerif. For this, we perform a manual translation (which could be automated) of the Genoa contract to the DSL of scVerif and adopt its front-end slightly to accept RISC-V assembly. The resulting tool allows to prove software compliance (Definition 4), as well as the weaker notion of probing security for assembly implementations against the IBEX contract. There are no software masking verifiers capable of verifying t–(S)NI while also accounting for the CPU netlists. Therefore, we compare our results against Coco [25], a tool that accounts for the CPU netlist, but only supports probing security.

All gadgets are hardened by adding the least amount of clearing instructions until they are threshold probing secure, *i.e.*, compliant with the contract under a weaker notion of security yet claimed secure against all gate-level leakage of IBEX. We have checked each gadget with both Coco and scVerif, and the results are shown in Table 1. The correctness of our methodology, hardware compliance checking tool and pen-and-paper model reduction (Theorem 2) are confirmed since there is no case where scVerif reports security while Coco rejects it.

Quality of contracts. We check that, whenever one of the clearing instructions that mitigate contract leakage is removed, Coco also rejects the program due to some gate-level leakage in the CPU netlist. If Coco would report that an implementation with less clearings is still secure, it would means that the leakage generalization in the contract was too broad and requires needless hardening of the program. In our tests, whenever we removed any of the 11 clearings from Table 1, Coco always reported some gate-level leakage that breaks probing security. This indicates that our contract does not cause wrong insecurity reports (false negatives).

7 CONCLUSION

We introduced a methodology for creating software leakage models and proving their completeness based on the netlist of a CPU. Our rigorous approach allows us to treat the model as contract between the software and the hardware which provably guarantees end-to-end security: any implementation secure w.r.t. a contract is also secure on any compliant processor for all leakages exposed

at gate-level. Overall the result significantly improve the secure construction of hardened software implementations.

Besides providing strong guarantees of side-channel resistance, easing the safe porting of programs to different CPUs and the most extensive modeling of different instructions's side-channel leakage, we think our approach could be beneficial for other applications as well. In particular, it could be used for leakage emulators or statistical security evaluations that can be derived from the executable Genoa contracts.

ACKNOWLEDGMENTS

This work was supported by the Austrian Research Promotion Agency (FFG) through projects FERMION (grant number 867542) and AWARE (grant number FO999891092). The work received funding from the Federal Ministry of Education and Research (BMBF) as part of the VE-Jupiter project (grant number 16ME0231K). This work was supported by the Graz University of Technology LEAD project "Dependable Internet of Things in Adverse Environments". Parts of this work received sponsoring from NXP Semiconductors Austria.

REFERENCES

[1] Arnold Abromeit, Florian Bache, Leon A. Becker, Marc Gourjon, Tim Güneysu, Sabrina Jorn, Amir Moradi, Maximilian Orlt, and Falk Schellenberg. 2021. Automated Masking of Software Implementations on Industrial Microcontrollers. In *Design, Automation & Test in Europe Conference & Exhibition, DATE 2021, Grenoble, France, February 1-5, 2021*. IEEE, 1006–1011. https://doi.org/10.23919/DATE51398.2021.9474183

[2] Alasdair Armstrong, Thomas Bauereiss, Brian Campbell, Kathryn E. Gray, Robert Norton-Wright, Christopher Pulte, Shaked Flur, and Peter Sewell. July 2021 (accessed January 12, 2022). https://raw.githubusercontent.com/rems-project/sail/sail2/manual.pdf.

[3] Alasdair Armstrong, Thomas Bauereiss, Brian Campbell, Alastair Reid, Kathryn E. Gray, Robert M. Norton, Prashanth Mundkur, Mark Wassell, Jon French, Christopher Pulte, Shaked Flur, Ian Stark, Neel Krishnaswami, and Peter Sewell. 2018 (accessed January 12, 2022). MiniSail: A core calculus for Sail. https://www.cl.cam.ac.uk/~mpew2/papers/minisail_anf.pdf.

[4] Alasdair Armstrong, Thomas Bauereiss, Brian Campbell, Alastair Reid, Kathryn E. Gray, Robert M. Norton, Prashanth Mundkur, Mark Wassell, Jon French, Christopher Pulte, Shaked Flur, Ian Stark, Neel Krishnaswami, and Peter Sewell. 2019. ISA Semantics for ARMv8-A, RISC-V, and CHERI-MIPS. In *Proc. 46th ACM SIGPLAN Symposium on Principles of Programming Languages*. https://doi.org/10.1145/3290384 Proc. ACM Program. Lang. 3, POPL, Article 71.

[5] Josep Balasch, Benedikt Gierlichs, Vincent Grosso, Oscar Reparaz, and François-Xavier Standaert. 2014. On the Cost of Lazy Engineering for Masked Software Implementations. In *Smart Card Research and Advanced Applications - 13th International Conference, CARDIS 2014, Paris, France, November 5-7, 2014. Revised Selected Papers (Lecture Notes in Computer Science, Vol. 8968)*, Marc Joye and Amir Moradi (Eds.). Springer, 64–81. https://doi.org/10.1007/978-3-319-16763-3_5

[6] Clark Barrett, Aaron Stump, and Cesare Tinelli. 2010. *The SMT-LIB Standard: Version 2.0.* Technical Report. Department of Computer Science, The University of Iowa. Available at `www.SMT-LIB.org`.

[7] Gilles Barthe, Sonia Belaïd, François Dupressoir, Pierre-Alain Fouque, Benjamin Grégoire, and Pierre-Yves Strub. 2015. Verified Proofs of Higher-Order Masking. In *EUROCRYPT 2015, Part I (LNCS, Vol. 9056)*, Elisabeth Oswald and Marc Fischlin (Eds.). Springer, Heidelberg, 457–485. https://doi.org/10.1007/978-3-662-46800-5_18

[8] Gilles Barthe, Sonia Belaïd, François Dupressoir, Pierre-Alain Fouque, Benjamin Grégoire, Pierre-Yves Strub, and Rébecca Zucchini. 2016. Strong Non-Interference and Type-Directed Higher-Order Masking. In *ACM CCS 2016*, Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi (Eds.). ACM Press, 116–129. https://doi.org/10.1145/2976749.2978427

[9] Gilles Barthe, Marc Gourjon, Benjamin Grégoire, Maximilian Orlt, Clara Paglialonga, and Lars Porth. 2021. Masking in Fine-Grained Leakage Models: Construction, Implementation and Verification. *IACR TCHES* 2021, 2 (2021), 189–228. https://doi.org/10.46586/tches.v2021.i2.189-228 https://tches.iacr.org/index.php/TCHES/article/view/8792.

[10] Omid Bazangani, Alexandre Iooss, Ileana Buhan, and Lejla Batina. 2021. ABBY: Automating the creation of fine-grained leakage models. Cryptology ePrint Archive, Report 2021/1569. https://ia.cr/2021/1569.

