

# Synthesis of Sound and Precise Leakage Contracts for Open-Source RISC-V Processors

Anonymous Author(s)

## Abstract

Leakage contracts have been proposed as a new security abstraction at the instruction set architecture level. Leakage contracts aim to capture the information that processors may leak via microarchitectural side channels. Recently, the first verification tools have emerged to check whether a processor satisfies a given contract. However, coming up with a contract that is both sound and precise for a given processor is challenging, time-consuming, and error-prone, as it requires in-depth knowledge of the timing side channels introduced by microarchitectural optimizations.

In this paper, we address this challenge by proposing LEASYN, the first tool for automatically synthesizing leakage contracts that are both *sound* and *precise* for processor designs at register-transfer level. Starting from a user-provided contract template that captures the space of possible contracts, LEASYN automatically constructs a contract, alternating between contract synthesis, which ensures precision based on an empirical characterization of the processor’s leaks, and contract verification, which ensures soundness.

Using LEASYN, we automatically synthesize contracts for six open-source RISC-V CPUs for a variety of contract templates. Our experiments indicate that the contracts synthesized by LEASYN are more precise (i.e., represent the actual leaks in the target processor more faithfully) than contracts derived manually in prior work.

## 1 Introduction

Microarchitectural attacks [1–15] rely on subtle but measurable hardware side effects on a computation’s execution time to compromise otherwise secure programs. The exploited side effects result from optimizations—such as caching and speculative execution—implemented in a CPU’s microarchitecture. Defending against such attacks is challenging since they are “invisible” at the level of the instruction set architecture (ISA), the traditional hardware-software interface familiar to programmers.

*Leakage contracts* [16] have recently been proposed as a new security abstraction at ISA-level to fill this gap and to serve as a rigorous foundation for writing programs that are resistant against microarchitectural attacks. In a nutshell, a leakage contract specifies at the ISA-level what information a program may leak via microarchitectural side channels. Leakage contracts have been successfully applied to reason about the security of hardware [17, 18] and software [19, 20] against microarchitectural leaks.

To prevent microarchitectural attacks, programmers can use leakage contracts as a guideline for secure programming by ensuring that secret information does not influence any of the leakage sources identified in the contracts. To be useful for this, however, leakage contracts should be *sound*, i.e., they should capture all leaks in the underlying processor, and *precise*, i.e., they should expose as part of the contract only the information that is actually leaked. Unsound contracts are problematic since they could lead to insecure code. At the same time, imprecise contracts may unnecessarily rule

out programs as “insecure” and introduce performance overheads, as they lead to unnecessarily defensive code.

Deriving a sound and precise leakage contract for a given processor is challenging. Manually constructing a leakage contract for a modern processor is a time-consuming [21] and error-prone endeavour: it requires a deep understanding of the processor’s microarchitecture and the leakage introduced by its optimizations. Automatic synthesis of leakage contracts promises to address the impracticalities of manually constructing a contract [22–24]. However, existing approaches for synthesis are limited: The approach by Mohr et al. [22] can synthesize a precise contract directly from a register-transfer level (RTL) processor design, but it cannot guarantee the contract’s soundness, i.e., the contract may miss leaks, compromising its usefulness for secure programming. The approach by Dinesh et al. [23, 24], on the other hand, synthesizes sound but (very) imprecise contracts, e.g., the synthesized contracts entirely rule out the use of any branch or memory instructions. Wang et al. [17] have proposed LEAVE, a tool that can verify the soundness of leakage contracts against RTL designs, and they have used LEAVE to synthesize sound contracts with a human in the loop. As we show in § 6, the resulting contracts are still fairly imprecise.

In this paper, we propose a methodology to automatically synthesize leakage contracts that are *sound* and *precise* given an RTL processor design. Next, we describe our main contributions.

**Synthesis methodology.** We propose a methodology (§ 4) to synthesize sound and precise contracts given an RTL processor design. Our methodology starts from (a) the RTL design of the target processor, and (b) a user-provided *contract template* that consists of a set of *contract atoms*, each capturing a potential instruction-level leak. Crucially, the template defines the search space for contract synthesis and any subset of the template’s atoms is a candidate contract. Figure 1 illustrates our approach, which consists of three phases:

- (1) *Empirical Leakage Characterization:* We start by characterizing a processor’s leakage based on a set of test cases (§ 4.1), where a test case consists of a pair of programs with their corresponding inputs. These test cases are automatically generated based on the given contract template, and they aim at exercising the processor to identify leaks. The test cases are classified in terms of attacker distinguishability, i.e., whether the attacker can distinguish the pair of executions or not, and template distinguishability, i.e., which sets of atoms from the template distinguish each test case.
- (2) *Contract Synthesis:* We build a candidate contract for the target processor by alternating between a synthesis phase (§ 4.2) and a verification phase (§ 4.3). We start by synthesizing the most precise candidate contract that captures all leaks exposed in the test cases, based on the attacker and template distinguishability relations. To ensure that the candidate contract is sound, i.e., it captures all leaks, even those not exercised by the initial test cases, we perform a bounded verification step. Whenever we

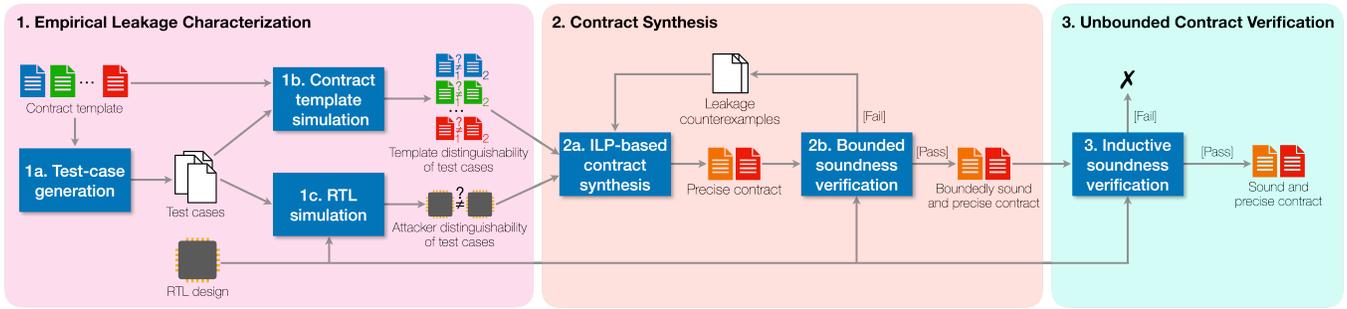


Figure 1: Overview of LEASYN synthesis approach

find a leakage counterexample, i.e., an attacker-distinguishable test case that is contract indistinguishable under the current candidate contract, this counterexample is added to the set of test cases, and a new contract candidate is synthesized.

- (3) *Unbounded Contract Verification*: When we cannot find further counterexamples with bounded verification, we perform a final unbounded verification step to ensure that indeed all possible leaks in the target processor are captured by the candidate contract (§ 4.4). If this passes, the contract is provably sound for the target processor, and we return it to the user.

In contrast to prior approaches for contract synthesis [22–24], our methodology ensures that the synthesized contract is both *sound*, i.e., it captures all leaks, and *precise*, i.e., it distinguishes as few attacker-indistinguishable test cases as possible.

**LEASYN synthesis tool.** We implement our methodology in LEASYN (§ 5), a tool for synthesizing leakage contracts from processor designs in Verilog. LEASYN implements all the steps of our methodology: it uses integer-linear programming for contract synthesis, a bounded model checker for bounded soundness verification, and the LEAVE contract verification tool [17] for unbounded soundness verification. LEASYN also implements a family of different contract templates capturing instruction-level leaks, which are implemented on top of the RISC-V Formal Interface [25] to simplify applying LEASYN to new targets.

**Evaluation.** We validate our methodology by synthesizing leakage contracts for six open-source RISC-V processors from three different core families (§ 6). We do so by (a) defining a variety of contract templates capturing different classes of instruction-level leaks (e.g., leaks through control flow, variable-latency instructions, and memory accesses), and (b) using LEASYN to synthesize the most precise and sound contracts satisfied by the studied target processors, against an attacker that observes when instructions retire. Our experiments confirm that LEASYN can successfully synthesize sound and precise leakage contracts for all of our targets in less than 48 hours. Furthermore, the contracts synthesized by LEASYN are significantly more precise than existing sound contracts derived either manually [17] or automatically [23].

**Summary.** To summarize, we make the following contributions:

- We propose a methodology for synthesizing sound and precise leakage contracts for RISC-V processors at RTL.
- We implement our methodology in a tool called LEASYN.

- We evaluate LEASYN on six open-source RISC-V processors, demonstrating that it can efficiently synthesize sound and precise contracts.
- We show that the contracts synthesized by LEASYN are significantly more precise than those derived in prior work.

## 2 Overview

We now present the core aspects of our approach with an example.

### 2.1 A simple processor

We start by introducing a simple ISA and a CPU implementing it.

**Instruction Set.** We consider a simple instruction set *ISA* that supports loading immediate values into registers, and performing (integer) divisions between them. The architectural state consists of  $n$  32-bit registers, representing signed integers, which we identify using  $R_1, \dots, R_n$ . It also contains a dedicated register PC holding the program counter. The instruction set consists of two instructions:

(1) The **load immediate** instruction `li RD, imm` sets the value of the destination register RD to `imm`. RD is a register identifier in  $R_1, \dots, R_n$  and `imm` is a 32-bit value.

(2) The **division** instruction `div RD, RS1, RS2` divides RS1 by RS2 and stores the result in RD. RD, RS1, RS2 are register identifiers in  $R_1, \dots, R_n$ . Inspired by the RISC-V ISA [26], upon a division by 0, i.e., when the value of RS2 is 0, the value of RD is set to `-1`.

**Processor.** We consider a simple two-stage pipelined implementation *IMPL* of *ISA*. The *fetch/decode stage* takes one cycle. It loads the current instruction at the address in PC, decodes it into instruction type and operands, and increments the program counter.

The *execute stage* works as follows:

- Load immediate instructions `li RD, imm` are executed in a single cycle by writing the value `imm` to destination register RD.
- Division instructions `div RD, RS1, RS2` are handled by a dedicated division unit. The unit takes 32 cycles to execute all divisions except when RS2 is 0 or 1, in which case it takes 1 cycle.

Whenever the execute stage takes more than one cycle, the *fetch/decode stage* gets stalled.

### 2.2 Modeling leaks with leakage contracts

Next, we show how leakage contracts can be used to capture—at ISA level—the leaks existing in the CPU *IMPL* from § 2.1.

**Leakage.** Executing instructions on *IMPL* can leak information about the value of some of the registers. For instance, consider an attacker

ATK that can observe at which cycles *IMPL* retires each instruction. In this case, the attacker can distinguish between the following pairs of program executions, where we represent each execution by a pair  $\langle p \mid \sigma \rangle$  consisting of a program  $p$  and an initial state  $\sigma$ :

- (1)  $\langle \text{div } R1, R2, R3 \mid \sigma_1 \rangle$  vs.  $\langle \text{div } R1, R2, R3 \mid \sigma_2 \rangle$ , for  $\sigma_1, \sigma_2$ , s.t.  $\sigma_1(R3) = 2$  and  $\sigma_2(R3) = 1$ . The two executions are attacker distinguishable since executing *div* takes 33 cycles (1 cycle for fetch/decode and 32 for the execution stage) starting from  $s_1$ , but only 2 cycles starting from  $s_2$ .
- (2)  $\langle \text{div } R1, R2, R3 \mid \sigma_1 \rangle$  vs.  $\langle \text{li } R1, 0 \times 0 \mid \sigma_3 \rangle$  for some  $\sigma_3$ , are also attacker distinguishable since *div* takes 33 cycles starting from  $s_1$  whereas *li* only takes 2 cycles.

The attacker ATK would, however, not be able to distinguish the following pairs of executions:

- (1)  $\langle \text{div } R1, R2, R3 \mid \sigma_1 \rangle$  vs.  $\langle \text{div } R1, R2, R3 \mid \sigma_4 \rangle$ , for  $\sigma_4$  with  $\sigma_4(R3) = 8$ , since *div* will take 33 cycles in both cases, and
- (2)  $\langle \text{div } R1, R2, R3 \mid \sigma_2 \rangle$  vs.  $\langle \text{li } R1, 0 \times 0 \mid \sigma_3 \rangle$ , since executing both *div* (since  $r_3$  is 0) and *li* take 2 cycles.

