

Synthesizing Hardware-Software Leakage Contracts for RISC-V Open-Source Processors

Gideon Mohr
Saarland University
Saarland Informatics Campus
Saarbrücken, Germany

Marco Guarnieri
IMDEA Software Institute
Madrid, Spain

Jan Reineke
Saarland University
Saarland Informatics Campus
Saarbrücken, Germany

Abstract—Microarchitectural attacks compromise security by exploiting software-visible artifacts of microarchitectural optimizations such as caches and speculative execution. Defending against such attacks at the software level requires an appropriate abstraction at the instruction set architecture (ISA) level that captures microarchitectural leakage. Hardware-software leakage contracts have recently been proposed as such an abstraction.

In this paper, we propose a semi-automatic methodology for synthesizing hardware-software leakage contracts for open-source microarchitectures. For a given ISA, our approach relies on human experts to (a) capture the space of possible contracts in the form of contract templates and (b) devise a test-case generation strategy to explore a microarchitecture’s potential leakage. For a given implementation of an ISA, these two ingredients are then used to automatically synthesize the most precise leakage contract that is satisfied by the microarchitecture.

We have instantiated this methodology for the RISC-V ISA and applied it to the Ibex and CVA6 open-source processors. Our experiments demonstrate the practical applicability of the methodology and uncover subtle and unexpected leaks.

I. INTRODUCTION

Microarchitectural attacks [16], [26], [27], [36], [41] compromise the security of programs by exploiting software-visible artifacts of microarchitectural optimizations such as caches and speculative execution. Defending against such attacks in software is challenging: instruction set architectures (ISAs), the traditional hardware-software interface, abstract from microarchitectural details and thus do not give any guarantees w.r.t. these attacks.

Hardware-software leakage contracts [23], [39] (short: leakage contracts) have recently been proposed as a new security abstraction at the ISA level to fill this gap. Such contracts aim to capture possible microarchitectural side-channel leaks by associating leakage traces, i.e., sequences of leakage observations, with ISA-level executions. For example, a contract could expose the addresses of memory instructions as leakage observations to capture data cache leaks. Similarly, a contract could expose the operands of division instructions to capture leakage via operand-dependent latencies. Given a contract that faithfully captures the microarchitectural leakage of a processor, it is then possible to program the hardware securely by making sure that all leakage observations are independent of secrets [23].

This project has received funding from the European Research Council under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 101020415), from the Spanish Ministry of Science, Innovation, and University under the projects PID2022-142290OB-I00 ESPADA and TED2021-132464B-I00 PRODIGY and under the Ramón y Cajal grant RYC2021-032614-I, and from a gift by Intel Corporation.

However, most of today’s processor designs lack precise specifications of microarchitectural leakage. In this work, we close this gap by proposing a semi-automatic methodology for synthesizing leakage contracts from open-source processor designs.

Our methodology consists of four steps:

1) **Definition of Contract Template.** A human expert determines a set of *contract atoms* that capture potential instruction-level leakage observations. For example, a contract atom may expose the value of the register operand of a memory instruction. The set of all contract atoms forms the *contract template*, and any of its subsets is a candidate *contract*.

2) **Test-Case Generation.** A human expert devises a test-case generation strategy. Each *test case* consists of two ISA-level programs with fixed data inputs. These test cases are used to exercise and analyze a processor’s microarchitectural leakage.

3) **Evaluation of Test Cases.** Test cases are automatically evaluated on the target processor design to determine which test cases are *distinguishable*, i.e., lead to distinguishable executions for a given microarchitectural attacker. Distinguishable test cases expose actual leaks that need to be accounted for at contract level. Test cases are also automatically evaluated on the contract template to derive the set of *distinguishing atoms*, i.e., those atoms that distinguish the two programs. For instance, an atom exposing the value of register operands for memory instructions will distinguish programs accessing different addresses.

4) **Automatic Contract Synthesis.** Based on the results of the test-case evaluation, a contract, i.e. a set of contract atoms, is automatically synthesized from the distinguishing atoms associated with attacker-distinguishable test case. Our approach ensures that the synthesized contract is *satisfied* by the processor design on all test cases, i.e., it captures all leaks exposed by the test cases on the processor. Moreover, it also ensures that the synthesized contract is the *most precise* such contract, i.e., it distinguishes the fewest attacker-indistinguishable test cases.

The proliferation of the open-source RISC-V ISA [13] and its growing ecosystem is promising in this context for two reasons: (1) The simplicity and modularity of the ISA provides a good foundation for the definition of contract atoms. (2) The growing body of open-source RISC-V processor designs of varying complexity provides natural targets for contract synthesis. Leveraging the RISC-V Formal Interface [7], we have instantiated our contract-synthesis methodology for the RISC-V ISA and applied it to two open-source processor designs: Ibex [5] and CVA6 [1]. Our experiments reveal subtle

previously undocumented cases of microarchitectural leakage and demonstrate the practical applicability of our methodology.

