

Hardware-Software Contracts for Secure Speculation

Marco Guarnieri*, Boris Köpf†, Jan Reineke‡, and Pepe Vila*

**IMDEA Software Institute* †*Microsoft Research* ‡*Saarland University*

Abstract—Since the discovery of Spectre, a large number of hardware mechanisms for secure speculation have been proposed. Intuitively, more defensive mechanisms are less efficient but can securely execute a larger class of programs, while more permissive mechanisms may offer more performance but require more defensive programming. Unfortunately, there are no hardware-software contracts that would turn this intuition into a basis for principled co-design.

In this paper, we put forward a framework for specifying such contracts, and we demonstrate its expressiveness and flexibility: On the hardware side, we use the framework to provide the first formalization and comparison of the security guarantees provided by a representative class of mechanisms for secure speculation.

On the software side, we use the framework to characterize program properties that guarantee secure co-design in two scenarios traditionally investigated in isolation: (1) ensuring that a benign program does not leak information while computing on confidential data, and (2) ensuring that a potentially malicious program cannot read outside of its designated sandbox.

Finally, we show how the properties corresponding to both scenarios can be checked based on existing tools for software verification, and we use them to validate our findings on executable code.

I. INTRODUCTION

Speculative execution avoids expensive pipeline stalls by predicting the outcome of branching (and other) decisions, and by continuing the execution based on these predictions. When a prediction turns out to be incorrect, the processor rolls back the effect of speculatively executed instructions on the architectural state consisting of registers, flags, and main memory.

However, the microarchitectural state, which includes the content of various caches and buffers, is not (or only partially) rolled back. This side effect can leak information about the speculatively accessed data and thus violate confidentiality, see Fig. 1a. Spectre attacks [1], [2] demonstrate that this vulnerability affects all modern general-purpose processors and poses a serious threat for platforms with multiple tenants.

A multitude of hardware mechanisms for secure speculation have been proposed. They are based on a number of basic ideas, such as delaying load operations until they cannot be squashed [3], delaying operations that depend on speculatively loaded data [4], [5], limiting the effect of speculatively executed instructions [6], [7], [8], [9], or rolling back the microarchitectural state when a misprediction is detected [10].

Intuitively, more defensive mechanisms are less efficient but can securely execute a larger class of programs, while more permissive mechanisms offer more performance but require more defensive programming. We refer to this intuition as (*).

For example, consider the variant of Spectre v1 shown in Fig. 1b, where array *A* is accessed before the bounds check.

<pre> 1 if (y < size_A) 2 x = A[y]; 3 temp &= B[x * 64]; </pre>	<pre> 1 x = A[y]; 2 if (y < size_A) 3 temp &= B[x * 64]; </pre>
(a) Program P_1	(b) Program P_2

Fig. 1: Program P_1 is the vanilla Spectre v1 example, where $A[y]$ can be speculatively read, and leaked into the data cache via an access to array *B*, for $y \geq \text{size_A}$. Program P_2 , is a variant where $A[y]$ is accessed non-speculatively before the bounds check but the leak occurs during speculative execution.

Mechanisms delaying loads until they cannot be squashed [3] prevent speculatively leaking $A[y]$, for $y \geq \text{size_A}$. In contrast, more permissive mechanisms that delay only loads depending on speculatively accessed data [4], [5] do *not* prevent the leak, because $A[y]$ is accessed non-speculatively.

While the performance characteristics of secure speculation mechanisms are well-studied, there has been little work on (1) characterizing the security guarantees they provide, and in particular on (2) investigating how these guarantees can be effectively leveraged by software to achieve global security guarantees.¹ That is, we lack hardware-software contracts that support principled co-design for secure speculation, and that would formalize the intuition (*) described above.

Contracts: In this paper, we put forward a framework for specifying such contracts, based on three basic building blocks: an ISA language, a model of the microarchitecture, and an adversary model specifying which microarchitectural components (such as caches or branch predictor state) are observable via side-channels.