**Leakage contracts.** Securely programming *IMPL* requires understanding which program executions an attacker might distinguish to ensure that secret data is not leaked. Leakage contracts [16] provide a way of characterizing such leaks at the ISA level, thereby enabling secure programming. For this, contracts map architectural executions, *i.e.*, executions according to the *ISA*, to *contract traces* that capture what information might be leaked through side channels.

A contract is *sound* for a CPU (or, equivalently, the CPU satisfies the contract), if the contract captures *all* leaks in the CPU, that is, any two program executions that are distinguishable by ATK must result in different contract traces. For instance, *IMPL* satisfies a contract that exposes all the operands of executed *div* instructions, since distinguishable executions like the ones mentioned above would be mapped to different contract traces. While sound contracts are a prerequisite for secure programming, not all sound contracts are useful. Indeed, consider a contract that exposes the operands of *all* instructions. Despite being sound for *IMPL*, this contract is useless from a secure programming perspective, since it only admits programs that do not process any secrets! This contract is imprecise since it over-approximates the actual leaks in *IMPL*.

### 2.3 Synthesizing sound and precise contracts with LEASYN

Figure 1 depicts LEASYN’s approach to synthesize sound and precise contracts given an RTL processor. The approach takes as input:

- (1) The RTL design of the processor under synthesis *IMPL*.
- (2) A user-provided *contract template*  $\mathbb{T}$  consisting of a set of *contract atoms* (§ 3.2), each capturing potential instruction-level leaks. The template defines the search space for contract synthesis and any subset of its atoms is a candidate contract.

Based on these inputs, LEASYN outputs a contract *CTR* that is (a) *sound*, *i.e.*, it captures *all* leaks in the target CPU *IMPL*, and (b) *as precise as possible*, *i.e.*, among all sound contracts in the template it distinguishes the fewest of the attacker-indistinguishable test cases explored during synthesis. Next, we describe the contract template for our example, and then the three phases of our approach.

**Contract template.** We consider a contract template that can expose as part of the contract trace three kinds of information: (1) the

value  $\llbracket \text{imm} \rrbracket$  of the immediate *imm* in load-immediate instructions, (2) the destination and source registers used in *div* instructions (denoted as  $\llbracket \text{RD} \rrbracket$ ,  $\llbracket \text{RS1} \rrbracket$ ,  $\llbracket \text{RS2} \rrbracket$ ), and (3) the values of these registers (denoted as  $\llbracket \text{Reg}[\text{RD}] \rrbracket$ ,  $\llbracket \text{Reg}[\text{RS1}] \rrbracket$ , and  $\llbracket \text{Reg}[\text{RS2}] \rrbracket$ ).

Furthermore, an observation corresponding to the instruction type  $ITYPE = \{\text{li}, \text{div}\}$  can be added to the trace whenever an instruction of that type is executed.

Therefore, the set of atoms in our template is as follows:

$$\begin{aligned} \text{Atoms} = \{ & (\text{li}, \langle \text{"li"}, - \rangle), (\text{div}, \langle \text{"div"}, - \rangle), (\text{li}, \langle \text{"imm"}, \llbracket \text{imm} \rrbracket \rangle), \\ & (\text{div}, \langle \text{"RD"}, \llbracket \text{RD} \rrbracket \rangle), (\text{div}, \langle \text{"Reg}[\text{RD}]", \llbracket \text{Reg}[\text{RD}] \rrbracket \rangle), \\ & (\text{div}, \langle \text{"RS1"}, \llbracket \text{RS1} \rrbracket \rangle), (\text{div}, \langle \text{"Reg}[\text{RS1}]", \llbracket \text{Reg}[\text{RS1}] \rrbracket \rangle), \\ & (\text{div}, \langle \text{"RS2"}, \llbracket \text{RS2} \rrbracket \rangle), (\text{div}, \langle \text{"Reg}[\text{RS2}]", \llbracket \text{Reg}[\text{RS2}] \rrbracket \rangle) \} \end{aligned}$$

For instance, atom  $(\text{li}, \langle \text{"li"}, - \rangle)$  adds to the trace the tuple  $\langle \text{"li"}, - \rangle$  whenever a load immediate instruction is executed, whereas atom  $(\text{div}, \langle \text{"Reg}[\text{RS2}]", \llbracket \text{Reg}[\text{RS2}] \rrbracket \rangle)$  exposes the tuple  $\langle \text{"Reg}[\text{RS2}]", \llbracket \text{Reg}[\text{RS2}] \rrbracket \rangle$  containing the value of the second source operand whenever a *div* instruction is executed.

**Phase 1: Empirical leakage characterization.** To empirically characterize a processor’s leakage, LEASYN first generates a set of test cases  $T$ , where each test case consists of a pairs of ISA-level programs with associated data inputs. Then, it simulates the execution of test cases on the target processor *IMPL* and determines which test cases are *attacker distinguishable* and which are not.

To allow the subsequent synthesis step to associate leakage with particular contract atoms in the template, LEASYN also simulates each test case on the contract template  $\mathbb{T}$ . The result of this simulation is condensed into the *template distinguishability* of each test case, which compactly captures which sets of contract atoms distinguish a particular test case and which do not.

In our example, we consider an attacker ATK that observes when instructions retire and we systematically generate the following test cases, each consisting of a pairs of program executions:

- (T<sub>1</sub>)  $\langle \text{li } R1, 0 \times 1234 \mid \sigma \rangle$  vs.  $\langle \text{li } R1, 0 \times 5678 \mid \sigma \rangle$  for some  $\sigma$ .
- (T<sub>2</sub>)  $\langle \text{li } R1, 0 \times 1234 \mid \sigma \rangle$  vs.  $\langle \text{li } R2, 0 \times 1234 \mid \sigma \rangle$  for some  $\sigma$ .
- (T<sub>3</sub>)  $\langle \text{li } R1, 0 \times 1234 \mid \sigma \rangle$  vs.  $\langle \text{li } R1, 0 \times 1234 \mid \sigma' \rangle$  for  $\sigma, \sigma'$ , s.t.  $\sigma(R1) \neq \sigma'(R1)$ , *i.e.*, for different initial values of  $R1$ .
- (T<sub>4</sub>)  $\langle \text{li } R1, 0 \times 1234 \mid \sigma \rangle$  vs.  $\langle \text{div } R1, R2, R3 \mid \sigma \rangle$  for some  $\sigma$ .
- (T<sub>5</sub>)  $\langle \text{div } R1, R2, R3 \mid \sigma \rangle$  vs.  $\langle \text{div } R1, R2, R3 \mid \sigma' \rangle$  for  $\sigma, \sigma'$ , s.t.  $\sigma(R1) \neq \sigma'(R1)$ ,  $\sigma(R2) \neq \sigma'(R2)$ , and  $\sigma(R3) \neq \sigma'(R3)$ . Further,  $\sigma(R3), \sigma'(R3) \notin \{0, 1\}$ , *i.e.*, neither of the initial states assigns 0 or 1 to  $R3$ .

For the attacker, all test cases except  $T_4$  are indistinguishable, since the program executions take the same amount of time. In terms of template distinguishability, for instance, the test case  $T_4$  is distinguished by including any of the template’s atoms in the contract.

**Phase 2a: Contract Synthesis.** Based on the empirical leakage characterization, LEASYN synthesizes, using integer linear programming, a candidate contract *CTR* that (a) distinguishes all attacker-distinguishable test cases, and (b) is as precise as possible given the template  $\mathbb{T}$ . In other words, it distinguishes as few of the attacker-indistinguishable test cases as possible.

In our example, based on the initial test cases, LEASYN might synthesize the following initial contract:

$$\text{CTR}_1 = \{(\text{li}, \langle \text{"li"}, - \rangle)\}$$

This contract  $\text{CTR}_1$  captures that the attacker  $\text{ATK}$  can distinguish  $\text{li}$  instructions from  $\text{div}$  instructions (as indicated by  $T_4$  above). Note that  $\text{CTR}_1$  is the most precise contract that captures all leaks exercised by the test cases. For instance, the contract exposing also the value of the immediate in  $\text{li}$  instructions, *i.e.*,  $(\text{li}, \langle \text{"imm"}, \llbracket \text{imm} \rrbracket \rangle)$ , would still capture the leak in  $T_4$ , but it also unnecessarily distinguishes  $T_5$  and it is, therefore, not a possible solution.

**Phase 2b: Bounded verification of soundness.** Since the test cases may miss some leaks, the synthesized contract is not guaranteed to be sound. To ensure soundness,  $\text{LEASYN}$  alternates the contract synthesis step with a bounded verification step, which either provides a counterexample to soundness or it proves that the contract is sound up to a given bound on the number of executed cycles. In case of unsoundness, the counterexample is used to refine the contract by looping back to the contract synthesis step (*Phase 2a*).

In our example,  $\text{CTR}_1$  is unsound since it does not capture the leaks generated by  $\text{div}$  instructions. The bounded verification might generate the following test case as counterexample to soundness:

$$(T_6) \langle \text{div R1, R2, R3} \mid \sigma \rangle \text{ vs. } \langle \text{div R1, R2, R3} \mid \sigma' \rangle \text{ for } \sigma, \sigma', \text{ s.t. } \\ \sigma(\text{R3}) = 0 \text{ and } \sigma'(\text{R3}) = 5.$$

Next,  $\text{LEASYN}$  simulates the contract template  $\mathbb{T}$  on the generated counterexamples to determine the distinguishing atoms, and re-executes the contract synthesis step (*Phase 2a*). This time  $\text{LEASYN}$  generates the following contract:

$$\text{CTR}_2 = \{(\text{div}, \langle \text{"Reg[RS2]"}, \llbracket \text{Reg[RS2]} \rrbracket \rangle)\}$$

$\text{CTR}_2$  captures that the value of the second source register influences the execution time of the  $\text{div}$  instruction, thereby distinguishing  $T_6$ . The presence of an atom that applies only to the  $\text{div}$  instruction also ensures that  $\text{div}$  instructions are distinguishable from  $\text{li}$  instructions, which distinguishes  $T_4$ .

**Phase 3: Unbounded verification of soundness.** As soon as the bounded verification does not find any more counterexamples,  $\text{LEASYN}$  attempts to verify the synthesized contract unboundedly. To this end,  $\text{LEASYN}$  generates inductive invariants using the Houdini algorithm, following the verification approach of Wang et al. [17]. If the verification succeeds,  $\text{LEASYN}$  returns the synthesized sound and precise contract to the user. In our example,  $\text{CTR}_2$  can be proved sound in an unbounded manner and  $\text{LEASYN}$  terminates.

## 2.4 Generating more precise contracts

$\text{LEASYN}$  synthesizes the most precise sound contract among all possible contracts given a specific contract template. The choice of the template, however, can significantly affect the precision of the synthesized contracts.

The sound contract  $\text{CTR}_2$  from § 2.3 is imprecise. For instance, it distinguishes the test case below, even though the CPU  $\text{IMPL}$  executes both instructions in the same number of cycles:

$$(T_7) \langle \text{div R1, R2, R3} \mid \sigma \rangle \text{ vs. } \langle \text{div R1, R2, R3} \mid \sigma' \rangle \text{ for } \sigma, \sigma', \text{ s.t. } \\ \sigma(\text{R3}) = 7 \text{ and } \sigma'(\text{R3}) = 8.$$

However,  $\text{LEASYN}$  cannot generate a more precise contract since the template from § 2.3 can only expose the entire value of operands for  $\text{div}$  instructions.