To summarize, our main contributions are:

- A methodology for synthesizing hardware-software leakage contracts from open-source processor designs.
- The definition of a concrete contract template for the RISC-V ISA and a corresponding test-case generation strategy.
- The implementation of a contract synthesis toolchain and its application to two open-source RISC-V processor designs.

II. PRELIMINARIES

A. Instruction Set Architectures

Instruction-set architectures (ISAs) define the software-visible interface of a processor. For instance, they define the architectural state, the set of supported instructions, and how these instructions modify the architectural state. We model an ISA as a state machine capturing the execution of programs one instruction at a time. Formally, an ISA is a function $ISA : ARCH \rightarrow ARCH$ that maps each architectural state $\sigma \in ARCH$ to its successor $ISA(\sigma)$, obtained by executing the instruction to be executed in σ . To capture an entire execution, we denote by $ISA^*(\sigma)$ the sequence of states reached from σ by successive application of ISA, i.e., $ISA^*(\sigma) := [\sigma_0, \sigma_1, \sigma_2, \dots]$ with $\sigma_0 = \sigma$ and $\sigma_{i+1} = ISA(\sigma_i)$ for all $i \geq 0$.

B. Microarchitectures

Microarchitectures are concrete implementations of an ISA and often contain complex performance-enhancing optimizations. They operate at cycle granularity rather than at instruction granularity. Thus, we model them as state machines that define how the processor's state evolves cycle-by-cycle. Formally, a microarchitecture is a function $IMPL : IMPLSTATE \rightarrow IMPLSTATE$ on the set of microarchitectural states $IMPLSTATE = ARCH \times \mu ARCH$ that captures how the microarchitectural state evolves from one cycle to the next. Each microarchitectural state $\sigma \in IMPLSTATE$ consists of an architectural part $\sigma_{ISA} \in ARCH$ and a microarchitectural part $\sigma_{IMPL} \in \mu ARCH$ modeling the state of microarchitectural components like caches and predictors. Similarly to $ISA^*(\sigma)$, $IMPL^*(\sigma)$ is the sequence of microarchitectural states reached from σ by successive applications of IMPL.

C. Microarchitectural Attackers

Programs that operate on secrets may leave traces of these secrets in the microarchitectural state. Microarchitectural attacks extract information about secrets from the microarchitectural state by leveraging software-visible side-effects [16], [26], [27], [36], [41], mostly affecting a program's execution time. In this paper, we consider (passive) microarchitectural attackers that can extract information from the microarchitectural state. Formally, we model a microarchitectural attacker as a function $\mu ATK : IMPLSTATE \rightarrow ATKOBS$ that maps microarchitectural states to attacker observations. Common attacker models, like the one exposing the timing of instruction retirement [35] or the one exposing the final state of caches [18], [41], can be instantiated in our setting. Given an attacker model, two executions $IMPL^*(\sigma)$ and $IMPL^*(\sigma')$ are *attacker distinguishable* if $\mu ATK(IMPL^*(\sigma)) \neq \mu ATK(IMPL^*(\sigma'))$, where we lift μATK to sequences by applying it to each element of the sequence.

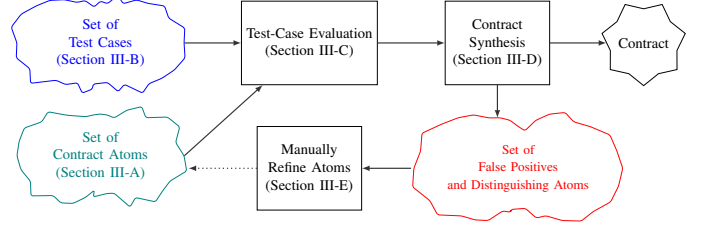


Fig. 1. High-level steps of our contract-synthesis methodology.

D. Hardware-Software Leakage Contracts

The goal of leakage contracts [23] is to capture microarchitectural side-channel leakage at the ISA level to allow reasoning about side-channel security of programs without having to explicitly consider their microarchitectural execution. To this end, contracts associate ISA-level executions with leakage traces, i.e., sequences of leakage observations.

A *contract* is a function $CTR : ARCH \rightarrow CTR OBS$ mapping architectural states to leakage observations. As an example, a contract can expose the addresses of memory instructions to capture leaks via the data cache. Two architectural states σ, σ' are *contract distinguishable* if $CTR(ISA^*(\sigma)) \neq CTR(ISA^*(\sigma'))$, where we lift CTR to sequences by applying it to all elements.