[11] Sonia Belaïd, Pierre-Évariste Dagand, Darius Mercadier, Matthieu Rivain, and Raphaël Wintersdorff. 2020. Tornado: Automatic Generation of Probing-Secure Masked Bitsliced Implementations. In *EUROCRYPT 2020, Part III (LNCS, Vol. 12107)*, Anne Canteaut and Yuval Ishai (Eds.). Springer, Heidelberg, 311–341. https://doi.org/10.1007/978-3-030-45727-3_11

[12] Gaëtan Cassiers and François-Xavier Standaert. 2020. Trivially and Efficiently Composing Masked Gadgets With Probe Isolating Non-Interference. *IEEE Trans. Inf. Forensics Secur.* 15 (2020), 2542–2555. https://doi.org/10.1109/TIFS.2020.2971153

[13] Suresh Chari, Charanjit S. Jutla, Josyula R. Rao, and Pankaj Rohatgi. 1999. Towards Sound Approaches to Counteract Power-Analysis Attacks. In *CRYPTO'99 (LNCS, Vol. 1666)*, Michael J. Wiener (Ed.). Springer, Heidelberg, 398–412. https://doi.org/10.1007/3-540-48405-1_26

[14] Thomas De Cnudde, Oscar Reparaz, Begül Bilgin, Svetla Nikova, Ventzislav Nikov, and Vincent Rijmen. 2016. Masking AES with d+1 Shares in Hardware. In *Cryptographic Hardware and Embedded Systems - CHES 2016 - 18th International Conference, Santa Barbara, CA, USA, August 17-19, 2016, Proceedings (Lecture Notes in Computer Science, Vol. 9813)*. Springer, 194–212.

[15] Jean-Sébastien Coron, Christophe Giraud, Emmanuel Prouff, Soline Renner, Matthieu Rivain, and Praveen Kumar Vadnala. 2012. Conversion of Security Proofs from One Leakage Model to Another: A New Issue. In *Constructive Side-Channel Analysis and Secure Design - Third International Workshop, COSADE 2012, Darmstadt, Germany, May 3-4, 2012. Proceedings (Lecture Notes in Computer Science, Vol. 7275)*. Springer, 69–81.

[16] Jean-Sébastien Coron. 2014. Higher Order Masking of Look-Up Tables. In *EUROCRYPT 2014 (LNCS, Vol. 8441)*, Phong Q. Nguyen and Elisabeth Oswald (Eds.). Springer, Heidelberg, 441–458. https://doi.org/10.1007/978-3-642-55220-5_25

[17] Thomas De Cnudde, Begül Bilgin, Benedikt Gierlichs, Ventzislav Nikov, Svetla Nikova, and Vincent Rijmen. 2017. Does Coupling Affect the Security of Masked Implementations?. In *COSADE 2017 (LNCS, Vol. 10348)*, Sylvain Guilley (Ed.). Springer, Heidelberg, 1–18. https://doi.org/10.1007/978-3-319-64647-3_1

[18] Wouter de Groot, Kostas Papagiannopoulos, Antonio de la Piedra, Erik Schneider, and Lejla Batina. 2016. Bitsliced Masking and ARM: Friends or Foes?. In *Lightweight Cryptography for Security and Privacy - 5th International Workshop, LightSec 2016, Aksaray, Turkey, September 21-22, 2016, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 10098)*. Springer, 91–109.

[19] Leonardo Mendonça de Moura and Nikolaj S. Bjørner. 2008. Z3: An Efficient SMT Solver. In *Tools and Algorithms for the Construction and Analysis of Systems, 14th International Conference, TACAS 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29-April 6, 2008. Proceedings (Lecture Notes in Computer Science, Vol. 4963)*, C. R. Ramakrishnan and Jakob Rehof (Eds.). Springer, 337–340. https://doi.org/10.1007/978-3-540-78800-3_24

[20] Sebastian Faust, Vincent Grosso, Santos Merino Del Pozo, Clara Paglialonga, and François-Xavier Standaert. 2018. Composable Masking Schemes in the Presence of Physical Defaults & the Robust Probing Model. *IACR TCHES* 2018, 3 (2018), 89–120. https://doi.org/10.13154/tches.v2018.i3.89-120 https://tches.iacr.org/index.php/TCHES/article/view/7270.

[21] Andreas Fröhlich, Gergely Kovásznai, and Armin Biere. 2013. More on the Complexity of Quantifier-Free Fixed-Size Bit-Vector Logics with Binary Encoding. In *Computer Science - Theory and Applications - 8th International Computer Science Symposium in Russia, CSR 2013, Ekaterinburg, Russia, June 25-29, 2013. Proceedings (Lecture Notes in Computer Science, Vol. 7913)*, Andrei A. Bulatov and Arseny M. Shur (Eds.). Springer, 378–390. https://doi.org/10.1007/978-3-642-38536-0_33

[22] Si Gao, Ben Marshall, Dan Page, and Elisabeth Oswald. 2019. Share-slicing: Friend or Foe? *IACR TCHES* 2020, 1 (2019), 152–174. https://doi.org/10.13154/tches.v2020.i1.152-174 https://tches.iacr.org/index.php/TCHES/article/view/8396.

[23] Si Gao and Elisabeth Oswald. 2021. A Novel Completeness Test and its Application to Side Channel Attacks and Simulators. Cryptology ePrint Archive, Report 2021/756. https://ia.cr/2021/756.

[24] Si Gao, Elisabeth Oswald, and Dan Page. 2021. Reverse Engineering the Micro-Architectural Leakage Features of a Commercial Processor. Cryptology ePrint Archive, Report 2021/794. https://ia.cr/2021/794.

[25] Barbara Gigerl, Vedad Hadzic, Robert Primas, Stefan Mangard, and Roderick Bloem. 2021. Coco: Co-Design and Co-Verification of Masked Software Implementations on CPUs. In *30th USENIX Security Symposium, USENIX Security 2021, August 11-13, 2021*, Michael Bailey and Rachel Greenstadt (Eds.). USENIX Association, 1469–1468. https://www.usenix.org/conference/usenixsecurity21/presentation/gigerl

[26] Barbara Gigerl, Robert Primas, and Stefan Mangard. 2021. Secure and Efficient Software Masking on Superscalar Pipelined Processors. In *Advances in Cryptology - ASIACRYPT 2021 - 27th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 6-10, 2021, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 13091)*, Mehdi Tibouchi and Huaxiong Wang (Eds.). Springer, 3–32. https://doi.org/10.1007/978-3-030-92075-3_1

[27] Hannes Gross, Stefan Mangard, and Thomas Korak. 2016. Domain-Oriented Masking: Compact Masked Hardware Implementations with Arbitrary Protection Order. In *Proceedings of the 2016 ACM Workshop on Theory of Implementation Security* (Vienna, Austria) *(TIS '16)*. ACM, New York, NY, USA, 3–3. https://doi.org/10.1145/2996366.2996426

[28] Hannes Gross, Stefan Mangard, and Thomas Korak. 2017. An Efficient Side-Channel Protected AES Implementation with Arbitrary Protection Order. In *CT-RSA 2017, San Francisco, CA, USA, February 14–17, 2017, Proceedings*, Helena Handschuh (Ed.). Springer International Publishing, Cham, 95–112. https://doi.org/10.1007/978-3-319-52153-4_6

[29] Marco Guarnieri, Boris Köpf, Jan Reineke, and Pepe Vila. 2021. Hardware-Software Contracts for Secure Speculation. In *42nd IEEE Symposium on Security and Privacy, SP 2021, San Francisco, CA, USA, 24-27 May 2021*. IEEE, 1868–1883. https://doi.org/10.1109/SP40001.2021.00036

[30] Vedad Hadzic and Roderick Bloem. 2021. COCOALMA: A Versatile Masking Verifier. In *Formal Methods in Computer Aided Design, FMCAD 2021, New Haven, CT, USA, October 19-22, 2021*. IEEE, 1–10. https://doi.org/10.34727/2021/isbn.978-3-85448-046-4_9

[31] Yuval Ishai, Amit Sahai, and David Wagner. 2003. Private Circuits: Securing Hardware against Probing Attacks. In *CRYPTO 2003 (LNCS, Vol. 2729)*, Dan Boneh (Ed.). Springer, Heidelberg, 463–481. https://doi.org/10.1007/978-3-540-45146-4_27

[32] Paul C. Kocher, Joshua Jaffe, and Benjamin Jun. 1999. Differential Power Analysis. In *Advances in Cryptology - CRYPTO '99, 19th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15-19, 1999, Proceedings (Lecture Notes in Computer Science, Vol. 1666)*, Michael J. Wiener (Ed.). Springer, 388–397. https://doi.org/10.1007/3-540-48405-1_25

[33] Paul C. Kocher, Joshua Jaffe, and Benjamin Jun. 1999. Differential Power Analysis. In *CRYPTO '99*. Springer-Verlag, London, UK, 10 pages. http://dl.acm.org/citation.cfm?id=646764.703989

[34] lowRISC. [n. d.]. Ibex RISC-V Core. https://github.com/lowRISC/ibex.