To derive more precise contracts, we need a more expressive contract template, which can be achieved by including additional

atoms of the following form, where  $k$  is a constant value:

$$(\text{div}, \langle \text{"Reg[RS2]=k?"}, \llbracket \text{Reg[RS2]} \rrbracket \rangle == k)$$

Running  $\text{LEASYN}$  again with the updated template and test cases  $T_1, \dots, T_7$  may result in generating the following contract at the end of *Phase 2a*:

$$\text{CTR}_2 = \{(\text{div}, \langle \text{"Reg[RS2]=0?"}, \llbracket \text{Reg[RS2]} \rrbracket \rangle == 0)\}$$

Even though this contract captures one of the leaks caused by  $\text{div}$ , it is not sound since it misses the other special case of the  $\text{div}$  instruction, *i.e.*, when  $\text{R2}$  is 1. When running the bounded verification (*Phase 2b*),  $\text{LEASYN}$  might produce the following counterexample:

$$(T_8) \langle \text{div R1, R2, R3} \mid \sigma \rangle \text{ vs. } \langle \text{div R1, R2, R3} \mid \sigma' \rangle \\ \text{for } \sigma, \sigma', \text{ s.t. } \sigma(\text{R2}) = 1 \text{ and } \sigma'(\text{R2}) = 5.$$

A new synthesis step, then, finally results in the following contract:

$$\text{CTR}_3 = \{(\text{div}, \langle \text{"Reg[RS2]=0?"}, \llbracket \text{Reg[RS2]} \rrbracket \rangle == 0), \\ (\text{div}, \langle \text{"Reg[RS2]=1?"}, \llbracket \text{Reg[RS2]} \rrbracket \rangle == 1)\}$$

This contract is sound and, again, it is the most precise given the updated template. Therefore,  $\text{LEASYN}$  terminates by returning  $\text{CTR}_3$  to the user after the final unbounded verification.

## 3 Formal model

Here, we present the core components of our formal model.

### 3.1 Architectures and microarchitectures

Leakage contracts act as a security abstraction between the instruction set architecture (short: architecture or ISA) and a microarchitecture, that is, a concrete implementation in a processor. Next, we first formalize both concepts and then what it means for a microarchitecture to correctly implement an architecture [17].

**Architectures.** We view an architecture as a state machine that defines how the execution progresses through a sequence of architectural states, where each transition corresponds to the execution of a single instruction. Formally, an *architecture* is a pair  $(\text{ARCHSTATE}, \text{ISA})$  where  $\text{ARCHSTATE}$  is a set of architectural states (each modeling the values of registers and data/instruction memory) and  $\text{ISA} : \text{ARCHSTATE} \rightarrow \text{ARCHSTATE}$  is a transition function that maps each state  $\sigma \in \text{ARCHSTATE}$  to its successor  $\text{ISA}(\sigma)$ , obtained by executing the next instruction to be executed in  $\sigma$ . To capture an execution, we denote by  $\text{ISA}^*(\sigma)$  the sequence of states reached from  $\sigma$  by successive applications of  $\text{ISA}$ , *i.e.*,  $\text{ISA}^*(\sigma) = \sigma_0, \sigma_1, \sigma_2, \dots$  where  $\sigma_0 = \sigma$  and  $\sigma_{i+1} = \text{ISA}(\sigma_i)$  for all  $i \geq 0$ .

**Microarchitectures.** We view microarchitectures as state machines modeling how the processor's state evolves at cycle-level. Thus, a *microarchitecture* is a triple  $(\text{IMPLSTATE}, \text{INITSTATE}, \text{IMPL})$  where  $\text{IMPLSTATE} = \text{ARCHSTATE} \times \mu\text{ARCHSTATE}$  is the set of microarchitectural states (some of which are initial states  $\text{INITSTATE} \subseteq \text{IMPLSTATE}$ ) and  $\text{IMPL} : \text{IMPLSTATE} \rightarrow \text{IMPLSTATE}$  is a transition function that maps each state  $\sigma \in \text{IMPLSTATE}$  to its successor  $\text{IMPL}(\sigma)$ , obtained by executing the processor for one cycle. Each microarchitectural state  $\sigma \in \text{IMPLSTATE}$  consists of an architectural part  $\sigma_{\text{ISA}} \in \text{ARCHSTATE}$  and a microarchitectural part  $\sigma_{\text{impl}} \in \mu\text{ARCHSTATE}$  modeling the state of microarchitectural components (such as caches and predictions). For simplicity, we require that there is a canonical initial microarchitectural component  $\mu_0 \in$

$\mu\text{ARCHSTATE}$  such that  $\text{INITSTATE} = \{(\sigma, \mu_0) \mid \sigma \in \text{ARCHSTATE}\}$ . Similarly to  $\text{ISA}^*(\sigma)$ ,  $\text{IMPL}^*(\sigma)$  denotes the sequence of states reached from  $\sigma$  by successive applications of  $\text{IMPL}$ . Moreover, given a predicate  $\phi$  over  $\text{IMPLSTATE}$ ,  $\text{IMPL}^*\phi(\sigma)$  denotes the sequence of elements from  $\text{IMPL}^*(\sigma)$  that satisfy  $\phi$ .

**ISA compliance.** To correctly implement an architecture  $\text{ISA}$ , an implementation  $\text{IMPL}$  needs to change the architectural state in a manner consistent with  $\text{ISA}$ . We capture this following the ISA compliance notion from Wang et al. [17], which relies on a *retirement predicate*  $\phi$  that indicates when  $\text{IMPL}$  retires instructions. Then, we say that a microarchitecture  $\text{IMPL}$  implements an architecture  $\text{ISA}$  if one can map changes of the architectural state in  $\text{IMPL}$  to  $\text{ISA}$ 's executions using  $\phi$ .

Formally, a microarchitecture  $\text{IMPL}$  *correctly implements* an architecture  $\text{ISA}$  [17] given a retirement predicate  $\phi$  over  $\text{ARCHSTATE}$ , written  $\text{IMPL} \vdash_{\phi} \text{ISA}$ , if for all valuations  $\sigma \in \text{INITSTATE}$ :

- (a)  $(\text{IMPL}^*\phi(\sigma))_{\text{ISA}} = \text{ISA}^*(\sigma_{\text{ISA}})$ , *i.e.*, all architectural changes witnessed by  $\phi$  agree with  $\text{ISA}$ , and
- (b)  $(\text{IMPL}^i(\sigma))_{\text{ISA}} = (\text{IMPL}^{i-1}(\sigma))_{\text{ISA}}$  whenever  $(\text{IMPL}^i(\sigma))_{\text{ISA}} \not\models \phi$ , *i.e.*, no architectural changes beyond those witnessed by  $\phi$ .

### 3.2 Leakage contracts

Next, we introduce how we model leakage contracts in  $\text{LEASYN}$ . We first introduce *contract atoms*, which are the basic building block for contracts. Next, we introduce *contract templates*, which capture the synthesis space for our approach as a set of possible atoms.

**Contract atoms.** Atoms capture potential leaks at the instruction level, and they are built from *applicability predicates* and *leakage functions*. The former determine whether a contract atom is applicable in a given architectural state, whereas the latter determine what information from the architectural state is leaked in case an atom is applicable. Formally, a *contract atom*  $A$  is a pair  $A = (\pi_A, \phi_A)$  where the applicability predicate  $\pi_A$  is a predicate over  $\text{ARCHSTATE}$  and the leakage function  $\phi_A$  is a function  $\phi_A : \text{ARCHSTATE} \rightarrow \text{CTROBS}$  mapping an architectural state to a contract observation in  $\text{CTROBS}$ .

**Example 1.** Addition and subtraction instructions `addi` and `subi` could both leak their immediate value. Contracts atoms for these instructions can be defined as follows:

- $(\text{addi}, \langle \text{"imm"}, \llbracket \text{imm} \rrbracket \rangle)$ , where predicate `addi` holds whenever the instruction is an addition instruction and  $\langle \text{"imm"}, \llbracket \text{imm} \rrbracket \rangle$  is a leakage function that leaks the immediate value of the instruction  $\llbracket \text{imm} \rrbracket$  together with a leakage identifier `"imm"`.
- $(\text{subi}, \langle \text{"imm"}, \llbracket \text{imm} \rrbracket \rangle)$ , where predicate `subi` holds whenever the instruction is a subtraction instruction and the leakage function is the same as above.

Importantly, note that the test case  $\langle \text{addi } R1, R2, 0x1234 \mid \sigma \rangle$  vs.  $\langle \text{subi } R1, R2, 0x1234 \mid \sigma \rangle$  would be indistinguishable by a contract including both atoms. Even though two different instructions are executed, both atoms would be applicable and expose  $\langle \text{"imm"}, 0x1234 \rangle$ .

**Contract templates.** A *contract template*  $\mathbb{T}$  is a set of contract atoms. We require that the applicability predicates of contract atoms with the same leakage function are mutually exclusive, *i.e.*, at most one of these atoms is applicable in any given state. This holds in Example 1, since instructions can either be additions or subtractions.

We also require that the images of different leakage functions are disjoint. This ensures that leakage from one function cannot "cancel out" leakage from another function. For instance, a leakage function exposing an immediate operand and another function exposing the instruction's opcode must not leak the same value. A simple way to satisfy this requirement is associating a unique identifier with each leakage function, like `"imm"` in Example 1.

**Leakage contracts.** Given a template  $\mathbb{T}$ , any subset  $S \subseteq \mathbb{T}$  of the template induces a *contract*  $LC_S : \text{ARCHSTATE} \rightarrow 2^{\text{CTROBS}}$ , which is a function from architectural states to sets of observations.<sup>1</sup> Formally,  $LC_S$  is defined by evaluating each applicable atom in  $S$  as follows:  $LC_S(\sigma) := \{\phi_A(\sigma) \mid (\pi_A, \phi_A) \in S \wedge \pi_A(\sigma)\}$ . Given a sequence  $\tau := \sigma_0, \sigma_1, \dots$  of architectural states,  $LC_S(\tau)$  denotes the corresponding sequence of observations  $LC_S(\sigma_0), LC_S(\sigma_1), \dots$ . Therefore, given an architectural state  $\sigma \in \text{ARCHSTATE}$ ,  $LC_S(\text{ISA}^*(\sigma))$  denotes the *contract trace*, *i.e.*, the sequence of contract observations associated with the execution  $\text{ISA}^*(\sigma)$ .

Finally, in  $\text{LEASYN}$ , a *test case*  $T = (\sigma, \sigma')$  is a pair of architectural states. We say that  $T$  is *contract distinguishable* for given contract  $LC_S$  if the corresponding contract traces are different, *i.e.*,  $LC_S(\text{ISA}^*(\sigma)) \neq LC_S(\text{ISA}^*(\sigma'))$ .

### 3.3 Contract satisfaction

We conclude by formalizing contract satisfaction [16] in our setting. For this, we model microarchitectural attackers and then characterize when a contract captures all leaks observable by the attacker.

**Attackers.** We consider (passive) microarchitectural attackers that can extract information from the microarchitectural state. Formally, we model a microarchitectural attacker as a function  $\text{ATK} : \text{IMPLSTATE} \rightarrow \text{ATKOBS}$  that maps microarchitectural states to attacker observations in  $\text{ATKOBS}$ . Common attacker models, like the one exposing the timing of instruction retirement [27] or the one exposing the final state of caches [3, 28], can be instantiated in this setting. Given an initial microarchitectural state  $\sigma \in \text{INITSTATE}$ ,  $\text{ATK}(\text{IMPL}^*(\sigma))$  denotes the *attacker trace*, *i.e.*, the sequence of attacker observations associated with the cycle-accurate microarchitectural execution  $\text{IMPL}^*(\sigma)$ . We say that a test case  $T = (\sigma, \sigma')$  is *attacker distinguishable* if the corresponding attacker traces are different, *i.e.*,  $\text{ATK}(\text{IMPL}^*(\sigma)) \neq \text{ATK}(\text{IMPL}^*(\sigma'))$ .

**Contract Satisfaction.** A microarchitecture  $\text{IMPL}$  satisfies the contract  $LC_S$  for an attacker  $\text{ATK}$  if  $\text{ATK}$  cannot learn more information about the initial architectural state by monitoring  $\text{IMPL}$ 's executions than what is exposed by the contract. That is, for any two initial states, whenever the contract traces are the same, then the attacker traces must be identical, *i.e.*,  $\text{ATK}$  cannot distinguish the two architectural executions. This is formalized in Definition 1.

**Definition 1.** Microarchitecture  $\text{IMPL}$  satisfies contract  $LC_S$  for attacker  $\text{ATK}$ , written  $\text{ISA}, \text{IMPL} \vdash \text{CTRSAT}(LC_S, \text{ATK})$ , if for all initial states  $\sigma, \sigma' \in \text{INITSTATE}$ , if  $LC_S(\text{ISA}^*(\sigma_{\text{ISA}})) = LC_S(\text{ISA}^*(\sigma'_{\text{ISA}}))$ , then  $\text{ATK}(\text{IMPL}^*(\sigma)) = \text{ATK}(\text{IMPL}^*(\sigma'))$ .