Microarchitecture IMPL *satisfies* contract CTR for instruction-set architecture ISA under attacker model μATK if all contract-indistinguishable executions are also attacker indistinguishable:

$$\begin{aligned} \forall \sigma, \sigma' \in IMPLSTATE : \sigma_{IMPL} = \sigma'_{IMPL} \implies \\ CTR(ISA^*(\sigma_{ISA})) = CTR(ISA^*(\sigma'_{ISA})) \implies \\ \mu ATK(IMPL^*(\sigma)) = \mu ATK(IMPL^*(\sigma')) \end{aligned}$$

Note that we require the microarchitectural parts of σ and σ' to initially be the same. Otherwise, the executions could be attacker distinguishable due to initial differences rather than due to leakage during the execution.

The benefit of contract satisfaction is that it allows reasoning about side-channel security directly at the ISA level: A program that does not leak any secret information according to the contract is also guaranteed not to leak any secret information to attackers on *any* microarchitecture that satisfies the contract [23].

III. CONTRACT-SYNTHESIS METHODOLOGY

Our goal is to synthesize contracts directly from a processor's register-transfer level (RTL) design. Figure 1 overviews our contract-synthesis methodology, whose steps we describe next.

A. Contract Templates

The basic building blocks of contracts are *contract atoms*. Contract atoms capture potential leakage at the instruction level. Formally, a contract atom A is a tuple (π_A, τ_A, ϕ_A) where $\pi_A : ARCH \rightarrow \mathbb{B}$ determines whether the contract observation is applicable in a particular architectural state, τ_A is an identifier for the leakage source, and $\phi_A : ARCH \rightarrow O_A$ is the observation function that maps an architectural state to the atom's observation. For example, a contract atom A could expose the divisor of division instructions: π_A would hold whenever the instruction to execute is a division, τ_A would identify the leakage source, and ϕ_A would return the divisor's value.

A *contract template* T is a set of contract atoms. A subset $S \subseteq T$ of the template defines a candidate *contract* CTR_S as follows: $\text{CTR}_S(\sigma) = \{(\tau_A, \phi_A(\sigma)) \mid A \in S \wedge \pi_A(\sigma)\}$, *i.e.*, it exposes the leakage source and the observation for each atom that is applicable in a given state. Note that multiple atoms may generate observations for the same leakage source.

B. Test-Case Generation

Our goal is to find a contract that is satisfied by the microarchitecture while being as precise as possible, *i.e.*, distinguishing as few executions as possible. This requires (1) to determine which contracts from the template are satisfied by the microarchitecture, and (2) to measure their precision.

Ideally, one would like to formally verify contract satisfaction. However, contract verification techniques [39] do not yet scale to processors of the complexity of, *e.g.*, the CVA6 core. Thus, we resort to a testing-based approach and evaluate contract satisfaction on systematically-generated test cases. Formally, a *test case* is a pair (σ, σ') of microarchitectural states such that $\sigma_{\text{IMPL}} = \sigma'_{\text{IMPL}}$. A test case is *attacker distinguishable* if $\mu\text{ATK}(\text{IMPL}^*(\sigma)) \neq \mu\text{ATK}(\text{IMPL}^*(\sigma'))$. A test case is *atom distinguishable* for an atom A if $\text{CTR}_{\{A\}}(\text{ISA}^*(\sigma_{\text{ISA}})) \neq \text{CTR}_{\{A\}}(\text{ISA}^*(\sigma'_{\text{ISA}}))$. Note that a test case is contract distinguishable for contract CTR_S if it is atom distinguishable for at least one atom in S . Thus, checking atom distinguishability for all atoms is sufficient to characterize contract distinguishability for the entire contract template.

Test cases also allow to measure the precision of a contract. For this, we adopt the standard notion of precision used in the evaluation of binary classifiers: $\text{Precision} = \frac{TP}{TP+FP}$, where TP is the number of true positives, *i.e.*, test cases that are contract distinguishable *and* attacker distinguishable, and FP is the number of false positives, *i.e.*, test cases that are contract distinguishable *but not* attacker distinguishable. Higher precision contracts are desirable as they rule out fewer programs at the contract level that could actually be executed securely on the processor.

Our methodology requires a human expert to devise a test-case generation strategy used to generate the set of test cases TC .

C. Test-Case Evaluation

Given a set of test cases TC , we need to determine for each test case (1) whether it is attacker distinguishable, and (2) which contract atoms distinguish it. Attacker distinguishability can be determined via simulation of the microarchitecture with a suitable attacker model. Similarly, atom distinguishability can be determined by evaluating all contract atoms in parallel on top of a simulation of the instruction set architecture.