Contracts specify which program executions a side-channel adversary can distinguish. A contract in our framework is defined in terms of *executions* and *observations* made on these executions, and is formalized in terms of a labelled ISA semantics. A CPU satisfies a contract if, whenever two program executions agree on all observations, they are guaranteed to be indistinguishable by the adversary at the microarchitectural level. The contract semantics can mandate exploration of mispredicted paths, effectively requiring agreement on observations corresponding to transient instructions.

Secrets at the program level must not affect contract observations, because then they can become visible to the adversary. Hence, contracts exposing more observations correspond to

¹A notable exception to (1) is STT [5], which is backed by a security property that guarantees the confidentiality of speculatively loaded data. However, this property alone does not provide an actionable basis for (2), as the code snippet in Fig. 1b is simply declared to be “out of scope” [5, Section 4].

hardware with weaker security guarantees, whereas contracts exposing fewer observations correspond to hardware with stronger guarantees. The extreme case is a contract with no observations, which is satisfied by an ideal side-channel resilient platform that can securely execute every program.

Software Side: Our framework provides a basis for deriving requirements that *software* needs to satisfy to run securely on a specific platform. For deriving such requirements, we consider two scenarios typically considered in the literature:

- In the first scenario, called “constant-time programming”, the goal is to ensure that a benign program, such as a cryptographic algorithm, does not leak information while computing on confidential data.
- In the second scenario, which we call “sandboxing”, the goal is to restrict the memory region that a potentially malicious program, such as a Web application, can read from.²

For each scenario, we identify program-level properties that guarantee security on hardware that satisfies a given contract. We stress that secure speculation approaches usually *either* consider constant-time programming [12], [13], [14], [15] *or* sandboxing [16], [17]. In contrast, our framework supports *both* goals through program-level properties.

We provide tool support for automatically checking if programs are secure in both scenarios. For this, we extend a static analysis tool for detecting speculative leaks [12] to cater for different contracts, and we use it to validate all examples used in the paper on x86 executable code.

Hardware Side: We use our framework to define contracts for a comprehensive set of recent hardware mechanisms for secure speculation: disabling speculation, delaying speculative load operations [3], and speculative taint tracking [4], [5].

To this end, we formalize each mechanism in the context of a variant of the simple speculative out-of-order processor from [14] and we prove that it satisfies specific contracts against an adversary that observes caches, predictors, and (part of) the reorder buffer during execution. We show that the contracts we define form a lattice, and we use this to give, for the first time, a rigorous comparison of the security guarantees offered by different secure speculation mechanisms.

Our analysis highlights that the analyzed mechanisms [3], [4], [5] prevent leaks of speculatively accessed data, and confirms the results of [5]. For software, this means that “sandboxing” is supported out-of-the-box, in the sense that programs only need to place appropriate bounds checks, but no speculation barriers.

Our analysis also shows that the mechanisms offer no support for “constant-time programming”. This means that programs that are constant-time in the traditional sense [18] still require additional checks [12], [14] or insertion of speculation barriers [19], even if hardware mechanisms for secure speculation are deployed.

Summary of contributions: We propose a novel framework for expressing security contracts between hardware and software. Our framework is expressive enough to (1) characterize the security guarantees provided by recent proposals for secure

speculation, and (2) provide program-level properties formalizing how to leverage these hardware guarantees to achieve global, end-to-end security for different scenarios. From a theoretical perspective, we provide the first characterization of security for a comprehensive class of hardware mechanisms for secure speculation. From a practical perspective, we show how to automate checks for programs to run securely on top of these mechanisms.

II. ISA LANGUAGE, SEMANTICS, AND ADVERSARIES

We introduce the foundations for specifying hardware-software contracts: an ISA language (§II-A), its architectural semantics (§II-B), a general notion of hardware semantics (§II-C), and an adversary model capturing which aspects of the microarchitecture are observable via side channels (§II-D).