Definition 1 refers to 4 different traces: two contract traces defined over the architecture  $\text{ISA}$  and two attacker traces defined over the microarchitecture  $\text{IMPL}$ . Wang et al. [17] proposed the notion of microarchitectural contract satisfaction, formalized below, which is

<sup>1</sup>We often refer to a contract  $LC_S$  directly using the set of atoms inducing it.

expressed only over a pair of microarchitectural traces and allows to decouple reasoning about leakage and about ISA compliance.

**Definition 2.** Microarchitecture  $IMPL$  microarchitecturally-satisfies contract  $LC_S$  for attacker  $ATK$  and predicate  $\phi$ , written  $IMPL \vdash_\phi \text{CTRSAT}(LC_S, ATK)$ , if for all initial states  $\sigma, \sigma'$ , if  $LC_S(IMPL^*|\phi(\sigma)) = LC_S(IMPL^*|\phi(\sigma'))$ , then  $ATK(IMPL^*(\sigma)) = ATK(IMPL^*(\sigma'))$ .

As stated in [17, Theorem 1], whenever  $IMPL$  correctly implements  $ISA$ , then contract satisfaction and microarchitectural contract satisfaction are equivalent. Hence, Definition 2 is the notion targeted by LEASYN and, for conciseness, we will refer to it simply as “contract satisfaction” throughout the rest of the paper.

## 4 Contract synthesis

In this section, we describe LEASYN’s synthesis approach. First, we describe the empirical leakage characterization (*Phase 1* in Fig. 1; § 4.1). Next, we describe how LEASYN synthesizes a candidate contract using integer linear programming (*Phase 2a* in Fig. 1; § 4.2). Then, we describe how LEASYN checks whether a contract is sound, *i.e.*, it capture all leaks, with bounded verification (*Phase 2b* in Fig. 1; § 4.3) and unbounded verification (*Phase 3* in Fig. 1; § 4.4).

### 4.1 Phase 1: Empirical leakage characterization

LEASYN generates a set of test cases  $T$  to characterize a processor’s leakage. As an input to the ILP-based contract synthesis (*Phase 2a*), LEASYN needs to determine for each test case: (a) whether it is *attacker distinguishable*, and (b) which contracts from the contract template would make it *contract distinguishable*.

**Test-case generation:** Test cases aid contract synthesis in two opposing ways: (1) Attacker-distinguishable test cases uncover leaks and force the inclusion in the contract of atoms that expose these leaks. (2) Attacker-indistinguishable test cases show which atoms should not be included in the contract to avoid imprecision.

Based on this observation, our goal is to generate test cases that differ in exactly one atom. This has two benefits. First, if a test case is attacker distinguishable, there is only one possible responsible atom. Second, as test cases “differ” only in one atom, this strategy is likely to generate many attacker indistinguishable tests, in particular when targeting atoms that do not capture leakage in the processor.

To generate such test cases, we proceed as follows. We start by generating an architectural state  $\sigma_0$  that consists of a random instruction sequence and a random initial valuation of the registers. Then, we randomly pick an instruction from the sequence and an applicable leakage function and generate a pair of architectural states  $(\sigma, \sigma')$  such that the leakage function will evaluate differently in  $\sigma$  and  $\sigma'$  on the chosen instruction.

To implement this, we define a *modifier function* for each leakage function in our template. Modifier functions may, for instance, change the initial value of a register, inject new instruction to adapt the architectural state, or modify one of the instructions.

Recall the template introduced in § 2.3. Given this template, the modifier function for the leakage function `imm` would for example transform a state  $\sigma_0$  whose first instruction is `li R1, 0x1234` to the pair  $(\sigma, \sigma')$  where  $\sigma = \sigma_0$  and the only difference in  $\sigma'$  is that its first instruction is `li R1, 0x5678`. Similarly, the modifier function for the leakage function `RD` could substitute the instruction `div`

`R1, R2, R3` in  $\sigma_0$  with `div R4, R2, R3` in  $\sigma'$ . For other leakage functions such as `Reg[RS1]`, the modifier function does not alter any instructions but modifies the value of a specific register in the initial architectural state of  $\sigma'$ . Depending on the context of an instruction, modifier functions are not guaranteed to generate a test case that only differs in the chosen leakage function, *e.g.*, if dependent instructions are present in the instruction sequence.

**Attacker distinguishability:** Each of the test cases  $t = (\sigma, \sigma')$  in  $T$  is executed on the target CPU  $IMPL$ . To derive the attacker traces for  $\sigma$  and  $\sigma'$ , we construct the two initial states  $(\sigma, \mu_0)$  and  $(\sigma', \mu_0)$  for the CPU  $IMPL$ . To determine attacker indistinguishability, we then simulate the executions  $IMPL^*((\sigma, \mu_0))$  and  $IMPL^*((\sigma', \mu_0))$  and check whether the attacker can distinguish them, *i.e.*, we check if  $ATK(IMPL^*((\sigma, \mu_0))) \neq ATK(IMPL^*((\sigma', \mu_0)))$ . In general, the executions  $IMPL^*((\sigma, \mu_0))$  and  $IMPL^*((\sigma', \mu_0))$  might be infinite, so LEASYN simulates them up to a fixed number of cycles  $n$ , and we pick an  $n$  that is large enough to fully simulate all our test cases.

**Template distinguishability:** To succinctly characterize the contract distinguishability of each test case for all possible contract induced by the template  $\mathbb{T}$ , we rely on the following observation.

Let  $(\sigma, \sigma') \in T$  be a test case and let  $LC_S$  be a contract. Also, let  $ISA^*(\sigma) = \sigma_0, \sigma_1, \dots$  and  $ISA^*(\sigma') = \sigma'_0, \sigma'_1, \dots$  be the architectural executions from  $\sigma$  and  $\sigma'$ . Test case  $(\sigma, \sigma')$  can be contract distinguishable under  $LC_S$  for two distinct reasons:

- **Leakage Mismatch:** There are atoms  $A, B \in S$  with  $\phi_A = \phi_B$  and an index  $i$ , s.t.  $\pi_A(\sigma_i) \wedge \pi_B(\sigma'_i)$  and  $\phi_A(\sigma_i) \neq \phi_B(\sigma'_i)$ , *i.e.*, at some point, the atoms  $A$  and  $B$  are applicable in the two executions but evaluate differently. Note that  $A$  and  $B$  can be the same atom.
- **Applicability Mismatch:** There is an atom  $A \in S$  and an index  $i$ , s.t.  $\pi_A(\sigma_i)$  (or  $\pi_A(\sigma'_i)$ ) and no atom  $B \in S$  with  $\phi_A = \phi_B$  is applicable in  $\sigma'_i$  (or  $\sigma_i$ ), *i.e.*, there is an observation only in one executions.

**Example 2.** Consider the following test cases:

- Test case `li R1, 0x1234` vs. `li R1, 0x5678` would be distinguishable due to a leakage mismatch by a contract including the atom `(li, (<"imm", [[imm]]))`, since the immediate value exposed by the atom is different in the two executions.
- Test case `li R1, 0x1234` vs. `add R1, R2, R3` would be distinguishable due to an applicability mismatch by a contract consisting only of the atom `(li, ("imm", [[imm]]))` as no atom with the same leakage function is applicable in the second execution.

We characterize the template distinguishability of each test case  $t = (\sigma, \sigma')$  via its *strongly-distinguishing atoms*  $SD_t \subseteq \mathbb{T}$  and its *xor-distinguishing pairs*  $XOR_t \subseteq \binom{\mathbb{T}}{2}$ :

A *strongly-distinguishing atom*  $A \in \mathbb{T}$  is one for which there exists an index  $i$  such that  $\pi_A(\sigma_i)$  (or  $\pi_A(\sigma'_i)$ ) holds, and no atom  $B \in \mathbb{T}$  satisfies  $\pi_B(\sigma'_i) \wedge \phi_A(\sigma_i) = \phi_B(\sigma'_i)$  (or  $\pi_B(\sigma_i) \wedge \phi_B(\sigma_i) = \phi_A(\sigma'_i)$ ). Including any strongly-distinguishing atom  $A \in SD_t$  in a contract will make the test case  $t$  contract distinguishable independently of which other atoms are included. This contract distinguishability may be caused by a leakage mismatch *or* an applicability mismatch.

For example, in the test case `li R1, 0x1234` vs. `li R1, 0x5678` the atom `(li, (<"imm", [[imm]]))` is strongly distinguishing, as it is applicable in both executions and the immediate values are different.

A *xor-distinguishing atom pair*  $\{A, B\} \in \binom{\mathbb{T}}{2}$  with  $A \neq B$  is one for which there exists an index  $i$  such that  $\pi_A(\sigma_i) \wedge \pi_B(\sigma'_i)$  holds

and  $\phi_A(\sigma_i) = \phi_B(\sigma'_i)$ . Due to our assumption that the applicability predicates of contract atoms that share a leakage function are mutually exclusive, there can be no other atom  $C \in \mathbb{T}$  with  $\phi_A = \phi_B = \phi_C$  that is applicable in  $\sigma_i$  or  $\sigma'_i$ . Thus, including exactly one of the two atoms  $A$  and  $B$  in a xor-distinguishing pair  $\{A, B\} \in XOR_t$  will make the test case  $t$  contract distinguishable due to an applicability mismatch. For example, in the test case `addi R1, R2, 0x1234 vs. subi R1, R2, 0x1234` the pair of atoms `(addi, ("imm", [imm]))` and `(subi, ("imm", [imm]))` is a xor-distinguishing pair.

The two sets  $SD_t$  and  $XOR_t$  fully characterize the contract distinguishability of a test case  $t$  under any possible contract:

**Proposition 1.** *Let  $t = (\sigma, \sigma')$  be a test case and let  $LC_S$  be a contract. Then,  $t$  is contract distinguishable under  $LC_S$  iff*

- $S$  includes a strongly-distinguishing atom, i.e.,  $SD_t \cap S \neq \emptyset$ , or
- $S$  includes exactly one of the two atoms in a xor-distinguishing pair, i.e.,  $\exists \{A, B\} \in XOR_t$  with  $|S \cap \{A, B\}| = 1$ .

To compute the sets  $SD_t$  and  $XOR_t$  for a test case  $t$ , we simulate the execution of the test case w.r.t. the contract template  $\mathbb{T}$  and check the conditions for strongly-distinguishing and xor-distinguishing atoms for each index  $i$ . To simulate  $ISA^*(\sigma)$  (which might be infinite), LEASYN only simulates executions up to a fixed number of steps.

Prior work by Mohr et al. [22] relies solely on *distinguishing atoms*. An atom  $A \in \mathbb{T}$  is distinguishing if there exists an index  $i$  such that  $\pi_A(\sigma_i) \oplus \pi_A(\sigma'_i)$  or  $\pi_A(\sigma_i) \wedge \pi_A(\sigma'_i) \wedge \phi_A(\sigma_i) \neq \phi_A(\sigma'_i)$ . For templates that contain multiple atoms that share the same leakage function, the set of distinguishing atoms is not sufficient to fully characterize a test case's contract distinguishability. In particular, the inclusion of two distinguishing atoms may result in contract indistinguishability as Example 1 demonstrates.

## 4.2 Phase 2a: ILP-based contract synthesis

Given a template  $\mathbb{T}$ , the empirical leakage characterization from § 4.1 computes the following information for each test case  $t \in T$ :

- $SD_t \subseteq \mathbb{T}$  the set of strongly-distinguishing atoms,
- $XOR_t \subseteq \binom{\mathbb{T}}{2}$  the set of xor-distinguishing atom pairs, and
- $d_t \in \mathbb{B}$  whether the test case is attacker distinguishable.

For convenience, we denote the set of attacker-distinguishable test cases as  $T_d = \{t \in T \mid d_t = 1\}$  and the set of attacker-indistinguishable test cases as  $T_{nd} = T \setminus T_d$ .

Based on this information, LEASYN uses integer linear programming (ILP) to synthesize a contract. That is, we compute a set of atoms  $S \subseteq \mathbb{T}$ , such that  $LC_S$  distinguishes all attacker-distinguishable test cases and as few attacker-indistinguishable test cases as possible. Next, we detail the ILP formulation that achieves this.

**Variables.** For each atom  $A \in \mathbb{T}$ , we introduce a boolean variable  $s_A$  that is true iff the atom  $A$  is included in the synthesized contract.