The test-case evaluation phase has two outputs: (1) The set of attacker-distinguishable test cases $\text{Dist} \subseteq TC$. (2) For each test case t , the set of distinguishing atoms $\text{distinguishing}(t) \subseteq T$.

D. Contract Synthesis

We now show how to use integer linear programming (ILP) to synthesize a contract from the template that distinguishes all attacker-distinguishable test cases TC and maximizes precision.

The ILP uses a boolean variable s_A for each atom A in the contract template T to encode whether the atom is selected to be

part of the synthesized contract or not. Further, the ILP uses a boolean variable c_t for each attacker-indistinguishable test case $t \in \text{Indist} = TC \setminus \text{Dist}$, which, using constraints detailed below, is forced to be 1 for test cases that are contract distinguishable and thus are false positives.

Maximizing precision is equivalent to minimizing the number of false positives, as the number of true positives is the same for all correct contracts. Thus, the objective function of the ILP is $\min \sum_{t \in \text{Indist}} c_t$. To ensure that only correct contracts are considered, we introduce the following constraint for each attacker-distinguishable test case $t \in \text{Dist}$: $\sum_{A \in \text{distinguishing}(t)} s_A \geq 1$, *i.e.*, at least one atom that distinguishes t must be selected for the contract. For each attacker-indistinguishable test case $t \in \text{Indist}$ and each contract atom $A \in \text{distinguishing}(t)$ we further introduce the constraint: $s_A \leq c_t$. This ensures that c_t can only be 0 for test cases that are not contract distinguishable.

From the ILP's solution, we extract the synthesized contract via the variables s_A and the false-positive test cases via the variables c_t .

E. Refinement of the Contract Template

In addition to returning a contract that maximizes precision, our implementation also returns a ranking of the contract atoms according to the number of false positives caused by their inclusion in the contract and the corresponding test cases. Inspecting these test cases allows a human expert to identify contract atoms that should be refined to obtain a more precise contract.

IV. INSTANTIATING THE METHODOLOGY FOR RISC-V

We instantiated the above methodology for the RISC-V instruction set, more precisely its I and M subsets.

A. RISC-V Contract Template

Instructions of different type often show different leakage behavior, *e.g.*, while divisions may leak the second operand, additions usually do not. Thus, our contract template contains a contract atom for each instruction type and each applicable potential leakage source. For each atom, π_A determines whether the current instruction has the given type, τ_A is an identifier for the leakage source, and ϕ_A extracts the corresponding leakage from the architectural state. For example, the contract atom $(\pi^{\text{DIV}}, \text{REG_RS2}, \phi^{\text{REG_RS2}})$, where π^{DIV} detects divisions and $\phi^{\text{REG_RS2}}$ determines the value of register RS2, allows to capture leakage of the second operand of division instructions. Similarly, the contract atom $(\pi^{\text{ADD}}, \text{REG_RS2}, \phi^{\text{REG_RS2}})$ allows to capture leakage of the second operand of addition instructions, sharing the same leakage source as the corresponding division atom.

We first defined a base template capturing the architectural state that directly influences the execution of an instruction:

- **Instruction leakages (IL)** expose values from an instruction's encoding: The operation OP , the destination and source registers RD , RS1 , and RS2 , and the immediate value IMM .
- **Register leakages (RL)** expose the values of registers: The values of the source registers REG_RS1 and REG_RS2 before execution and the final value of the destination register REG_RD .
- **Memory leakages (ML)** expose the memory addresses and memory contents accessed by an instruction: MEM_R_ADDR

and MEM_W_ADDR expose the accessed addresses, whereas MEM_R_DATA and MEM_W_DATA expose the accessed content.

During evaluation, we noticed that, while the above categories are sufficient, the precision of synthesized contracts can be improved by adding the following contract atoms:

- **Alignment leakages (AL)** expose the alignment of a memory access: IS_WORD_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 address are not 11.
- **Branch leakages (BL)** expose whether a branch is taken or not: This can be BRANCH_TAKEN for taken branches and NEW_PC to indicate the target of the branch. The latter one also applies to unconditional jumps.
- **Data-dependency leakages (DL)** expose data dependencies in the program: RAW_RS1_n, RAW_RS2_n, RAW_RD_n, and WAW_n, respectively, indicate Read-After-Write dependencies on register RS1, RS2, and RD and Write-After-Write dependencies within a distance of n instructions.

We remark that not all atoms are applicable to every instruction type, *e.g.*, some instructions do not have an immediate value. With a maximum distance of $n = 4$, the above template results in a total of 762 atoms.