[35] Stefan Mangard, Thomas Popp, and Berndt M. Gammel. 2005. Side-Channel Leakage of Masked CMOS Gates. In *CT-RSA 2005 (LNCS, Vol. 3376)*, Alfred Menezes (Ed.). Springer, Heidelberg, 351–365. https://doi.org/10.1007/978-3-540-30574-3_24

[36] Stefan Mangard, Norbert Pramstaller, and Elisabeth Oswald. 2005. Successfully Attacking Masked AES Hardware Implementations. In *CHES 2005 (LNCS, Vol. 3659)*, Josyula R. Rao and Berk Sunar (Eds.). Springer, Heidelberg, 157–171. https://doi.org/10.1007/11545262_12

[37] Stefan Mangard, Norbert Pramstaller, and Elisabeth Oswald. 2005. Successfully Attacking Masked AES Hardware Implementations. In *CHES (Lecture Notes in Computer Science, Vol. 3659)*. Springer, 157–171.

[38] Ben Marshall, Dan Page, and James Webb. 2021. MIRACLE: MIcRo-ArChitectural Leakage Evaluation: A study of micro-architectural power leakage across many devices. *IACR Transactions on Cryptographic Hardware and Embedded Systems* 2022, 1 (Nov. 2021), 175–220. https://doi.org/10.46586/tches.v2022.i1.175-220

[39] David McCann, Elisabeth Oswald, and Carolyn Whitnall. 2017. Towards Practical Tools for Side Channel Aware Software Engineering: 'Grey Box' Modelling for Instruction Leakages. In *USENIX Security 2017*, Engin Kirda and Thomas Ristenpart (Eds.). USENIX Association, 199–216.

[40] Thomas S. Messerges. 2000. Using Second-Order Power Analysis to Attack DPA Resistant Software. In *CHES 2000*. https://doi.org/10.1007/3-540-44499-8_19

[41] Lauren De Meyer, Elke De Mulder, and Michael Tunstall. 2020. On the Effect of the (Micro)Architecture on the Development of Side-Channel Resistant Software. *IACR Cryptol. ePrint Arch.* 2020 (2020), 1297. https://eprint.iacr.org/2020/1297

[42] Prashanth Mundkur, Rishiyur S. Nikhil, Bluespec Inc, Jon French, Brian Campbell, Robert Norton-Wright, Alasdair Armstrong, Thomas Bauereiss, Shaked Flur, Christopher Pulte, Peter Sewell, Alexander Richardson, Hesham Almatary, Jessica Clarke, Microsoft, Nathaniel Wesley Filardo, Peter Rugg, and Aril Computer Corp. August 2021 (accessed January 17, 2022). RISCV Sail Model. https://github.com/riscv/sail-riscv.

[43] Svetla Nikova, Christian Rechberger, and Vincent Rijmen. 2006. Threshold Implementations Against Side-Channel Attacks and Glitches. In *ICICS 06 (LNCS, Vol. 4307)*, Peng Ning, Sihan Qing, and Ninghui Li (Eds.). Springer, Heidelberg, 529–545.

[44] Svetla Nikova, Vincent Rijmen, and Martin Schläffer. 2011. Secure Hardware Implementation of Nonlinear Functions in the Presence of Glitches. *Journal of Cryptology* 24, 2 (April 2011), 292–321. https://doi.org/10.1007/s00145-010-9085-7

[45] Kostas Papagiannopoulos and Nikita Veshchikov. 2017. Mind the Gap: Towards Secure 1st-Order Masking in Software. In *COSADE 2017 (LNCS, Vol. 10348)*, Sylvain Guilley (Ed.). Springer, Heidelberg, 282–297. https://doi.org/10.1007/978-3-319-64647-3_17

[46] Emmanuel Prouff and Matthieu Rivain. 2013. Masking against Side-Channel Attacks: A Formal Security Proof. In *EUROCRYPT 2013 (LNCS, Vol. 7881)*, Thomas Johansson and Phong Q. Nguyen (Eds.). Springer, Heidelberg, 142–159. https://doi.org/10.1007/978-3-642-38348-9_9

[47] Jean-Jacques Quisquater and David Samyde. 2001. ElectroMagnetic Analysis (EMA): Measures and Counter-Measures for Smart Cards. In *Smart Card Programming and Security, International Conference on Research in Smart Cards, E-smart 2001, Cannes, France, September 19-21, 2001, Proceedings (Lecture Notes in Computer Science, Vol. 2140)*, Isabelle Attali and Thomas P. Jensen (Eds.). Springer, 200–210. https://doi.org/10.1007/3-540-45418-7_17

[48] Oscar Reparaz, Begül Bilgin, Svetla Nikova, Benedikt Gierlichs, and Ingrid Verbauwhede. 2015. Consolidating Masking Schemes. In *CRYPTO 2015*. https://doi.org/10.1007/978-3-662-47989-6_37

[49] Peter Sewell. June 2020 (accessed January 14, 2022). ISA Formal Spec Public Review. https://github.com/riscvarchive/ISA_Formal_Spec_Public_Review.

[50] Madura A. Shelton, Niels Samwel, Lejla Batina, Francesco Regazzoni, Markus Wagner, and Yuval Yarom. 2021. Rosita: Towards Automatic Elimination of Power-Analysis Leakage in Ciphers. In *NDSS*. The Internet Society.

APPENDIX

7.1 RISC-V Model for IBEX

The missing parts of our contract for IBEX are depicted in Listing 6.