We also introduce a boolean variable  $fp_t$  for each  $t \in T_{nd}$  to capture whether the attacker-indistinguishable test case  $t$  is a false positive, i.e., it is contract distinguishable in the synthesized contract.

Finally, we introduce a boolean variable  $x_{\{A,B\}}$  for each pair  $\{A, B\}$  that is xor-distinguishing for some test case  $t$ .

**Objective function.** The objective function of the ILP then is to minimize the number of false positives:  $\min \sum_{t \in T_{nd}} fp_t$ . As a secondary objective we minimize the number of atoms in the contract,

i.e., we minimize  $\sum_{A \in \mathbb{T}} s_A$ , which promises to yield a smaller contract that is likely more precise on unseen test cases.

**Constraints.** It remains to ensure that all attacker-distinguishable test cases are contract distinguishable and that the  $fp_t$  variables are consistent with the contract encoded by the  $s_A$  variables. From Proposition 1, we know that a test case  $t$  is contract distinguishable iff the contract contains (a) a strongly-distinguishing atom or (b) exactly one of the two atoms of one of its xor-distinguishing pairs.

We start by adding the following constraints enforcing  $x_{\{A,B\}} = s_A \oplus s_B$  for any xor-distinguishing pair  $\{A, B\} \in \bigcup_{t \in T_d} XOR_t$ :

$$\begin{aligned} x_{\{A,B\}} &\leq s_A + s_B, & x_{\{A,B\}} &\geq s_A - s_B, \\ x_{\{A,B\}} &\geq s_B - s_A, & x_{\{A,B\}} &\leq 2 - s_A - s_B. \end{aligned}$$

Following Proposition 1, we add the following constraint for every  $t \in T_d$  to ensure that every attacker-distinguishable test case is contract distinguishable:

$$\sum_{A \in SD_t} s_A + \sum_{\{A,B\} \in XOR_t} x_{\{A,B\}} \geq 1.$$

Finally, to ensure that the  $fp_t$  variables are consistent with the contract distinguishability of the test cases we add the following constraints for each attacker-indistinguishable test case  $t \in T_{nd}$ :

$$\begin{aligned} fp_t &\geq s_A && \text{for every } A \in SD_t, \\ fp_t &\geq x_{\{A,B\}} && \text{for every } \{A, B\} \in XOR_t. \end{aligned}$$

**Proposition 2.** *Any solution to the ILP corresponds to a set of atoms  $S \subseteq \mathbb{T}$  such that  $LC_S$  distinguishes all attacker-distinguishable test cases and as few attacker-indistinguishable test cases as possible.*

## 4.3 Phase 2b: Bounded soundness verification

The ILP-based synthesis from § 4.2 generates a new contract based on the attacker and template distinguishability derived from the test cases (§ 4.1). Since the test cases may miss some leaks in the target processor, the contract generated from *Phase 2a* may be unsound. To discover leaks that are missed by the candidate contract, LEASYN performs a bounded verification step, which we describe next.

**Intuition.** LEASYN verifies whether the candidate contract  $LC_S$  captures all leaks visible by the attacker ATK on target CPU IMPL when considering all possible executions up to a given length  $k$ . That is, LEASYN verifies a bounded variant of Definition 2. For this, LEASYN encodes the contract verification task as a bounded model checking (BMC) problem. Since microarchitectural contract satisfaction is a property defined over pairs of IMPL executions, our encoding relies on self-composition [29], which reduces reasoning about pairs of executions to reasoning about a single execution, by first constructing a product circuit consisting of two copies of IMPL. The bounded contract satisfaction property  $\phi_{ctrsat}$  is then encoded on top of this circuit. The BMC either (a) falsifies  $\phi_{ctrsat}$  and provides a counterexample—a test case with the same contract traces but different attacker traces—which can be used to synthesize a better contract, or (b) proves that the contract is sound up to bound  $k$ .

**Constructing the product circuit.** We construct a product circuit consisting of two copies of IMPL executing in parallel. To compare contract observations, which are produced only when instructions retire (i.e., whenever the user-provided retirement predicate  $\psi$  holds), we need to synchronize the two copies of IMPL on retirement.

That is, whenever one copy of *IMPL* retires an instruction but the other does not (*i.e.*,  $\psi$  holds only on one of the two executions), the copy of *IMPL* where  $\psi$  holds is “paused” until the other one catches up. To do so, LEASYN uses the stuttering product circuit construction introduced by Wang et al. [17], which constructs a product circuit synchronized on the retirement predicate  $\psi$ .

**Property.** We encode microarchitectural contract satisfaction on top of the product circuit into the  $\phi_{ctrsat}^{C,\psi}$  property, where  $C$  is the contract and  $\psi$  is the retirement predicate, as follows. Our encoding is parametric in: (a) the total bound  $k$ , and (b) the attacker bound  $1 \leq b \leq k$ . Intuitively, the formula  $\phi_{ctrsat}^{C,\psi}$  works as follows:

- (1) it assumes that, at cycle 0, the two copies of *IMPL* start from a valid initial state and with the same microarchitectural state,
- (2) it assumes that, for cycles 0 to  $k$ , the two executions produce the same contract traces, *i.e.*, whenever the retirement predicate  $\psi$  holds on both executions, the contract observations are the same,
- (3) it asserts that, for cycles 0 to  $b$ , the two executions have the same attacker observations, *i.e.*, they are attacker-indistinguishable.

Note that the attacker bound  $b$  is less than the total bound  $k$ , since contract observations are “slower” than attacker observations, as the former are produced only when instructions retire. This requires looking ahead to check whether future contract observations may “declassify” a difference in attacker observations. Since these differences are often caused by in-flight instructions,  $k - b$  needs to be large enough to account the retirement of these instructions to retire to ensure that the associated observations are produced. However, sizing  $k - b$  to the worst-case, result in large bounds that (a) slow down checking contract satisfaction in simpler cases, and (b) may result in counterexamples with many instructions.

To account for this, LEASYN uses a simple optimization, which is parametric in an instruction bound  $i$ . As soon as a difference in attacker observations is detected (in the first  $b$  cycles), we start decrementing  $i$  whenever an instruction is retired by both executions. LEASYN then terminates the BMC as soon as either  $i$  reaches 0 or the cycle bound  $k$  is reached, whichever comes first. This optimization allows us to dynamically adjust the number of cycles explored by the BMC. As we show in § 6.2, this has two benefits: (a) it speeds up bounded verification since the BMC finds many simple counterexamples faster, without having to always explore all  $k$  cycles, and (b) it often results in shorter counterexamples with fewer instructions, which restrict the synthesizer’s search space. The encoding of the property checked by LEASYN in terms of `assume` and `assert` statements in Verilog is given in § B.

**Proposition 3.** *Whenever the BMC disproves  $\phi_{ctrsat}^{C,\psi}$ , with retirement predicate  $\psi$ , BMC bound  $k$ , attacker bound  $1 \leq b \leq k$ , and instruction bound  $i$ , then there is an attacker-distinguishable test case  $T$ , of at most  $i$  instructions, that has the same  $C$ -traces.<sup>2</sup>*

#### 4.4 Phase 3: Unbounded soundness verification

As soon as Phase 2b cannot discover further counterexamples, LEASYN performs an unbounded verification step to ensure that the synthesized contract is *sound*, *i.e.*, all possible leaks in the target processor are captured by the contract.

For this, LEASYN uses the LEAVE contract verifier [17] to prove that (unbounded) microarchitectural contract satisfaction holds for the current contract and the target CPU. In a nutshell, LEAVE verifies unbounded contract satisfaction by (a) learning the strongest possible inductive invariant over pairs of contract-equivalent executions, and (b) use this invariant to prove that contract-equivalent executions are also attacker-equivalent.

If LEAVE’s verification passes, the synthesized contract is sound (from [17, Theorem 2]) and as precise as possible (from Proposition 2) for the target processor. If it fails, LEASYN terminates and notifies the user. This failure can be due to two reasons:

- If the contract is unsound, the user can re-run LEASYN with a larger bound for Phase 2b to identify the missed leaks.
- If contract is sound but LEAVE fails to verify it, the user can provide additional information to LEAVE (e.g., additional relational invariants) to make the unbounded proof go through [17].

**Proposition 4.** *If Phase 3 passes, the contract synthesized by LEASYN is sound and distinguishes as few attacker-distinguishable test cases as possible among the test cases explored during the synthesis process.*

## 5 Implementation

In this section, we present LEASYN, a prototype that implements our synthesis approach from § 4 for CPU designs in Verilog. Our prototype relies on the Yosys Open Synthesis Suite [30] for processing Verilog circuits, the Icarus Verilog simulator [31] for simulating test cases, the Google OR-Tools [32] for the ILP, the Yices SMT solver [33] for the bounded verification, and the LEAVE contract verification tool [17] for the unbounded verification. LEASYN’s implementation is available at [34]. Furthermore, we will release LEASYN as open source together with benchmarks and scripts for reproducing the experiments from § 6.

**Inputs.** LEASYN takes as input (1) the target processor *IMPL* implemented in Verilog, (2) a Verilog expression  $e$  over *IMPL* expressing the retirement predicate, (3) a Verilog expression *ATK* over *IMPL* modeling the microarchitectural attacker, (4) the bounds for the bounded verification, and (5) the additional inputs for the unbounded verification in LEAVE (e.g., relational invariants). The user can also provide as input a set of invariants for fixing the initial state of *IMPL*, e.g., for the initial value of microarchitectural components.

**Contract templates.** LEASYN implements multiple contract templates (described in § 6.1) that capture different classes of instruction-level leaks. Even though these templates are defined in terms of the RISC-V ISA, they need to be instantiated for each new target processor since different processors often implement the same architectural elements (e.g., the register file) with different signals in Verilog. To address this, all our templates and atoms (across the entire LEASYN’s pipeline) are implemented in terms of the RISC-V formal interface [25] (RVFI), which provides a common interface to the architectural state of a RISC-V core for testing and verification. This interface, which is implemented by several open-source cores, significantly simplifies applying LEASYN to a new CPU that already implements the RVFI.

**Test generation.** LEASYN can accept user-provided test cases as input. Additionally, LEASYN implements a test generation strategy for automatically synthesizing a given number of test cases. For

<sup>2</sup>We require  $k - b$  to be large enough for *IMPL* to execute and retire at least  $i$  instructions.

each leakage function in the implemented templates, we also implement the associated modifier functions, which LEASYN uses to generate test cases. LEASYN’s test generation strategy follows the process described in § 4.1: (a) we first randomly generate an architectural state  $\sigma$  (consisting of a RISC-V program and initial values for registers), (b) identify one of the instructions in the program, and (c) use one of the applicable modifier functions to the selected instruction to generate the other initial state  $\sigma'$ .

**Workflow.** LEASYN implements the workflow described in § 4. First, it generates a set of test cases, simulates them, and compute their attacker and template-distinguishability w.r.t. the attacker and selected template. Then, it alternates between synthesis and bounded verification to come up with precise candidate contracts. Finally, it verifies whether the candidate contract is sound with the LEAVE (unbounded) contract verifier and terminates.

## 6 Evaluation

In this section, we report on the use of LEASYN to synthesize sound and precise leakage contracts. We first introduce the methodology we followed for our evaluation (§ 6.1): the processors we analyze, the contract templates and attacker we consider, and the experimental setup. In our evaluation (§ 6.2), we address these research questions:

- **RQ1:** Can LEASYN synthesize sound and precise leakage contracts from processor designs at RTL?
- **RQ2:** What is the impact of the number of initial test cases on the synthesized contract’s precision and on the synthesis time?
- **RQ3:** What is the impact of different contract templates on the synthesized contract’s precision?
- **RQ4:** Does LEASYN synthesize contracts that are more precise than sound leakage contracts derived in prior work?
- **RQ5:** What is the impact of different property encodings on the verification time of Phase 2b (§ 4.3)?

### 6.1 Methodology

**Benchmarks.** We analyze the following CPUs:

- **DarkRISCV:** An in-order RISC-V core that implements the RV32E and RV32I instruction sets [35]. We analyzed the 2-stage **DarkRISCV-2** and the 3-stage **DarkRISCV-3** implementations.
- **Sodor-2:** An educational 2-stage RISC-V core [36] that implements the RV32I instruction set.
- **Ibex:** An open-source, production-quality 32-bit RISC-V CPU core [37].<sup>3</sup> We target the default “small” configuration, which has two stages and supports the RV32IMC instruction set. We study three variants: (1) **Ibex-small** is the default “small” configuration with constant-time multiplication (three cycles) and without caches, (2) **Ibex-cache** is **Ibex-small** extended with a simple (single-line) cache, and (3) **Ibex-mult-div** employs a non-constant-time multiplication unit whose execution time depends on the operands [38].