B. Test-Case Generation

The set of test cases used for the simulation affects the synthesized contract, since the synthesis algorithm cannot account for leakages that are not exposed by the test cases. Thus, ideally the test-case generation algorithm should consider potential microarchitectural implementations and their potential leakage.

The test-case generation method we use is simple but has shown effective on the tested targets. Note that we only generate the architectural state and fix the initial microarchitectural part of the microarchitectural state to allow reusing test cases for different microarchitectures. Each test case consists of two programs which aim to be differentiable by one specific contract atom.

Each program consists of three parts, of which only the second part differs among the two programs in the test case: (1) First, we initialize every architectural register to a randomly selected value. (2) Next, we try to trigger the leakage of the contract atom we want to test for. For this, we generate a random instance of the given instruction type. We then derive two sequences of instructions that could be differentiated by the given atom, *e.g.*, if we want to test whether the immediate value leaks, we alter the immediate value or if we want to test whether the memory contents that are read leak, we insert an earlier instruction in both programs writing different values to the same address. (3) Finally, we append randomly selected instructions that aim to surface the leakage and to make sure all instructions from (2) are executed completely.

C. Attacker Model

We consider an attacker that can observe the timing of instruction retirements at cycle granularity. Our implementation of this attacker model relies on the RISC-V Formal Interface (RVFI) [7], a standard interface for RISC-V processors exposing information about the execution of programs, including the cycles at which instructions retire. RVFI simplifies the

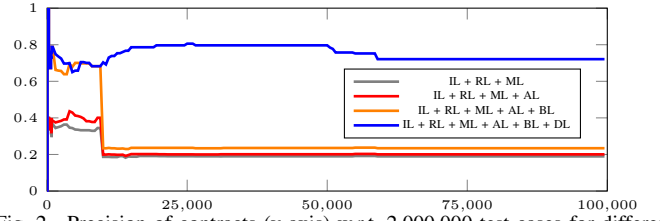


Fig. 2. Precision of contracts (y-axis) w.r.t. 2,000,000 test cases for different contract templates starting from the base contract (IL+RL+ML) depending on the number of test case (x-axis) used for contract synthesis.

interaction with the microarchitecture and allows reusing most of our implementation for any processor implementing it.

By instantiating two instances of the respective core and simulating the execution of both programs of a test case in parallel, our implementation determines attacker distinguishability by comparing the cycles at which instructions retire. Other attacker models could be implemented either using the RVFI or by exposing the required signals directly from the design.

D. Identifying Distinguishing Atoms

To determine the distinguishing atoms for a test case, we need to evaluate all contract atoms on top of an architectural simulation of both programs of a test case.

However, as long as the microarchitecture correctly implements the ISA, the sequence of architectural states can also be extracted from the microarchitectural states upon instruction retirement [39]. Thus, our implementation piggy backs on the microarchitectural simulation from §IV-C and it uses the RVFI to extract the architectural state whenever an instruction retires. The simulation produces a VCD waveform, from which we derive a test case’s distinguishing atoms. Adaptation to other cores that support the RVFI should be straightforward.

V. EXPERIMENTAL EVALUATION

We evaluate our methodology on two open-source RISC-V cores, Ibex [5] and CVA6 [1], both in a configuration implementing RV32IMC, for which we obtain precise contracts. Additionally, we analyze the impact of the test-case generation method and the granularity of the contract template.

A. Experimental Setup

To enable the simulation of test cases we first convert the processor sources from SystemVerilog to standard Verilog using sv2v [9] for Ibex and Yosys [11] with the Yosys-systemverilog plugin [10] for CVA6. and then compile and execute the two cores embedded in a testbench using Icarus Verilog [6]. In both cases, some manual adjustments to the processors were needed to eliminate unsupported SystemVerilog constructs.

Both cores support the RVFI, which we use to extract the attacker observations and distinguishing atoms as described in §IV-C and §IV-D. The contract synthesis algorithm is implemented in Java using Google’s OR-Tools [4] ILP solver.

B. Ibex Core

We first synthesized a contract for the Ibex core using the base contract template described in Section IV-A. To this end, we used 100,000 test cases for synthesis and evaluated the performance

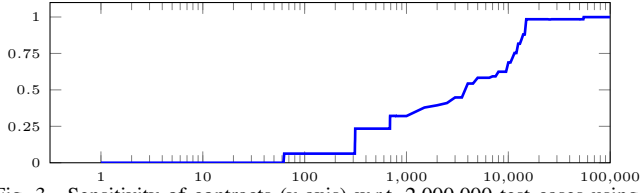


Fig. 3. Sensitivity of contracts (y-axis) w.r.t. 2,000,000 test cases using the full contract template (IL+RL+ML+AL+BL+DL) depending on the number of test cases (x-axis) used for contract synthesis. Note the logarithmic x-axis.

of the obtained contract using an independent set of 2,000,000 test cases. The results allowed us to refine the contract template (as described in Section IV-A) by adding alignment leakages (AL), branch leakages (BL), and data-dependency leakages (DL).