Listing 6: Contract model of remaining instructions for IBEX.

```
/*==============================*/
/* RISCV Sail Model */
/* This Sail RISC-V architecture model, comprising all files
    ↪  and directories except for the snapshots of the
    ↪ Lem and Sail libraries in the prover_snapshots
    ↪ directory (which include copies of their licences),
    ↪  is subject to the BSD two-clause licence below. */
/* Copyright (c) 2017-2021 Prashanth Mundkur, Rishiyur S.
    ↪ Nikhil and Bluespec Inc., Jon French, Brian
    ↪ Campbell, Robert Norton-Wright, Alasdair Armstrong,
    ↪  Thomas Bauereiss, Shaked Flur, Christopher Pulte,
    ↪ Peter Sewell, Alexander Richardson, Hesham Almatary,
    ↪  Jessica Clarke, Microsoft, for contributions by
    ↪ Robert Norton-Wright and Nathaniel Wesley Filardo,
    ↪ Peter Rugg and Aril Computer Corp., for
    ↪ contributions by Scott Johnson */
/* Copyright 2020-2022 - TUHH, TU Graz */
/* All rights reserved. */
/* This software was developed by the above within the
    ↪ Rigorous Engineering of Mainstream Systems (REMS)
    ↪ project, partly funded by EPSRC grant EP/K008528/1,
    ↪  at the Universities of Cambridge and Edinburgh. */
/* This software was developed by SRI International and the
    ↪ University of Cambridge Computer Laboratory (
    ↪ Department of Computer Science and Technology)
    ↪ under DARPA/AFRL contract FA8650-18-C-7809 ("CIFV"),
    ↪  and under DARPA contract HR0011-18-C-0016 ("ECATS")
    ↪  as part of the DARPA SSITH research programme. */
/* This project has received funding from the European
    ↪ Research Council (ERC) under the European Union's
    ↪ Horizon 2020 research and innovation programme (
    ↪ grant agreement 789108, ELVER). */
/* This software has received funding from the Federal
    Ministry of Education and Research
    (BMBF) as part of the VE-Jupiter project grant
    16ME0231K. */
/* This work was supported by the Austrian Research
    Promotion Agency (FFG) through the FERMION project
    (grant number 867542). */
```

```
/* Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: */
/* 1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. */
/* 2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. */
/* THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. */
/*==============================*/

val common_leakage : (xlenbits, xlenbits) -> unit effect {rreg, wreg, leakage}
function common_leakage(rs1_val : xlenbits, rs2_val : xlenbits) =
{
  leak(rs1_val, rs2_val, rf_pA, rf_pB, mem_last_addr, mem_last_read);
  rf_pA = rs1_val;
  rf_pB = rs2_val;
  mem_last_read = 0x00000000;
}

val overwrite_leakage : (regidx, xlenbits) -> unit effect {rreg, leakage}
function overwrite_leakage(dest_idx : regidx, res : xlenbits ) = {
  let x1_n = if (dest_idx == 0b00001)
    then {res} else {x1} in leak(x1 , x1_n );
  let x2_n = if (dest_idx == 0b00010)
    then {res} else {x2} in leak(x2 , x2_n );
  let x3_n = if (dest_idx == 0b00011)
    then {res} else {x3} in leak(x3 , x3_n );
  let x4_n = if (dest_idx == 0b00100)
    then {res} else {x4} in leak(x4 , x4_n );
  let x5_n = if (dest_idx == 0b00101)
    then {res} else {x5} in leak(x5 , x5_n );
  let x6_n = if (dest_idx == 0b00110)
    then {res} else {x6} in leak(x6 , x6_n );
  let x7_n = if (dest_idx == 0b00111)
    then {res} else {x7} in leak(x7 , x7_n );
  let x8_n = if (dest_idx == 0b01000)
    then {res} else {x8} in leak(x8 , x8_n );
  let x9_n = if (dest_idx == 0b01001)
    then {res} else {x9} in leak(x9 , x9_n );
  let x10_n = if (dest_idx == 0b01010)
    then {res} else {x10} in leak(x10, x10_n);
  let x11_n = if (dest_idx == 0b01011)
    then {res} else {x11} in leak(x11, x11_n);
  let x12_n = if (dest_idx == 0b01100)
    then {res} else {x12} in leak(x12, x12_n);
  let x13_n = if (dest_idx == 0b01101)
    then {res} else {x13} in leak(x13, x13_n);
  let x14_n = if (dest_idx == 0b01110)
    then {res} else {x14} in leak(x14, x14_n);
  let x15_n = if (dest_idx == 0b01111)
    then {res} else {x15} in leak(x15, x15_n);
}

/* ***************************** */

enum uop = {RISCV_LUI, RISCV_AUIPC}
union clause ast = UTYPE : (bits(20), regidx, uop)

mapping encdec_uop : uop <-> bits(7) = {
  RISCV_LUI <-> 0b0110111,
  RISCV_AUIPC <-> 0b0010111
}

mapping clause encdec = UTYPE(imm, rd, op)
  <-> imm @ rd @ encdec_uop(op)
  if (rd[4] == bitzero)

function clause execute UTYPE(imm, rd, op) = {
  let rs1_val = X(0b0 @ imm[6 .. 3]);
  let rs2_val = X(0b0 @ imm[11 .. 8]);
  common_leakage(rs1_val, rs2_val);

  let off : xlenbits = EXTS(imm @ 0x000);
  let ret : xlenbits = match op {
    RISCV_LUI => off,
    RISCV_AUIPC => get_arch_pc() + off
  };

  X(rd) = ret;
  RETIRE_SUCCESS
}

/* ***************************** */

union clause ast = RISCV_JAL : (bits(21), regidx)

mapping clause encdec =
  RISCV_JAL(imm_19 @ imm_7_0 @ imm_8 @ imm_18_13 @ imm_12_9 @ 0b0, rd)
```

```
  <-> imm_19 : bits(1) @ imm_18_13 : bits(6) @ imm_12_9 :
       ↪ bits(4) @ imm_8 : bits(1) @ imm_7_0 : bits(8) @ rd
       ↪  @ 0b1101111
  if (rd[4] == bitzero)

function clause execute (RISCV_JAL(imm, rd)) = {
  let rs1_val = X(0b0 @ imm[18 .. 15]);
  let rs2_val = X(0b0 @ imm[3 .. 1]
                  @ subrange_bits(imm, 11, 11));
  common_leakage(rs1_val, rs2_val);

  let t : xlenbits = PC + EXTS(imm);
  let rd_next = get_next_pc();

  overwrite_leakage(rd, rd_next);
  X(rd) = rd_next;
  if t[1 .. 0] == 0b00 then {
    set_next_pc(t);
    RETIRE_SUCCESS
  } else RETIRE_FAIL
}

/* ***************************** */

union clause ast =
  RISCV_JALR : (bits(12), regidx, regidx)

mapping clause encdec = RISCV_JALR(imm, rs1, rd)
  <-> imm @ rs1 @ 0b000 @ rd @ 0b1100111
  if (rs1[4] == bitzero & rd[4] == bitzero)