**Contract templates.** We consider the following basic templates:

- **Instruction Leaks (I):** atoms exposing information about an instruction’s encoding (op code, destination and source registers, immediate values).
- **Register Leaks (R):** atoms exposing the value of an instruction’s source and destination registers, and the program counter.

- **Memory Leaks (M):** atoms exposing the accessed memory address and content read from/written to memory.

- **Alignment Leaks (A):** two atoms that expose the alignment of a memory access: *IS\_ALIGNED* exposes whether the last two bits of the memory address are 00 and *IS\_HALF\_ALIGNED* exposes whether the last two bits of the memory address are not 11.

- **Branch Leaks (BT):** atoms that expose the outcome of branch instructions. For conditional jumps, the atoms expose if the branch is taken. For unconditional ones, the atoms return the constant 1.

- **Value Leaks (V):** atoms exposing selected information about values of source and destination registers: (a) whether these values are equal to 0, and (b) the logarithm of the value. These atoms are inspired by known value-dependent leaks in the **Ibex** core [17, 22].

Additionally, we consider templates obtained by combining the basic ones listed above, and we denote such combinations with +. In particular, the **Base Leaks (B)** template consists of **I + R + M**, that is, the template containing all atoms from **I**, **R**, and **M**.

**Attacker.** We consider an attacker that observes when instructions retire, i.e., when the retirement predicate holds.

**Unbounded verification setup.** For the unbounded verification using LEAVE, for all benchmarks, we re-use the state invariants (and additional candidate inductive invariants) manually constructed by Wang et al. [17] for verifying the same cores.

**Experimental setup.** All our experiments ran on an AMD Ryzen Threadripper PRO 5995WX with 64 cores and 512 GB of RAM. Each experiment ran in a dockerized Ubuntu 24.04 with OpenJDK 21. LEASYN used Google OR-Tools version 9.10, Yosys version 0.48, Icarus Verilog version s20250103, and Yices version 2.6.5.

**Evaluating precision.** Leakage contracts classify pairs of executions as distinguishable or indistinguishable. To measure the precision of synthesized contracts, we use the standard notion of precision used for binary classifiers, following [22]. Given a target processor, a contract  $C$ , and a validation set  $V$  of test cases, the precision of  $C$  w.r.t.  $V$  is  $\frac{FP}{TP+FP}$  where  $TP$  is the number of test cases that are contract and attacker distinguishable, and  $FP$  is the number of test cases that are contract distinguishable but attacker indistinguishable. For each benchmark, we generate a validation set consisting of 2048K test cases.

### 6.2 Experimental results

**RQ1 – Synthesis for open-source RISC-V cores.** To evaluate whether LEASYN can synthesize sound and precise leakage contracts from open-source cores, we apply LEASYN to all benchmarks from § 6.1. For each core, we use LEASYN to synthesize a contract against the **B + A + BT + V** template with a set of 128K test cases.

Table 1 summarizes the results of our evaluation. The table reports (a) the number of iterations for the synthesis loop, (b) the static BMC bound  $k$ , the attacker bound  $b$  and the instruction bound  $i$  used, (c) the total synthesis time split between Phase 1, Phase 2a, Phase 2b, and Phase 3, (d) the number of atoms in the final synthesized contract, and (e) the precision of the synthesized contract.

We highlight the following findings:

- LEASYN successfully synthesizes sound contracts for all target CPUs, and all contracts pass the unbounded verification phase.

<sup>3</sup>We translate Ibex (written in SystemVerilog) into Verilog for the analysis with LEASYN.

**Table 1: Results of LEASYN’s synthesis campaign.**

Processor	Iterations	BMC bound	ATK bound	Instr. bound	Synthesis time (hh:mm:ss)				# atoms	Precision
					Phase 1	Phase 2a	Phase 2b	Phase 3		
<b>DarkRISCV-2</b>	0	15	10	1	00:03:06	00:00:00	01:22:44	43:16:44	8	0.998
<b>DarkRISCV-3</b>	0	20	15	1	00:03:56	00:00:00	04:52:56	18:51:26	8	0.999
<b>Sodor-2</b>	0	6	5	1	00:05:17	00:00:00	00:13:31	03:01:59	8	0.998
<b>Ibex-small</b>	3	51	12	1	00:08:39	00:01:44	02:38:33	22:09:32	24	0.999
<b>Ibex-cache</b>	23	51	12	1	00:16:17	00:14:18	02:11:28	21:43:59	27	0.579
<b>Ibex-mult-div</b>	1	51	12	1	00:09:30	00:00:49	03:28:17	24:38:21	24	0.999

- The synthesized contracts have precision higher than 0.99 for all targets except for **Ibex-cache**. The low precision in the **Ibex-cache** case is due to the core having a single-line cache, which leaks whether a given address has been accessed consecutively or not. This leak cannot be precisely captured in the **B+A+BT+V** template, and LEASYN falls back to exposing the values of all memory accesses.

- For all cores, LEASYN can synthesize a contract that passes the bounded verification (without phase 3) in less than 5 hours. In several cases, the unbounded verification of the contract using LEAVE [17] dominates the overall execution time.

- For **DarkRISCV-2**, **DarkRISCV-3**, and **Sodor-2**, LEASYN can synthesize a sound contract (each one with 8 atoms) directly from the initial test cases. For **Ibex**, instead, LEASYN needs to iterate between Phases 2a and 2b, before converging on the final contract. Furthermore, the synthesized contracts contain more atoms than those for **DarkRISCV** and **Sodor-2**, which is consistent with **Ibex** being more complex than the other benchmarks.

- While **Sodor-2**, **DarkRISCV-2**, and **DarkRISCV-3** only leak through control-flow operations, contracts for the **Ibex** cores capture additional leaks. The contracts for all **Ibex** cores expose: (a) if the second operand of `div`, `divu`, `rem`, `remu` instructions is 0, (b) if memory accesses are aligned or not (only for stores for **Ibex-cache**), (c) if the core is executing a memory read or write, and (d) if the core is executing a `mul`, `mulh`, `mulhu`, or `mulhsu`. Furthermore, the contract for **Ibex-cache** exposes the entire address of memory loads (capturing the cache leak), whereas the contract for **Ibex-mult-div** exposes the logarithm of the second operand for `mul` instructions (capturing the timing leak introduced by the slow multiplication unit). Table 4 in Appendix A summarizes these leaks.

**RQ2 – Impact of the number of test cases.** LEASYN uses test cases to generate the candidate contract. To evaluate the impact of number of test cases on the synthesis process, we use LEASYN to synthesize a contract for all our benchmarks using the **B+A+BT+V** template and different number of test cases (from 0 to 2048K).

Figure 2 reports the results of our experiments. In particular, Figure 2a reports the execution time (excluded the unbounded verification)<sup>4</sup> needed to synthesize a sound contract for the different configurations, whereas Figure 2b reports the precision of the synthesized contract. We highlight the following findings:

- A very small number of initial test cases often results in a high synthesis time. The reason for this is that small sets of initial test cases often exercise only a limited behavior of the target processor, which results in a greater number of iterations of the synthesis loop.

- Increasing the number of test cases is beneficial as long as the simulation time required to find a new leak is shorter than the time required to find a new counterexample via bounded verification. At around 1024k/2048k test cases the probability of finding new leaks via simulation is low and the simulation time starts to dominate.

- A very small number of test cases often results in low precision contracts. This is because phase 2a (§ 4.2) relies on attacker-indistinguishable test cases to optimize the precision of the synthesized contract. Small sets of test cases do not provide enough information to guide synthesis and often leads to imprecise contracts.

- Increasing the number of test cases thus usually increases the precision of the synthesized contract, even though there are cases where the contract gets worse with more test cases. For all cases except **Ibex-cache**, the precision of the synthesized contract stabilizes at 1 quickly (at most after reaching 32K test cases). For **Ibex-cache**, the precision stabilizes at 0.579. This is due to the limitations of the template. All synthesized contracts for **Ibex-cache** expose the accessed memory address (to capture cache leaks), even though **Ibex-cache** implements a single-line cache so the actual leak only exposes whether a given address has been accessed consecutively or not, which cannot be captured in our template.

**RQ3 – Impact of contract templates.** The user-provided contract template defines the search space explored by LEASYN. To study the impact of different templates, we use LEASYN to synthesize a contract with increasingly more expressive templates and measure the contract’s precision. For each core and template, we run LEASYN with the same bounds as in Table 1 and with 64K test cases.

Figure 3 summarizes our results. We highlight these findings:

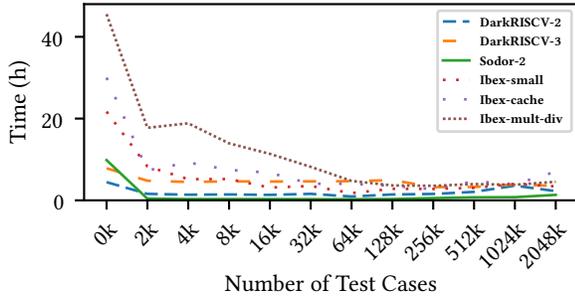
- Template **B = I + R + M** is the first template for which LEASYN can synthesize sound contracts for all target CPUs. For templates **I**, **R**, **M** in isolation, LEASYN cannot synthesize a sound contract, since no sound contract exists in these templates.

- **B**  $\rightarrow$  **B + A** : For the **Ibex** cores, additionally considering template **A** results in more precise contracts. Executing aligned and non-aligned memory accesses take different times. In template **B**, this can only be captured by exposing the entire memory address. In **B + A**, LEASYN synthesizes a more precise contract that exposes only whether memory accesses are aligned or not.

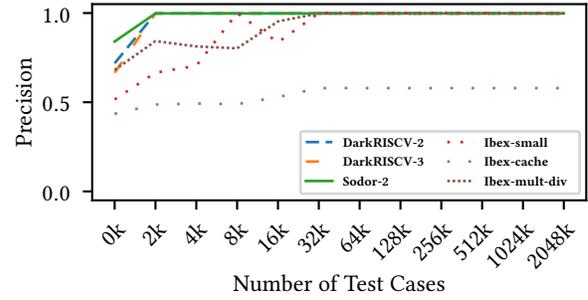
- **B + A**  $\rightarrow$  **B + A + BT** : Additionally considering the **BT** template increases precision. The branch-taken and branch-non-taken cases have different execution times. Without **BT**, LEASYN is forced to synthesize contracts that expose the values of the registers used to compute the branch. With **BT**, LEASYN synthesizes more precise contracts that only expose whether the branch is taken.

- **B + A + BT**  $\rightarrow$  **B + A + BT + V** : For the **Ibex** cores, additionally considering the template **V** further increases precision. These cores

<sup>4</sup>We omit the unbounded verification time because (a) it is roughly independent on the number of test cases and (b) it often dominates the rest of the process.



(a) Execution time (excluding unbounded verification)



(b) Precision of the synthesized contract

Figure 2: Impact of the number of test cases on contract synthesis

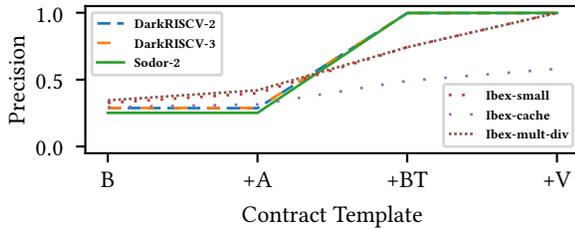


Figure 3: Impact of template expressiveness on precision

rely on a multiplication/division unit where (a) the execution time of multiplications is proportional to the logarithm of one of its operands, and (b) there is a fast path for division if the divisor equals to 0. With **V**, LEASYN can synthesize contracts that characterize such leaks, rather than having to expose entire operands.

**RQ4 – Comparison with prior sound contracts.** Prior work have manually constructed sound leakage contracts for some RISC-V cores. We focus on contracts from [17] and [23] for the **Ibex-small** CPU and we compare them (in terms of precision) to the contract synthesized by LEASYN using **B + A + BT + V** as template, an initial set of 64K test cases, and the bounds from Table 1. Like LEASYN, [17] and [23] target an attacker that observes retirement time.