Figure 2 shows the precision of the obtained contracts for different contract templates in terms of the number of test cases used for synthesis (from 0 to 100,000). Some leakage sources are only discovered after a while, *e.g.*, the first data-dependency leakages are discovered after about 10,000 test cases, which explains the drop in precision for templates that do not include DL at that point. We observe that the refined templates lead to an increase in precision as they allow for more fine-grained leakage contracts. Data-dependency leakages (DL), in particular, enable significant precision improvements.

Figure 3 shows the sensitivity of the synthesized contracts for the full contract template depending on the number of test cases used for synthesis. We adopt the standard notion of sensitivity used in the evaluation of binary classifiers: $Sensitivity = \frac{TP}{TP+FN}$, where FN is the number of false negatives. Initially, sensitivity increases rapidly, as additional test cases frequently reveal new sources of leakage. After the first 15,000 test cases, however, the curve flattens. The final contract has a sensitivity of 99.93%.

Finally, we synthesized a contract, summarized in Table I, using the larger set of 2,000,000 test cases. In total, the synthesized contract includes 82 atoms. ● indicates that all instructions in this category show leakages in the respective category (as introduced in Section IV-A), ○ indicates that no leakages were found, ◐ indicates that there are some leakages in this category, but not all instructions have these leakages, and - indicates that the atom does not apply to this category. The distance n in the DL category is always 1.

Next, we overview some of our findings: (1) Our experiments and the synthesized contracts show that the Ibex core leaks whether memory accesses are aligned or not. Indeed, the documentation of the Ibex core confirms that requests on the memory interface are always aligned to a word boundary, which is compatible with our observation, but the impact on information leakage had not been established before. (2) We also discovered that the timing of branch instructions depend on whether the branch is taken or not taken even if the branch target is the same in both cases, *e.g.*, BEQ r1 r2 4 jumps to the next instruction independently of r1 and r2.

C. CVA6 Core

We similarly analyzed the CVA6 core. A summary of the contract obtained using 500,000 test cases is shown in Table II. In total the synthesized contract consists of 77 atoms.

TABLE I
SYNTHESIZED CONTRACT FOR THE IBEX PROCESSOR.

	IL	RL	ML	AL	BL	DL
Arithmetic instructions	●	●	-	-	-	●
Division, Remainder	○	●	-	-	-	●
Multiplication	●	○	-	-	-	●
Loads	●	○	○	●	-	○
Stores	●	○	○	○	-	○
Branch instructions	●	○	-	-	●	○

TABLE II
SYNTHESIZED CONTRACT FOR THE CVA6 PROCESSOR.

	IL	RL	ML	AL	BL	DL
Arithmetic instructions	●	●	-	-	-	●
Division, Remainder	●	●	-	-	-	●
Multiplication	○	●	-	-	-	●
Loads	●	○	○	○	-	●
Stores	○	●	○	○	-	○
Branch instructions	○	○	-	-	●	●

The CVA6 core uses a more complex memory interface that does not expose anything about a specific memory access in the analyzed setup. We remark that, even though the CVA6 features a simple branch predictor, the contract template originally composed for the Ibex, was sufficient to capture the CVA6's leakage.

As the CVA6 core has a deeper and more complex pipeline than the Ibex, data and control dependencies can have a more pronounced effect in CVA6, which is reflected in the synthesized contract: While the data dependencies all have a distance of $n = 1$ as forwarding is effective here, we observe distances of up to $n = 4$ due to control dependencies upon branch instructions.

D. Computation Time

Table III compares the runtime of the contract synthesis algorithm for the two cores. The test-case simulation for CVA6 is much slower than on Ibex due to the higher complexity of the design and due to the translation into standard Verilog with Yosys. However, the contract synthesis algorithm is still able to synthesize a contract for CVA6 in a reasonable amount of time.

VI. RELATED WORK

Leakage contracts: Leakage contracts [23], which our methodology builds on, are an abstraction for capturing timing leaks at ISA level. The leaks captured in a contract are formally connected with the actual timing leaks in a hardware implementation. There are several tools for reasoning about contracts: for verifying contract satisfaction against RTL processor designs [14], [39], for detecting contract violations for black-box

TABLE III
PERFORMANCE MEASUREMENTS FOR CONTRACT SYNTHESIS USING 100,000 TEST CASES ON AN AMD RYZEN THREADRIPPER PRO 5995WX WITH 512 GB OF RAM USING UP TO 128 THREADS.