function clause execute (RISCV_JALR(imm, rs1, rd)) = {
  let rs1_val = X(rs1);
  let rs2_val = X(0b0 @ imm[3..0]);
  common_leakage(rs1_val, rs2_val);

  let t : xlenbits =
    [(rs1_val + EXTS(imm)) with 0 = bitzero];
  if t[1 .. 0] == 0b00 then {
    overwrite_leakage(rd, get_next_pc());
    X(rd) = get_next_pc();
    set_next_pc(t);
    RETIRE_SUCCESS
  } else RETIRE_FAIL
}

/* ***************************** */

enum bop = {RISCV_BEQ, RISCV_BNE, RISCV_BLT, RISCV_BGE,
     ↪ RISCV_BLTU, RISCV_BGEU}
union clause ast =
  BTYPE : (bits(13), regidx, regidx, bop)

mapping encdec_bop : bop <-> bits(3) = {
  RISCV_BEQ <-> 0b000,
  RISCV_BNE <-> 0b001,
  RISCV_BLT <-> 0b100,
  RISCV_BGE <-> 0b101,
  RISCV_BLTU <-> 0b110,
  RISCV_BGEU <-> 0b111
}

mapping clause encdec = BTYPE(imm7_6 @ imm5_0 @ imm7_5_0 @
     ↪ imm5_4_1 @ 0b0, rs2, rs1, op)
  <-> imm7_6 : bits(1) @ imm7_5_0 : bits(6) @ rs2 @ rs1 @
       ↪ encdec_bop(op) @ imm5_4_1 : bits(4) @ imm5_0 :
       ↪ bits(1) @ 0b1100011
  if (rs1[4] == bitzero & rs2[4] == bitzero)

function clause execute (BTYPE(imm, rs2, rs1, op)) = {
  let rs1_val = X(rs1);
  let rs2_val = X(rs2);
  common_leakage(rs1_val, rs2_val);
  let taken : bool = match op {
    RISCV_BEQ => rs1_val == rs2_val,
    RISCV_BNE => rs1_val != rs2_val,
    RISCV_BLT => rs1_val <_s rs2_val,
    RISCV_BGE => rs1_val >=_s rs2_val,
    RISCV_BLTU => rs1_val <_u rs2_val,
    RISCV_BGEU => rs1_val >=_u rs2_val
  };

  overwrite_leakage(0b00000, 0
       ↪ b00000000000000000000000000000000);

  let t : xlenbits = PC + EXTS(imm);
  if (t[1 .. 0] != 0b00) then
    return RETIRE_FAIL;
  if taken then { set_next_pc(t); };
  return RETIRE_SUCCESS
}

/* ***************************** */

enum iop = {RISCV_ADDI, RISCV_SLTI, RISCV_SLTIU, RISCV_XORI,
     ↪  RISCV_ORI, RISCV_ANDI}
union clause ast =
  ITYPE : (bits(12), regidx, regidx, iop)

mapping encdec_iop : iop <-> bits(3) = {
  RISCV_ADDI <-> 0b000,
  RISCV_SLTI <-> 0b010,
  RISCV_SLTIU <-> 0b011,
  RISCV_ANDI <-> 0b111,
  RISCV_ORI <-> 0b110,
  RISCV_XORI <-> 0b100
}

mapping clause encdec = ITYPE(imm, rs1, rd, op)
  <-> imm @ rs1 @ encdec_iop(op) @ rd @ 0b0010011
  if (rs1[4] == bitzero) & (rd[4] == bitzero)

function clause execute (ITYPE (imm, rs1, rd, op)) = {
  let rs1_val = X(rs1);
```

```
  let rs2_val = X(0b0 @ imm[3 .. 0]);
  common_leakage(rs1_val, rs2_val);
  let immext : xlenbits = EXTS(imm);
  let result : xlenbits = match op {
    RISCV_ADDI => rs1_val + immext,
    RISCV_SLTI =>
      EXTZ(bool_to_bits(rs1_val <_s immext)),
    RISCV_SLTIU =>
      EXTZ(bool_to_bits(rs1_val <_u immext)),
    RISCV_ANDI => rs1_val & immext,
    RISCV_ORI => rs1_val | immext,
    RISCV_XORI => rs1_val ^ immext
  };
  overwrite_leakage(rd, result);
  X(rd) = result;
  RETIRE_SUCCESS
}

/* ***************************** */

enum sop = {RISCV_SLLI, RISCV_SRLI, RISCV_SRAI}
union clause ast =
  SHIFTIOP : (bits(6), regidx, regidx, sop)

mapping encdec_sop : sop <-> bits(3) = {
  RISCV_SLLI <-> 0b001,
  RISCV_SRLI <-> 0b101,
  RISCV_SRAI <-> 0b101
}

mapping clause encdec = SHIFTIOP(shamt, rs1, rd, RISCV_SLLI)
  <-> 0b000000 @ shamt @ rs1 @ 0b001 @ rd @ 0b0010011
  if (shamt[5] == bitzero) &(rs1[4] == bitzero) & (rd[4] ==
      ↪ bitzero)
mapping clause encdec = SHIFTIOP(shamt, rs1, rd, RISCV_SRLI)
  <-> 0b000000 @ shamt @ rs1 @ 0b101 @ rd @ 0b0010011
  if (shamt[5] == bitzero) &(rs1[4] == bitzero) & (rd[4] ==
      ↪ bitzero)
mapping clause encdec = SHIFTIOP(shamt, rs1, rd, RISCV_SRAI)
  <-> 0b010000 @ shamt @ rs1 @ 0b101 @ rd @ 0b0010011
  if (shamt[5] == bitzero) &(rs1[4] == bitzero) & (rd[4] ==
      ↪ bitzero)

function clause execute (SHIFTIOP(shamt, rs1, rd, op)) = {
  let rs1_val = X(rs1);
  let rs2_val = X(0b0 @ shamt[3..0]);
  common_leakage(rs1_val, rs2_val);
  /* the decoder guard ensures that shamt[5] = 0 for RV32E
      ↪ */
  let result : xlenbits = match op {
    RISCV_SLLI => if sizeof(xlen) == 32
                  then rs1_val << shamt[4..0]
                  else rs1_val << shamt,
    RISCV_SRLI => if sizeof(xlen) == 32
                  then rs1_val >> shamt[4..0]
                  else rs1_val >> shamt,
    RISCV_SRAI => if sizeof(xlen) == 32
                  then shift_right_arith32(rs1_val, shamt
                      ↪ [4..0])
                  else shift_right_arith64(rs1_val, shamt)};
  overwrite_leakage(rd, result);
  X(rd) = result;
  RETIRE_SUCCESS
}

/* ***************************** */

enum rop = {RISCV_ADD, RISCV_SUB, RISCV_SLL, RISCV_SLT,
            RISCV_SLTU, RISCV_XOR, RISCV_SRL, RISCV_SRA,
            RISCV_OR, RISCV_AND}
union clause ast = RTYPE : (regidx, regidx, regidx, rop)

mapping clause encdec = RTYPE(rs2, rs1, rd, RISCV_ADD)
  <-> 0b0000000 @ rs2 @ rs1 @ 0b000 @ rd @ 0b0110011
  if (rs1[4] == bitzero) & (rs2[4] == bitzero) & (rd[4] ==
      ↪ bitzero)

mapping clause encdec = RTYPE(rs2, rs1, rd, RISCV_SLT)
  <-> 0b0000000 @ rs2 @ rs1 @ 0b010 @ rd @ 0b0110011
  if (rs1[4] == bitzero) & (rs2[4] == bitzero) & (rd[4] ==
      ↪ bitzero)
mapping clause encdec = RTYPE(rs2, rs1, rd, RISCV_SLTU)
  <-> 0b0000000 @ rs2 @ rs1 @ 0b011 @ rd @ 0b0110011
  if (rs1[4] == bitzero) & (rs2[4] == bitzero) & (rd[4] ==
      ↪ bitzero)
mapping clause encdec = RTYPE(rs2, rs1, rd, RISCV_AND)
  <-> 0b0000000 @ rs2 @ rs1 @ 0b111 @ rd @ 0b0110011
  if (rs1[4] == bitzero) & (rs2[4] == bitzero) & (rd[4] ==
      ↪ bitzero)
mapping clause encdec = RTYPE(rs2, rs1, rd, RISCV_OR)
  <-> 0b0000000 @ rs2 @ rs1 @ 0b110 @ rd @ 0b0110011
  if (rs1[4] == bitzero) & (rs2[4] == bitzero) & (rd[4] ==
      ↪ bitzero)
mapping clause encdec = RTYPE(rs2, rs1, rd, RISCV_XOR)
  <-> 0b0000000 @ rs2 @ rs1 @ 0b100 @ rd @ 0b0110011
  if (rs1[4] == bitzero) & (rs2[4] == bitzero) & (rd[4] ==