We measure the precision of the three contracts against two different validation set: **different-programs** and **same-programs**. The **different-programs** validation set consists of 64K test cases where the two programs in each test case might be different, which we generate with LEASYN’s default test-generation strategy. The **same-programs** validation set consists of 64K test cases where the two programs in each test case are the same and only the initial inputs differ. We consider also **same-programs** since [17, 23] assume that the program itself is public and known to the attacker.

Table 2 summarizes the precision of all contracts. The contract generated by LEASYN is significantly more precise than those from [17, 23] when considering the **different-programs** validation set. This is unsurprising as the contract templates underlying [17, 23] are unable to express indistinguishability between different programs. However, the contract produced by LEASYN is also more precise than the ones from [17, 23] under the **same-programs** validation set. This is due to the fact that [17, 23] consider coarser

Table 2: Comparison with contracts from previous works

	LeaVe [17]	ConjunCT [23]	Our work
<b>same-programs</b>	0.774	0.544	1.000
<b>different-programs</b>	0.054	0.056	0.999

Table 3: Impact of property encoding on bounded verification

# test cases	2K	4K	8K	16K	32K	64K	128K	256K
Time (hh:mm)	<b>o</b> 8:18	5:10	5:12	3:07	3:30	1:46	2:48	2:47
	<b>b</b> 19:27	9:37	7:39	6:19	3:58	2:32	1:41	2:43
Iterations	<b>o</b> 61	40	27	13	10	7	3	3
	<b>b</b> 110	48	37	33	15	9	2	3
avg. # instr. in cex	<b>o</b> 1.1	1.3	1.1	1.2	1.2	1.3	1.7	1.7
	<b>b</b> 5.8	4.4	5.4	5.3	10	7.6	3	7

templates than **B + A + BT + V**, which require them to expose more information than needed to capture the actual leaks in **Ibex-small**.

**RQ5 – Impact of property encoding on bounded verification.** We use LEASYN to synthesize contracts for the **Ibex-small** core with **B + A + BT + V** as template. We implement and compare two encodings: the base one (denoted **b**) where the BMC always explores up to 51 cycles (same bound as in Table 1), and the optimized one (denoted **o** and described in § 4.3) where the BMC stops the exploration as soon as: (a) one instruction retires after observing a difference on the attacker trace, or (b) it reaches 51 cycles.

Table 3 summarizes the results. We highlight these findings:

- The optimized encoding results in faster (bounded) synthesis. However, the speedup is more significant with a lower number of initial test cases. Part of this speedup is due to a reduced number of iterations along the synthesis loop.

- The optimized encoding results in shorter counterexamples, i.e., programs with less instructions. Intuitively, shorter counterexamples are better since they reduce the synthesizer’s search space.

## 7 Discussion

**Verification bounds.** The choice of bounds for the bounded verification step (Phase 2b) can affect the success of the synthesis process and its runtime. On the one hand, too short a bound can result in a candidate contract that passes the bounded verification but fails the final unbounded verification step (Phase 3), since the

bounded verification failed in discovering some leaks. On the other hand, too large a bound can significantly slow down the entire synthesis process, in particular when the synthesis loops requires many iterations. In practice, in our experiments starting with a total bound of  $i * k$  and an instruction bound of  $i - 1$ , where  $k$  is the maximum number of cycles for retiring an instruction and  $i$  is the maximum number of in-flight instructions, was sufficient to successfully synthesize the contracts.

**Limitations on supported contracts.** Our formal model (§ 3) is restricted to *sequential* leakage contracts that only capture leaks generated by architectural instructions. Using the terminology of [16, 39], LEASYN only synthesizes “leakage clauses” for the sequential execution model. Furthermore, given that the leakage functions in our atoms map single architectural states to contract observations, LEASYN is geared towards synthesizing contracts for instruction-level leaks, *i.e.*, those leaks that can be attributed directly to a specific instruction rather than to a sequence of instructions. We leave the synthesis of other classes of contracts, *e.g.*, those capturing speculative leaks [16, 21], as future work.

**Dependence on user-provided contract atoms.** Our approach critically relies on user-provided contract atoms that are able to capture the leakage of the target processor soundly and precisely. For non-speculative leaks, it is fairly straightforward to define atoms able to soundly capture a processor’s leakage, by exposing any information an instruction may depend on. It is less obvious how to define atoms able to capture leakage precisely, as this requires understanding the actual leakage of the processor. It is future work to investigate how to automatically derive such atoms.

**Termination of the synthesis loop.** The contract synthesis algorithm generates a sequence of candidate contracts in a non-monotonic fashion: each new candidate need not be a superset of the previous one. For instance,  $CTR_2$  from § 2.1 is incomparable to  $CTR_1$ . Nevertheless, the synthesis loop is guaranteed to terminate, since each counterexample rules out at least one incorrect candidate, and the number of possible contracts given a template is finite.

**(Non-)determinism of synthesis.** For a fixed set of test cases, there may be multiple incomparable sound contracts that have the same optimal precision. Beyond this inherent non-determinism, however, the synthesis process is deterministic: While the bounded verification step may generate different counterexamples in different runs, the algorithm will always converge to one of the sound contract that is optimal with respect to the set of test cases generated in Phase 1.

## 8 Related work

**Leakage verification:** LeaVe [17] *verifies* a given leakage contract against a RTL processor by learning a set of inductive relational invariants. Contract Shadow Logic [40] follows a similar approach, but it uses an explicit state model checker instead of relying on invariants. While both approaches ensure soundness of contracts, these contracts are user-provided and their precision crucially relies on careful modelling by the user. In contrast, our work *synthesizes* sound and precise leakage contracts automatically, starting from a high-level contract template describing possible leakage sources.

Given a hardware design, ConjunCT [23] and its follow-up H-Houdini [24] identify a set of *safe* instructions that can be invoked

with secret arguments without timing leaks. While their focus on identifying safe instructions allows these tools to scale remarkably well, it also limits the guarantees they provide. In particular, they only guarantee that programs consisting exclusively of safe instructions will not result in leaks. However, since branching and memory instructions affect timing, they are not safe, which severely limits set of programs that can be run securely on the processor.

Ceesay-Seitz et al. [41] introduce a property called microarchitectural control-flow integrity ( $\mu$ -CFI), which checks for undue influences from instruction operands to the (micro-architectural) program counter using a taint-tracking approach. While  $\mu$ -CFI captures a number of security violations of processors (*e.g.*, certain speculative execution attacks), it can only check for violations of a fixed property and cannot express general leakage contracts.

Bloem et al. [42] verify leakage contracts for power side channels. Their technique can verify the soundness of a given contract, but it does not perform synthesis and offers no precision guarantees.

**Leakage Testing:** Revizor [39] and ScamV [43] check black-box CPUs against leakage contracts via testing. White box fuzzers like Phantom Trails [44], IntroSpectre [45], and SpecDoctor [46] can detect speculative execution bugs in CPUs by fuzzing an RTL design, but they cannot capture general contracts. While these approaches can find violations for certain classes of leaks, they cannot prove their absence and do not synthesize leakage descriptions.

**Synthesis of leakage descriptors:** RTL2muPath [47, 48] synthesizes happens-before models for microarchitectural leaks. It enumerates the paths that an instruction can take through the microarchitecture. If different paths have a different timing profiles, the instruction may leak via timing. The technique therefore records which other instructions may influence the choice of microarchitectural path, thereby obtaining a form of leakage descriptor. While the leakage descriptors in [47] are synthesized semi-automatically, unlike LEASYN, their synthesis is not guaranteed to be precise, and it is not obvious how to relate their descriptors back to software. Mohr et al. [22] synthesize leakage contracts that are precise, but cannot guarantee their soundness differently from LEASYN. Moreover, they cannot handle contract templates where a leakage function is shared by multiple atoms, which LEASYN uses when synthesizing contracts for all our benchmarks in § 6.

## 9 Conclusion

To use leakage contracts to write software secure against microarchitectural attacks, programmers require sound and precise contracts for existing processors, which we currently lack. To address this gap, we presented an approach to automatically derive such contracts directly from an RTL processor design, with limited user intervention. We implemented our approach in the LEASYN synthesis tool, which we used to synthesize sound and precise contracts for six open-source RISC-V cores. Our results indicate that contracts synthesized by LEASYN are more precise than sound contracts constructed in prior works.

## References

- [1] O. Aciğer, “Yet another microarchitectural attack: Exploiting I-cache,” in *Proceedings of the 2007 ACM workshop on Computer Security Architecture, CSAW 2007, Fairfax, VA, USA, November 2, 2007*, P. Ning and V. Atluri, Eds. ACM, 2007, pp. 11–18. [Online]. Available: <https://doi.org/10.1145/1314466.1314469>