	Ibex	CVA6
Compilation of the testbench	7.2 s	2123 s
Simulation of a single test case	0.2 s	88 s
Extraction of distinguishing atoms	175 ms	75 ms
Computation of the contract	15.6 s	16 s
Overall computation time	5.3 min	1175 min

CPUs [15], [24], [30]–[32], and for verifying program security w.r.t. a given contract [19], [22]. Contracts synthesized using our methodology may serve as inputs to these tools.

Beyond leakage contracts, several proposals extend ISAs with leakage guarantees. Yu et al. [42] propose a data-oblivious ISA extension that allows to specify whether or not instruction operands are safe (i.e., never leaked microarchitecturally). Similarly, CPU vendors implement data-oblivious execution modes that, when activated, ensure that the execution time of specific instructions is independent of their operands [2], [3], [8]. All these guarantees can be formalized as leakage contracts.

Modeling timing leaks: There are multiple formal approaches for studying timing leaks. Most of them capture leaks at program level: from simple models associated with “constant-time programming” [12], [29] to more complex ones capturing leaks of transient instructions [17], [19], [22], [33], [37]. Other approaches, instead, capture leaks on simplified processor models [21], [38]. Our work enables synthesizing program-level models from actual processor implementations, while previous approaches rely on manually written models that have no formal connection to particular concrete processors.

Detecting leaks through testing: Revizor [24], [31], [32] and Scam-V [15], [30] search for contract violations for black-box CPUs and can only be applied post-silicon. Other approaches [20], [28], [40] detect leaks by analyzing hardware measurements without relying on leakage contracts but, again, apply only post-silicon. Finally, SpecDoctor [25] and SIGFuzz [34] focus on detecting microarchitectural leaks in RTL designs in the pre-silicon phase. In contrast to our work, they do not aim to characterize leakage at ISA level.

VII. CONCLUSION

ISA-level models of the information leaked microarchitecturally by processors are a critical component for the development of systems resistant to microarchitectural attacks. In this paper, we showed how to semi-automatically synthesize such models, in the form of hardware-software leakage contracts, directly from RTL processor designs. This allowed us to derive ISA-level descriptions of leakage for two open-source RISC-V cores, Ibex and CVA6, which so far lacked a formal specification of their microarchitectural leakage properties.

The contract synthesis toolchain is available open source at <https://github.com/hw-sw-contracts/riscv-contract-synthesis> and archived at <https://doi.org/10.5281/zenodo.10491534>.