      ↪ bitzero)
mapping clause encdec = RTYPE(rs2, rs1, rd, RISCV_SLL)
  <-> 0b0000000 @ rs2 @ rs1 @ 0b001 @ rd @ 0b0110011
  if (rs1[4] == bitzero) & (rs2[4] == bitzero) & (rd[4] ==
      ↪ bitzero)
mapping clause encdec = RTYPE(rs2, rs1, rd, RISCV_SRL)
  <-> 0b0000000 @ rs2 @ rs1 @ 0b101 @ rd @ 0b0110011
  if (rs1[4] == bitzero) & (rs2[4] == bitzero) & (rd[4] ==
      ↪ bitzero)
mapping clause encdec = RTYPE(rs2, rs1, rd, RISCV_SUB)
  <-> 0b0100000 @ rs2 @ rs1 @ 0b000 @ rd @ 0b0110011
  if (rs1[4] == bitzero) & (rs2[4] == bitzero) & (rd[4] ==
      ↪ bitzero)
mapping clause encdec = RTYPE(rs2, rs1, rd, RISCV_SRA)
  <-> 0b0100000 @ rs2 @ rs1 @ 0b101 @ rd @ 0b0110011
  if (rs1[4] == bitzero) & (rs2[4] == bitzero) & (rd[4] ==
      ↪ bitzero)
```

```
function clause execute (RTYPE(rs2, rs1, rd, op)) = {
  let rs1_val = X(rs1);
  let rs2_val = X(rs2);
  common_leakage(rs1_val, rs2_val);

  let result : xlenbits = match op {
    RISCV_ADD => rs1_val + rs2_val,
    RISCV_SLT => EXTZ(bool_to_bits(rs1_val <_s rs2_val)),
    RISCV_SLTU => EXTZ(bool_to_bits(rs1_val <_u rs2_val)),
    RISCV_AND => rs1_val & rs2_val,
    RISCV_OR => rs1_val | rs2_val,
    RISCV_XOR => rs1_val ^ rs2_val,
    RISCV_SLL => if sizeof(xlen) == 32
                 then rs1_val << (rs2_val[4..0])
                 else rs1_val << (rs2_val[5..0]),
    RISCV_SRL => if sizeof(xlen) == 32
                 then rs1_val >> (rs2_val[4..0])
                 else rs1_val >> (rs2_val[5..0]),
    RISCV_SUB => rs1_val - rs2_val,
    RISCV_SRA => if sizeof(xlen) == 32
                 then shift_right_arith32(rs1_val, rs2_val
                      ↪ [4..0])
                 else shift_right_arith64(rs1_val, rs2_val
                      ↪ [5..0])
  };
  // leak(X(rd), result);
  overwrite_leakage(rd, result);

  X(rd) = result;
  RETIRE_SUCCESS
}

/* ***************************** */

enum word_width = {BYTE, HALF, WORD, DOUBLE}
union clause ast = LOAD :
  (bits(12), regidx, regidx, bool, word_width, bool, bool)

mapping clause encdec = LOAD(imm, rs1, rd, is_unsigned, size
     ↪ , false, false)
  if ((word_width_bytes(size) < sizeof(xlen_bytes)) | (
       ↪ not_bool(is_unsigned) & word_width_bytes(size) <=
       ↪ sizeof(xlen_bytes))) & (rs1[4] == bitzero) & (rd
       ↪ [4] == bitzero)
  <-> imm @ rs1 @ bool_bits(is_unsigned) @ size_bits(size) @
       ↪  rd @ 0b0000011
  if ((word_width_bytes(size) < sizeof(xlen_bytes)) | (
       ↪ not_bool(is_unsigned) & word_width_bytes(size) <=
       ↪ sizeof(xlen_bytes))) & (rs1[4] == bitzero) & (rd
       ↪ [4] == bitzero)

function aligned(vaddr : xlenbits, width : word_width) ->
     ↪ bool =
  { width == BYTE | (width == HALF & vaddr[0] == bitzero) |
       ↪ (width == WORD & vaddr[1 .. 0] == 0b00) }

val load_leakage : (xlenbits, xlenbits, xlenbits, xlenbits)
  -> unit effect {rreg, wreg, leakage}
function load_leakage(rs1_val : xlenbits, rs2_val : xlenbits
     ↪ , addr: xlenbits, req_data: xlenbits) = {
  leak(rf_pA, rf_pB, rs1_val, rs2_val);
  leak(rf_pA, rf_pB, mem_last_addr, mem_last_read);
  leak(addr, req_data, mem_last_addr);
  rf_pA = rs1_val;
  rf_pB = rs2_val;
  mem_last_read = req_data;
  mem_last_addr = addr;
}

function clause execute(LOAD(imm, rs1, rd, is_unsigned,
     ↪ width, aq, rl)) = {
  let offset : xlenbits = EXTS(imm);
  let rs1_val = X(rs1);
  let rs2_val = X(0b0 @ imm[3 .. 0]);
  let addr = rs1_val + offset;
  let req_addr = addr[(sizeof(xlen) - 1) .. 2] @ 0b00;
  let req_data = read_mem(Read_plain, sizeof(xlen), req_addr
       ↪ , 4);
  load_leakage(rs1_val, rs2_val, addr, req_data);
  let req_byte : bits(8) = match (addr[1 .. 0]) {
    0b00 => req_data[ 7 .. 0],
    0b01 => req_data[15 .. 8],
    0b10 => req_data[23 .. 16],
    0b11 => req_data[31 .. 24]};
  let req_half : bits(16) = match (addr[1]) {
    bitzero => req_data[15 .. 0],
    bitone => req_data[31 .. 16]};
  match (width, addr[1 .. 0]) {
    (BYTE, _) => process_load(rd, addr, req_byte,
         ↪ is_unsigned),
    (HALF, 0b00) => process_load(rd, addr, req_half,
         ↪ is_unsigned),
    (HALF, 0b10) => process_load(rd, addr, req_half,
         ↪ is_unsigned),
    (WORD, 0b00) => process_load(rd, addr, req_data,
         ↪ is_unsigned),
    (_, _) => RETIRE_FAIL // takes care of misaligned}
}

/* ***************************** */

union clause ast = STORE :
  (bits(12), regidx, regidx, word_width, bool, bool)

mapping clause encdec = STORE(imm7 @ imm5, rs2, rs1, size,
     ↪ false, false)
  if (word_width_bytes(size) <= sizeof(xlen_bytes)) & (rs1
       ↪ [4] == bitzero) & (rs2[4] == bitzero)
  <-> imm7 : bits(7) @ rs2 @ rs1 @ 0b0 @ size_bits(size) @
       ↪ imm5 : bits(5) @ 0b0100011
  if (word_width_bytes(size) <= sizeof(xlen_bytes)) & (rs1
       ↪ [4] == bitzero) & (rs2[4] == bitzero)
```

```
function clause execute (STORE(imm, rs2, rs1, width, aq, rl) ↪ ) = {
  let offset : xlenbits = EXTS(imm);
  let rs1_val = X(rs1);
  let rs2_val = X(rs2);
  common_leakage(rs1_val, rs2_val);
  let addr = rs1_val + offset;
  // address comptation and register file access leakage
  leak(mem_last_addr, addr);
  mem_last_addr = addr;
  if aligned(addr, width) then {
    let result = rs2_val;
    overwrite_leakage(0b00000, result);
    let success : bool = match(width) {
      BYTE => write_mem(Write_plain, sizeof(xlen), addr, 1, ↪ result[7..0]),
      HALF => write_mem(Write_plain, sizeof(xlen), addr, 2, ↪ result[15..0]),
      WORD => write_mem(Write_plain, sizeof(xlen), addr, 4, ↪ result),
      _ => false};
    if success then {RETIRE_SUCCESS} else {RETIRE_FAIL}
  } else { RETIRE_FAIL }
}

/* ***************************** */

mapping clause encdec = ILLEGAL(s) <-> s
function clause execute (ILLEGAL(s)) =
  { return RETIRE_FAIL }
```