- [2] T. Ristenpart, E. Tromer, H. Shacham, and S. Savage, "Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds," in *Proceedings of the 2009 ACM Conference on Computer and Communications Security, CCS 2009, Chicago, Illinois, USA, November 9-13, 2009*, E. Al-Shaer, S. Jha, and A. D. Keromytis, Eds. ACM, 2009, pp. 199–212. [Online]. Available: <https://doi.org/10.1145/1653662.1653687>
- [3] Y. Yarom and K. Falkner, "Flush+Reload: A High Resolution, Low Noise, L3 Cache Side-channel Attack," in *USENIX Security*, 2014.
- [4] F. Liu, Y. Yarom, Q. Ge, G. Heiser, and R. B. Lee, "Last-level cache side-channel attacks are practical," in *2015 IEEE Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015*. IEEE Computer Society, 2015, pp. 605–622. [Online]. Available: <https://doi.org/10.1109/SP.2015.43>
- [5] Y. Yarom, D. Genkin, and N. Heninger, "CacheBleed: a timing attack on openssl constant-time RSA," *J. Cryptogr. Eng.*, vol. 7, no. 2, pp. 99–112, 2017. [Online]. Available: <https://doi.org/10.1007/s13389-017-0152-y>
- [6] P. Kocher, J. Horn, A. Fogh, D. Genkin, D. Gruss, W. Haas, M. Hamburg, M. Lipp, S. Mangard, T. Prescher, M. Schwarz, and Y. Yarom, "Spectre attacks: Exploiting speculative execution," in *2019 IEEE Symposium on Security and Privacy (SP)*. Los Alamitos, CA, USA: IEEE Computer Society, may 2019. [Online]. Available: <https://doi.ieeeecomputersociety.org/10.1109/SP.2019.00002>
- [7] G. Mairsuradze and C. Rossow, "Ret2spec: Speculative execution using return stack buffers," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '18. New York, NY, USA: ACM, 2018, pp. 2109–2122. [Online]. Available: <https://doi.acm.org/10.1145/3243734.3243761>
- [8] J. Stecklina and T. Prescher, "LazyFP: Leaking FPU register state using microarchitectural side-channels," *CoRR*, vol. abs/1806.07480, 2018.
- [9] A. C. Aldaya, B. B. Brumley, S. ul Hassan, C. P. Garcia, and N. Tuveri, "Port contention for fun and profit," in *2019 IEEE Symposium on Security and Privacy, SP 2019, San Francisco, CA, USA, May 19-23, 2019*. IEEE, 2019, pp. 870–887. [Online]. Available: <https://doi.org/10.1109/SP.2019.00066>
- [10] C. Canella, D. Genkin, L. Giner, D. Gruss, M. Lipp, M. Minkin, D. Moghimi, F. Piessens, M. Schwarz, B. Sunar, J. Van Bulck, and Y. Yarom, "Fallout: Leaking data on Meltdown-resistant CPUs," in *CCS*, 2019.
- [11] M. Schwarz, M. Lipp, D. Moghimi, J. Van Bulck, J. Stecklina, T. Prescher, and D. Gruss, "ZombieLoad: Cross-privilege-boundary data sampling," in *CCS*, 2019.
- [12] S. van Schaik, A. Milburn, S. Österlund, P. Frigo, G. Mairsuradze, K. Razavi, H. Bos, and C. Giuffrida, "RIDL: Rogue in-flight data load," in *S&P*, May 2019.
- [13] D. Weber, A. Ibrahim, H. Nemat, M. Schwarz, and C. Rossow, "Osiris: Automated discovery of microarchitectural side channels," in *30th USENIX Security Symposium, USENIX Security 2021, August 11-13, 2021*, M. D. Bailey and R. Greenstadt, Eds. USENIX Association, 2021, pp. 1415–1432. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity21/presentation/weber>
- [14] L. Gerlach, D. Weber, R. Zhang, and M. Schwarz, "A security RISC: microarchitectural attacks on hardware RISC-V CPUs," in *44th IEEE Symposium on Security and Privacy, SP 2023, San Francisco, CA, USA, May 21-25, 2023*. IEEE, 2023, pp. 2321–2338. [Online]. Available: <https://doi.org/10.1109/SP46215.2023.10179399>
- [15] A. Purnal, M. Bogner, F. Piessens, and I. Verbauwhede, "ShowTime: Amplifying arbitrary CPU timing side channels," in *Proceedings of the 2023 ACM Asia Conference on Computer and Communications Security, ASIA CCS 2023, Melbourne, VIC, Australia, July 10-14, 2023*, J. K. Liu, Y. Xiang, S. Nepal, and G. Tsudik, Eds. ACM, 2023, pp. 205–217. [Online]. Available: <https://doi.org/10.1145/3579856.3590332>
- [16] M. Guarnieri, B. Köpf, J. Reineke, and P. Vila, "Hardware-software contracts for secure speculation," in *Proceedings of the 42nd IEEE Symposium on Security and Privacy*. IEEE, 2021.
- [17] Z. Wang, G. Mohr, K. von Gleissenthall, J. Reineke, and M. Guarnieri, "Specification and verification of side-channel security for open-source processors via leakage contracts," in *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security, CCS 2023, Copenhagen, Denmark, November 26-30, 2023*, W. Meng, C. D. Jensen, C. Cremers, and E. Kirda, Eds. ACM, 2023, pp. 2128–2142. [Online]. Available: <https://doi.org/10.1145/3576915.3623192>
- [18] O. Oleksenko, M. Guarnieri, B. Köpf, and M. Silberstein, "Hide and seek with spectres: Efficient discovery of speculative information leaks with random testing," in *IEEE S&P*, 2023.
- [19] G. Barthe, M. Böhme, S. Cauigi, C. Chuengsatansup, D. Genkin, M. Guarnieri, D. Mateos Romero, P. Schwabe, D. Wu, and Y. Yarom, "Testing side-channel security of cryptographic implementations against future microarchitectures," in *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '24. New York, NY, USA: Association for Computing Machinery, 2024, p. 1076–1090. [Online]. Available: <https://doi.org/10.1145/3658644.3670319>
- [20] M. Guarnieri, B. Köpf, J. F. Morales, J. Reineke, and A. Sánchez, "SPECTECTOR: Principled detection of speculative information flows," in *Proceedings of the 41st IEEE Symposium on Security and Privacy*. IEEE, 2020.
- [21] J. Hofmann, E. Vannacci, C. Fournet, B. Köpf, and O. Oleksenko, "Speculation at fault: modeling and testing microarchitectural leakage of CPU exceptions," in *Proceedings of the 32nd USENIX Conference on Security Symposium*, ser. SEC '23. USA: USENIX Association, 2023.
- [22] G. Mohr, M. Guarnieri, and J. Reineke, "Synthesizing hardware-software leakage contracts for RISC-V open-source processors," in *Design, Automation & Test in Europe Conference & Exhibition, DATE 2024, Valencia, Spain, March 25-27, 2024*. IEEE, 2024, pp. 1–6. [Online]. Available: <https://doi.org/10.23919/DATE58400.2024.10546681>
- [23] S. Dinesh, M. Parthasarathy, and C. W. Fletcher, "ConjunCT: Learning inductive invariants to prove unbounded instruction safety against microarchitectural timing attacks," in *2024 IEEE Symposium on Security and Privacy (SP)*. Los Alamitos, CA, USA: IEEE Computer Society, May 2024, pp. 3735–3753. [Online]. Available: <https://doi.ieeeecomputersociety.org/10.1109/SP54263.2024.00180>
- [24] S. Dinesh, Y. Zhu, and C. W. Fletcher, "H-Houdini: Scalable invariant learning," in *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1*, ser. ASPLOS '25. New York, NY, USA: Association for Computing Machinery, 2025, p. 603–618. [Online]. Available: <https://doi.org/10.1145/3669940.3707263>
- [25] "RISC-V Formal Verification Framework." [Online]. Available: <https://github.com/SymbioticEDA/riscv-formal>
- [26] A. Waterman, Y. Lee, D. A. Patterson, and K. Asanović, "The RISC-V Instruction Set Manual, Volume I: User-Level ISA, Version 2.0," EECSS Department, University of California, Berkeley, Tech. Rep. UCB/ECS-2014-54, May 2014. [Online]. Available: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2014/EECS-2014-54.html>
- [27] Y. Tsunoo, T. Saito, T. Suzaki, M. Shigeri, and H. Miyachi, "Cryptanalysis of DES implemented on computers with cache," in *CHES*, 2003.
- [28] G. Doychev, B. Köpf, L. Mauborgne, and J. Reineke, "CacheAudit: A tool for the static analysis of cache side channels," *ACM Trans. Inf. Syst. Secur.*, vol. 18, no. 1, pp. 4:1–4:32, 2015.
- [29] G. Barthe, P. R. D'Argenio, and T. Rezk, "Secure information flow by self-composition," *Math. Struct. Comput. Sci.*, vol. 21, no. 6, pp. 1207–1252, 2011.
- [30] "Yosys Open Synthesis Suite." [Online]. Available: <https://github.com/YosysHQ/yosys>
- [31] "Icarus Verilog." [Online]. Available: <https://github.com/steveicarus/iverilog>
- [32] "Google OR-Tools." [Online]. Available: <https://github.com/google/or-tools>
- [33] "Yices 2 SMT solver." [Online]. Available: <https://yices.csl.sri.com>
- [34] "LeaSYN implementation." [Online]. Available: <https://anonymous.4open.science/r/LeaSyn-1FFF>
- [35] "DarkRISCv processor." [Online]. Available: <https://github.com/darklife/darkriscv>
- [36] "RISC-V Sodor processor." [Online]. Available: <https://github.com/ucb-bar/riscv-sodor>
- [37] "Ibex: An embedded 32 bit RISC-V CPU core." [Online]. Available: <https://github.com/lowRISC/ibex>
- [38] "Ibex RISC-V core - multiplier/divider block." [Online]. Available: [https://ibex-core.readthedocs.io/en/latest/03\\_reference/instruction\\_decode\\_execute.html#mult-div](https://ibex-core.readthedocs.io/en/latest/03_reference/instruction_decode_execute.html#mult-div)
- [39] O. Oleksenko, C. Fetzer, B. Köpf, and M. Silberstein, "Revizor: Testing black-box CPUs against speculation contracts," in *ASPLOS*, 2022.
- [40] Q. Tan, Y. Yang, T. Bourgeat, S. Malik, and M. Yan, "RTL verification for secure speculation using contract shadow logic," in *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1, ASPLOS 2025, Rotterdam, The Netherlands, 30 March 2025 - 3 April 2025*, L. Eeckhout, G. Smaragdakis, K. Liang, A. Sampson, M. A. Kim, and C. J. Rossbach, Eds. ACM, 2025, pp. 970–986. [Online]. Available: <https://doi.org/10.1145/3669940.3707243>
- [41] K. Ceessay-Seitz, F. Solt, and K. Razavi, " $\mu$ -CFI: Formal verification of microarchitectural control-flow integrity," in *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '24. New York, NY, USA: Association for Computing Machinery, 2024, p. 213–227. [Online]. Available: <https://doi.org/10.1145/3658644.3690344>
- [42] R. Bloem, B. Gigerl, M. Gourjon, V. Hadzic, S. Mangard, and R. Primas, "Power contracts: Provably complete power leakage models for processors," in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 381–395. [Online]. Available: <https://doi.org/10.1145/3548606.3560600>
- [43] H. Nemat, P. Buiras, A. Lindner, R. Guanciale, and S. Jacobs, "Validation of abstract side-channel models for computer architectures," in *Computer Aided Verification*, S. K. Lahiri and C. Wang, Eds. Cham: Springer International Publishing, 2020, pp. 225–248.
- [44] A. de Faveri Tron, R. Isemann, H. Ragab, C. Giuffrida, K. von Gleissenthall, and H. Bos, "Phantom trails: Practical pre-silicon discovery of transient data leaks," in *Proceedings of the 34th USENIX Security Symposium*. Washington, D.C., USA: USENIX Association, 2025.

- [45] M. Ghaniyoun, K. Barber, Y. Zhang, and R. Teodorescu, "IntroSpectre: a pre-silicon framework for discovery and analysis of transient execution vulnerabilities," in *Proceedings of the 48th Annual International Symposium on Computer Architecture*, ser. ISCA '21. IEEE Press, 2021, p. 874–887. [Online]. Available: <https://doi.org/10.1109/ISCA52012.2021.00073>
- [46] J. Hur, S. Song, S. Kim, and B. Lee, "SpecDoctor: Differential fuzz testing to find transient execution vulnerabilities," in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 1473–1487. [Online]. Available: <https://doi.org/10.1145/3548606.3560578>
- [47] Y. Hsiao, N. Nikoleris, A. Khyzha, D. P. Mulligan, G. Petri, C. W. Fletcher, and C. Trippel, "RTL2M $\mu$ PATH: Multi- $\mu$ path synthesis with applications to hardware security verification," in *2024 57th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2024, pp. 507–524.
- [48] Y. Hsiao, D. P. Mulligan, N. Nikoleris, G. Petri, and C. Trippel, "Synthesizing formal models of hardware from RTL for efficient verification of memory model implementations," in *MICRO '21: 54th Annual IEEE/ACM International Symposium on Microarchitecture, Virtual Event, Greece, October 18–22, 2021*. ACM, 2021, pp. 679–694. [Online]. Available: <https://doi.org/10.1145/3466752.3480087>

## **A Additional tables**

Table 4 summarizes the leaks captured by the synthesized contracts in the experiment from **RQ1**.

## **B Property encoding in Verilog**

Figure 4 depicts the Verilog property encoding bounded contract satisfaction checked by LEASYN. The property is expressed on top of the product circuit constructed as outlined in § 4.3.

**Table 4: The table shows the leaks captured by the synthesized contract.**

Processor	Branch & Jump instr.: BRANCH_TAKEN <sup>a</sup>	div, divu, rem, remu: REG_RS2_ZERO	Load & Store instr.: distinguishable	Load & Store instr.: alignment	mul, mulh, mulhsu, mulhsu: distinguishable	Load instr.: MEM_ADDR	mul: REG_RS2_LOG2
DarkRISCv-2	●						
DarkRISCv-3	●						
Sodor-2	●						
Ibex-small	●	●	●	●	●		
Ibex-cache	●	●	●	● <sup>b</sup>	●	●	
Ibex-mult-div	●	●	●	●	●		●

<sup>a</sup> For jump instructions, there are other atoms that yield the same precision (w.r.t. to our evaluation set). Therefore, in some runs (depending on the initial test cases) the contract may expose slightly different atoms.

<sup>b</sup> Store instructions only, load instructions leak MEM\_ADDR.

```

1  wire atk_equiv = ... // expression encoding that the attacker observations are the same on the current cycle
2  wire ctr_equiv = ... // expression encoding that if the retirement predicate holds in the current cycle
3  // in the two executions, then the contract observations must be the same
4
5  // property checked
6  reg [31:0] counter = b;
7  reg [31:0] instr_counter = i;
8  reg state_atk_equiv = 1 ;
9  always @ (posedge clock) begin
10     if (counter > 0) begin
11         counter <= counter - 1;
12         state_atk_equiv <= state_atk_equiv && atk_equiv;
13     end
14     // There is an attacker conflict and a new instruction is retired
15     if (state_atk_equiv == 0 && Retire_left && Retire_right) begin
16         instr_counter <= instr_counter - 1;
17     end
18 end
19 // initialize microarchitectural states
20 assume property (init_state_equiv);
21 // assume contract equivalence until the instruction counter equals to 0
22 assume property (instr_counter == 0 || ctr_equiv);
23 // configuration of the processor
24 assume property (state_invariant);
25 // check the attacker equivalence from cycle 0 to b
26 assert property (instr_counter != 0 || state_atk_equiv);

```

**Figure 4: Encoding of the bounded contract satisfaction property in Verilog**