REFERENCES

- [1] “CVA6: A 6-stage RISC-V CPU core capable of booting linux.” [Online]. Available: <https://github.com/openhwgroup/cva6>
- [2] “Data operand independent timing instruction set architecture (ISA) guidance.” [Online]. Available: <https://www.intel.com/content/www/us/en/developer/articles/technical/software-security-guidance/best-practices/data-operand-independent-timing-isa-guidance.html>
- [3] “DIT, data independent timing.” [Online]. Available: <https://developer.arm.com/documentation/ddi0595/2020-12/AArch64-Registers/DIT-Data-Independent-Timing>
- [4] “Google’s Operations Research tools.” [Online]. Available: <https://github.com/google/or-tools>
- [5] “Ibex: An embedded 32 bit RISC-V CPU core.” [Online]. Available: <https://github.com/lowRISC/ibex>
- [6] “The Icarus Verilog compilation system.” [Online]. Available: <https://github.com/steveicarus/iverilog>
- [7] “RISC-V Formal Verification Framework.” [Online]. Available: <https://github.com/SymbioticEDA/riscv-formal>
- [8] “RISC-V Zkt “constant time” instruction list.” [Online]. Available: <https://github.com/rvkrypto/riscv-zkt-list/blob/main/zkt-list.adoc>
- [9] “sv2v: SystemVerilog to Verilog.” [Online]. Available: <https://github.com/zachjs/sv2v>
- [10] “systemverilog-plugin – SystemVerilog support for Yosys.” [Online]. Available: <https://github.com/chipsalliance/systemverilog-plugin>
- [11] “Yosys Open SYnthesis Suite.” [Online]. Available: <https://github.com/YosysHQ/yosys>
- [12] J. B. Almeida, M. Barbosa, G. Barthe, F. Dupressoir, and M. Emmi, “Verifying constant-time implementations,” in *USENIX Security*, 2016.
- [13] K. Asanović and D. A. Patterson, “Instruction sets should be free: The case for RISC-V,” Tech. Rep. UCB/EECS-2014-146, Aug 2014.
- [14] R. Bloem et al., “Power contracts: Provably complete power leakage models for processors,” in *CCS*, 2022.
- [15] P. Buiras, H. Nemati, A. Lindner, and R. Guanciale, “Validation of side-channel models via observation refinement,” in *MICRO*. ACM, 2021.
- [16] J. V. Bulck et al., “Foreshadow: Extracting the keys to the Intel SGX kingdom with transient out-of-order execution,” in *USENIX Security*, 2018.
- [17] S. Cauligi et al., “Constant-time foundations for the new Spectre era,” in *PLDI*, 2020.
- [18] G. Doychev et al., “CacheAudit: A tool for the static analysis of cache side channels,” *ACM TISSEC*, vol. 18, no. 1, pp. 4:1–4:32, 2015.
- [19] X. Fabian, M. Patrignani, and M. Guarnieri, “Automatic detection of speculative execution combinations,” in *CCS*, 2022.
- [20] B. Gras et al., “ABSynthe: Automatic blackbox side-channel synthesis on commodity microarchitectures,” in *NDSS*, 2020.
- [21] R. Guanciale, M. Balliu, and M. Dam, “InSpectre: Breaking and fixing microarchitectural vulnerabilities by formal analysis,” in *CCS*, 2020.
- [22] M. Guarnieri et al., “SPECTECTOR: Principled detection of speculative information flows,” in *IEEE S&P*, 2020.
- [23] M. Guarnieri, B. Köpf, J. Reineke, and P. Vila, “Hardware-software contracts for secure speculation,” in *IEEE S&P*, 2021.
- [24] J. Hofmann et al., “Speculation at fault: Modeling and testing microarchitectural leakage of CPU exceptions,” in *USENIX Security*, 2023.
- [25] J. Hur, S. Song, S. Kim, and B. Lee, “SpecDoctor: Differential fuzz testing to find transient execution vulnerabilities,” in *CCS*, 2022.
- [26] P. Kocher et al., “Spectre attacks: Exploiting speculative execution,” in *IEEE S&P*, 2019.
- [27] M. Lipp et al., “Meltdown: Reading kernel memory from user space,” in *USENIX Security*, 2018.
- [28] D. Moghimi, M. Lipp, B. Sunar, and M. Schwarz, “Medusa: Microarchitectural data leakage via automated attack synthesis background superscalar memory architecture,” in *USENIX Security*, 2020.
- [29] D. Molnar, M. Piotrowski, D. Schultz, and D. Wagner, “The program counter security model: Automatic detection and removal of control-flow side channel attacks,” in *ICISC*, 2005.
- [30] H. Nemati et al., “Validation of Abstract Side-Channel Models for Computer Architectures,” in *CAV*, 2020.
- [31] O. Oleksenko et al., “Hide and seek with spectres: Efficient discovery of speculative information leaks with random testing,” in *IEEE S&P*, 2023.
- [32] O. Oleksenko, C. Fetzer, B. Köpf, and M. Silberstein, “Revizor: Testing black-box CPUs against speculation contracts,” in *ASPLOS*, 2022.
- [33] M. Patrignani and M. Guarnieri, “Exorcising spectres with secure compilers,” in *CCS*, 2021.
- [34] C. Rajapaksha et al., “SIGFuzz: A framework for discovering microarchitectural timing side channels,” in *DATE*, 2023.
- [35] Y. Tsunoo, T. Saito, T. Suzaki, M. Shigeri, and H. Miyauchi, “Cryptanalysis of DES implemented on computers with cache,” in *CHES*, 2003.
- [36] S. van Schaik et al., “RIDL: Rogue in-flight data load,” in *S&P*, 2019.
- [37] M. Vassena et al., “Automatically eliminating speculative leaks from cryptographic code with Blade,” in *POPL*, 2021.
- [38] J. R. S. Vicarte et al., “Opening Pandora’s box: A systematic study of new vulnerabilities microarchitecture can leak private data,” in *ISCA*, 2021.
- [39] Z. Wang et al., “Specification and verification of side-channel security for open-source processors via leakage contracts,” in *CCS*, 2023.
- [40] D. Weber et al., “Osiris: Automated discovery of microarchitectural side channels,” in *USENIX Security*, 2021.
- [41] Y. Yarom and K. Falkner, “Flush+Reload: A High Resolution, Low Noise, L3 Cache Side-channel Attack,” in *USENIX Security*, 2014.
- [42] J. Yu et al., “Data oblivious ISA extensions for side channel-resistant and high performance computing,” in *NDSS*, 2019.