7.2 IBEX Configuration

In the following, we give the configuration file that specifies the mapping and normal operating conditions simultaneously.

Listing 7: IBEX configuration file

```
// Power Contract for IBEX
//
// Copyright (c) 2020-2022 - TUHH, TU Graz
//
// All rights reserved.
//
// This software has received funding from the Federal Ministry of Education and Research (BMBF) as part of the VE-Jupiter project grant 16ME0231K.
//
// This work was supported by the Austrian Research Promotion Agency (FFG) through the FERMION project (grant number 867542).
//
// Redistribution and use in source and binary forms,
// with or without modification, are permitted provided
// that the following conditions are met:
// 1. Redistributions of source code must retain the
// above copyright notice, this list of conditions
// and the following disclaimer.
// 2. Redistributions in binary form must reproduce the
// above copyright notice, this list of conditions
// and the following disclaimer in the documentation
// and/or other materials provided with the
// distribution.
//
// THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND
// CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR
// IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED
// TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY
// AND FITNESS FOR A PARTICULAR PURPOSE ARE
// DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR
// CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT,
// INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
// DAMAGES (INCLUDING, BUT NOT LIMITED TO,
// PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS
// OF USE, DATA, OR PROFITS; OR BUSINESS
// INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF
// LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY,
// OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
// ARISING IN ANY WAY OUT OF THE USE OF THIS
// SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
// SUCH DAMAGE.

////////////////////////////////////////////////////
// This file contains the configuration of our tool.
// It specifies
// - the state modeled in the contract (registers, memory, ↪ leakage state)
// - the registers of the IBEX processor (registers and ↪ memory)
// - a mapping between the states
// - which states may contain sensitive data
// - conditions which have to hold before/during execution ↪ of an instruction
// - which HW and contract state is printed in ↪ counterexamples
////////////////////////////////////////////////////

////////////////////////////////////////////////////
// specification of architectural registers and HW/CT ↪ mapping
////////////////////////////////////////////////////
// PC
contract register PC BitVec 32
hardware public u_ibex_core.pc_id
mapping register PC u_ibex_core.pc_id

// next PC
hardware public u_ibex_core.if_stage_i.gen_prefetch_buffer. ↪ prefetch_buffer_i.fifo_i.instr_addr_q
contract register nextPC BitVec 32
mapping register nextPC u_ibex_core.if_stage_i. ↪ gen_prefetch_buffer.prefetch_buffer_i.fifo_i. ↪ instr_addr_q

// REGISTERS
contract register x1 BitVec 32
```

```
contract register x2 BitVec 32
contract register x3 BitVec 32
contract register x4 BitVec 32
contract register x5 BitVec 32
contract register x6 BitVec 32
contract register x7 BitVec 32
contract register x8 BitVec 32
contract register x9 BitVec 32
contract register x10 BitVec 32
contract register x11 BitVec 32
contract register x12 BitVec 32
contract register x13 BitVec 32
contract register x14 BitVec 32
contract register x15 BitVec 32
hardware variable register_file_i.rf_reg_q[1]
hardware variable register_file_i.rf_reg_q[2]
hardware variable register_file_i.rf_reg_q[3]
hardware variable register_file_i.rf_reg_q[4]
hardware variable register_file_i.rf_reg_q[5]
hardware variable register_file_i.rf_reg_q[6]
hardware variable register_file_i.rf_reg_q[7]
hardware variable register_file_i.rf_reg_q[8]
hardware variable register_file_i.rf_reg_q[9]
hardware variable register_file_i.rf_reg_q[10]
hardware variable register_file_i.rf_reg_q[11]
hardware variable register_file_i.rf_reg_q[12]
hardware variable register_file_i.rf_reg_q[13]
hardware variable register_file_i.rf_reg_q[14]
hardware variable register_file_i.rf_reg_q[15]
mapping register x1 register_file_i.rf_reg_q[1]
mapping register x2 register_file_i.rf_reg_q[2]
mapping register x3 register_file_i.rf_reg_q[3]
mapping register x4 register_file_i.rf_reg_q[4]
mapping register x5 register_file_i.rf_reg_q[5]
mapping register x6 register_file_i.rf_reg_q[6]
mapping register x7 register_file_i.rf_reg_q[7]
mapping register x8 register_file_i.rf_reg_q[8]
mapping register x9 register_file_i.rf_reg_q[9]
mapping register x10 register_file_i.rf_reg_q[10]
mapping register x11 register_file_i.rf_reg_q[11]
mapping register x12 register_file_i.rf_reg_q[12]
mapping register x13 register_file_i.rf_reg_q[13]
mapping register x14 register_file_i.rf_reg_q[14]
mapping register x15 register_file_i.rf_reg_q[15]

contract opcode op BitVec 32
// instruction bits for the instruction whose last execution
    ↪ cycle is this cycle
// only true if the assertion for the valid_d is present
hardware opcode u_ibex_core.instr_rdata_id

// memory request name_contract name_hardware
contract register read_val_1 BitVec 32
contract register read_addr_1 BitVec 32
hardware variable data_rdata_i

memory raddr u_ibex_core.load_store_unit_i.adder_result_ex_i
    ↪ read_addr_1
memory rdata data_rdata_i read_val_1
memory req data_req_o
memory gnt data_gnt_i
memory ack data_rvalid_i
memory we data_we_o

////////////////////////////////////////////////////
// Non-regport signals that must be constrained
////////////////////////////////////////////////////
// fetching next instruction was successful, ready to
    ↪ continue execution
hardware const@end-1 u_ibex_core.if_stage_i.fetch_valid 0b1
hardware const@pre u_ibex_core.if_stage_i.fetch_valid 0b1
// do not load a new instruction until last cycle
hardware const@start:end-1 u_ibex_core.id_stage_i.
    ↪ id_in_ready_o 0b0
// make sure that nothing retires before the end of the last
    ↪ cycle
hardware const@start:end-1 u_ibex_core.instr_id_done 0b0
// make sure that an instruction has its last cycle in our
    ↪ last cycle
hardware const@end-1 u_ibex_core.instr_id_done 0b1
hardware const@pre u_ibex_core.instr_id_done 0b1
hardware const@start u_ibex_core.id_stage_i.decoder_i.
    ↪ illegal_insn 0b0

////////////////////////////////////////////////////
// important signals that must be constrained
////////////////////////////////////////////////////
// never trigger a reset of the core
hardware public rst_ni
hardware const@pre: rst_ni 0b1
// make sure that initially, the ID FSM is in state
    ↪ instr_first_cycle_i
// this means that we look at the case where we started
    ↪ executing in 0th cycle
hardware public u_ibex_core.id_stage_i.id_fsm_q
hardware const@start u_ibex_core.id_stage_i.id_fsm_q 0b0
// no compressed (valid or invalid) instructions at the
    ↪ output of instruction fetch stage
hardware public u_ibex_core.if_stage_i.instr_new_id_q
hardware public u_ibex_core.if_stage_i.
    ↪ instr_is_compressed_id_o
hardware const@start u_ibex_core.if_stage_i.
    ↪ instr_is_compressed_id_o 0b0
hardware public u_ibex_core.if_stage_i.illegal_c_insn_id_o
hardware const@start u_ibex_core.if_stage_i.
    ↪ illegal_c_insn_id_o 0b0
// this is a hidden assumption made by IBEX developers
hardware public u_ibex_core.if_stage_i.instr_rdata_alu_id_o
hardware public u_ibex_core.if_stage_i.instr_rdata_id_o
// this encodes pre cycle and first cycle assumptions
hardware equiv@pre:start+1 u_ibex_core.if_stage_i.
    ↪ instr_rdata_id_o u_ibex_core.if_stage_i.
    ↪ instr_rdata_alu_id_o
```

```
////////////////////////////////////////////////////
// annotation of input ports of ibex_top
////////////////////////////////////////////////////
hardware public clk_i
hardware public ram_cfg_i
hardware public test_en_i
hardware public hart_id_i
hardware public boot_addr_i
// Instruction memory interface
hardware public instr_gnt_i
hardware public instr_rvalid_i
hardware public instr_rdata_i
hardware public instr_err_i
hardware public data_gnt_i
hardware public data_rvalid_i
hardware public data_err_i
hardware const@pre: data_err_i 0b0
// interrupts
hardware public irq_software_i
hardware public irq_timer_i
hardware public irq_external_i
hardware public irq_fast_i
hardware public irq_nm_i
// disabling interrupts
hardware const@pre: irq_software_i 0b0
hardware const@pre: irq_timer_i 0b0
hardware const@pre: irq_external_i 0b0
hardware const@pre: irq_fast_i 0b000000000000000
hardware const@pre: irq_nm_i 0b0
// core debug
hardware public debug_req_i
hardware const@pre: debug_req_i 0b0
hardware public fetch_enable_i
hardware public scan_rst_ni

//////////////////////////////////////
// annotate internal state of ibex
//////////////////////////////////////
hardware public core_busy_q
hardware public u_ibex_core.instr_fetch_err
hardware public u_ibex_core.instr_fetch_err_plus2
// instructions in the prefetch fifo
hardware public u_ibex_core.if_stage_i.instr_valid_id_q
hardware public u_ibex_core.if_stage_i.instr_rdata_c_id_o
hardware public u_ibex_core.if_stage_i.gen_prefetch_buffer.
    ↪ prefetch_buffer_i.fifo_i.rdata_q0
hardware public u_ibex_core.if_stage_i.gen_prefetch_buffer.
    ↪ prefetch_buffer_i.fifo_i.rdata_q1
hardware public u_ibex_core.if_stage_i.gen_prefetch_buffer.
    ↪ prefetch_buffer_i.fifo_i.rdata_q2
hardware public u_ibex_core.if_stage_i.gen_prefetch_buffer.
    ↪ prefetch_buffer_i.fifo_i.err_q
hardware public u_ibex_core.if_stage_i.gen_prefetch_buffer.
    ↪ prefetch_buffer_i.fifo_i.valid_q
hardware public u_ibex_core.if_stage_i.gen_prefetch_buffer.
    ↪ prefetch_buffer_i.rdata_pmp_err_q
hardware public u_ibex_core.if_stage_i.gen_prefetch_buffer.
    ↪ prefetch_buffer_i.discard_req_q
hardware public u_ibex_core.if_stage_i.gen_prefetch_buffer.
    ↪ prefetch_buffer_i.branch_discard_q
hardware public u_ibex_core.if_stage_i.gen_prefetch_buffer.
    ↪ prefetch_buffer_i.rdata_outstanding_q
hardware public u_ibex_core.if_stage_i.gen_prefetch_buffer.
    ↪ prefetch_buffer_i.fetch_addr_q
hardware public u_ibex_core.if_stage_i.gen_prefetch_buffer.
    ↪ prefetch_buffer_i.stored_addr_q

// instruction decode
hardware public u_ibex_core.id_stage_i.controller_i.
    ↪ ctrl_fsm_cs
hardware public u_ibex_core.id_stage_i.controller_i.
    ↪ load_err_q
hardware public u_ibex_core.id_stage_i.controller_i.
    ↪ store_err_q
hardware public u_ibex_core.id_stage_i.controller_i.
    ↪ exc_req_q
hardware public u_ibex_core.id_stage_i.branch_set_raw
hardware const@start u_ibex_core.id_stage_i.branch_set_raw 0
    ↪ b0
hardware public u_ibex_core.id_stage_i.
    ↪ branch_jump_set_done_q
hardware public u_ibex_core.load_store_unit_i.data_we_q

// since data_pmp_err_i is 0, this should not be 1
hardware public u_ibex_core.load_store_unit_i.pmp_err_q
hardware const@start u_ibex_core.load_store_unit_i.pmp_err_q
    ↪  0b0
// since data_err_i is never 1 and pmp_err_q is also not 1,
    ↪ this must be 0
hardware public u_ibex_core.load_store_unit_i.lsu_err_q
hardware const@start u_ibex_core.load_store_unit_i.lsu_err_q
    ↪  0b0
hardware public u_ibex_core.load_store_unit_i.
    ↪ handle_misaligned_q
hardware public u_ibex_core.id_stage_i.controller_i.
    ↪ illegal_insn_q
hardware public u_ibex_core.id_stage_i.controller_i.
    ↪ do_single_step_q
hardware public u_ibex_core.id_stage_i.controller_i.
    ↪ enter_debug_mode_prio_q
// LSU register handling misaligned memory accesses which
    ↪ are not allowed by the contract
hardware public u_ibex_core.load_store_unit_i.rdata_q
contract leakagestate mem_last_read BitVec 32
// Must be idle when new instruction reaches ID/EX
hardware public u_ibex_core.load_store_unit_i.ls_fsm_cs
hardware const@start u_ibex_core.load_store_unit_i.ls_fsm_cs
    ↪  0b000
hardware variable u_ibex_core.load_store_unit_i.addr_last_q
hardware variable u_ibex_core.load_store_unit_i.
    ↪ rdata_offset_q
contract leakagestate mem_last_addr BitVec 32
```

```
mapping leakagestate mem_last_addr u_ibex_core.
    ↪ load_store_unit_i.addr_last_q
mapping leakagestate mem_last_addr u_ibex_core.
    ↪ load_store_unit_i.rdata_offset_q
hardware public u_ibex_core.load_store_unit_i.data_type_q
hardware public u_ibex_core.load_store_unit_i.
    ↪ data_sign_ext_q

// system registers
hardware public u_ibex_core.cs_registers_i.mie_q
hardware public u_ibex_core.cs_registers_i.mtval_q
hardware public u_ibex_core.cs_registers_i.mcause_q
hardware public u_ibex_core.cs_registers_i.mscratch_q
hardware public u_ibex_core.cs_registers_i.dscratch0_q
hardware public u_ibex_core.cs_registers_i.dscratch1_q
hardware public u_ibex_core.cs_registers_i.mstack_q
hardware public u_ibex_core.cs_registers_i.mstack_cause_q
hardware public u_ibex_core.cs_registers_i.mstack_epc_q
hardware public u_ibex_core.cs_registers_i.mstatus_q
hardware public u_ibex_core.cs_registers_i.dcsr_q
hardware const@start u_ibex_core.cs_registers_i.dcsr_q 0
    ↪ b00000000000000000000000000000000
hardware public u_ibex_core.cs_registers_i.mhpmcounter[0]
hardware public u_ibex_core.cs_registers_i.mhpmcounter[1]
hardware public u_ibex_core.cs_registers_i.mhpmcounter[2]
hardware public u_ibex_core.cs_registers_i.mcountinhibit
hardware public u_ibex_core.csr_depc
hardware public u_ibex_core.csr_mtvec
hardware public u_ibex_core.csr_mepc
hardware public u_ibex_core.dummy_instr_en
hardware public u_ibex_core.dummy_instr_mask
hardware public u_ibex_core.data_ind_timing
hardware public u_ibex_core.icache_enable
hardware public u_ibex_core.debug_mode
hardware public u_ibex_core.priv_mode_id
hardware public u_ibex_core.nmi_mode

contract leakagestate rf_pA BitVec 32
contract leakagestate rf_pB BitVec 32
```