A. ISA language

For modeling the ISA we rely on μ ASM, a simple assembly language from [12] with the following syntax:

Basic Types

$$\begin{aligned} (\text{Registers}) \quad x &\in \text{Regs} \\ (\text{Values}) \quad n, \ell &\in \text{Vals} = \mathbb{N} \cup \{\perp\} \end{aligned}$$

Syntax

$$\begin{aligned} (\text{Expressions}) \quad e &:= n \mid x \mid \ominus e \mid e_1 \otimes e_2 \mid \mathbf{ite}(e_1, e_2, e_3) \\ (\text{Instructions}) \quad i &:= \mathbf{skip} \mid x \leftarrow e \mid \mathbf{load} \ x, e \mid \mathbf{store} \ x, e \\ &\quad \mid \mathbf{jmp} \ e \mid \mathbf{beqz} \ x, \ell \mid \mathbf{spbarr} \\ (\text{Programs}) \quad p &:= i \mid p_1; p_2 \end{aligned}$$

- μ ASM expressions are built from a set of register identifiers *Regs*, which contains a designated element **pc** representing the program counter, and a set of values *Vals*, which consists of the natural numbers and \perp .
- μ ASM instructions include assignments, load and store instructions, indirect jumps, branching instructions, and a speculation barrier **spbarr**.
- μ ASM programs are sequences of instructions.

B. Architectural semantics \rightarrow

The architectural semantics models the execution of μ ASM programs at the architectural level. It is defined in terms of *architectural states* (arch. states for short) $\sigma = \langle m, a \rangle$ consisting of a *memory* m and a *register assignment* a . Memories m map memory addresses, represented by natural numbers, to values in *Vals*. Register assignments a map register identifiers to values in *Vals*. We signal program termination by assigning the special value \perp to the program counter **pc**.

The architectural semantics is a deterministic binary relation $\sigma \rightarrow \sigma'$ mapping an arch. state σ to its successor σ' . We present the arch. semantics in Appendix A. A *run* is a finite sequence of states $\sigma_0, \dots, \sigma_n$ with $\sigma_0 \rightarrow \dots \rightarrow \sigma_n$ such that σ_0 is initial (that is, all registers including **pc** have value 0) and σ_n is final (that is, $\sigma_n(\mathbf{pc}) = \perp$).

C. Hardware semantics \Rightarrow

A hardware semantics models the execution of μ ASM programs at the microarchitectural level. Here we describe

²In the terminology of [11], sandboxing aims to block disclosure gadgets.

a general notion of hardware semantics with the key aspects necessary for explaining hardware-software contracts; we provided multiple, concrete hardware semantics modeling different processors and countermeasures in §V–VI.

Hardware semantics are defined in terms of *hardware states* $\langle \sigma, \mu \rangle$ consisting of an arch. state σ (as before) and a *microarchitectural state* (μ arch. state for short) μ , which models the state of components like predictors, caches, and reorder buffer.

A hardware semantics is a deterministic relation \Rightarrow mapping hardware states $\langle \sigma, \mu \rangle$ to their successors $\langle \sigma', \mu' \rangle$. A *hardware run* is a sequence $\langle \sigma_0, \mu_0 \rangle \Rightarrow \dots \Rightarrow \langle \sigma_n, \mu_n \rangle$ such that $\langle \sigma_0, \mu_0 \rangle$ is initial and $\langle \sigma_n, \mu_n \rangle$ is final. For this, we assume that there is a fixed, initial μ arch. state μ_0 , where, for instance, the reorder buffer is empty and all caches have been invalidated.

D. Adversary model

We consider adversaries that can observe parts of the μ arch. state during execution. We model hardware observations as projections to parts of the μ arch. state. For instance, a cache-adversary can be modeled as a function \mathcal{A} projecting μ to its cache component. In the paper, we consider an adversary \mathcal{A} that has access to the state of caches, predictors, and (part of) the reorder buffer; we formalize \mathcal{A} in Section V-C.

Given a program p , $\llbracket p \rrbracket(\sigma)$ denotes the trace $\mathcal{A}(\mu_0) \cdot \dots \cdot \mathcal{A}(\mu_n)$ of hardware observations produced in the run $\langle \sigma, \mu_0 \rangle \Rightarrow \dots \Rightarrow \langle \sigma_n, \mu_n \rangle$. We refer to $\llbracket p \rrbracket$ as the *hardware trace semantics* (hardware semantics for short) of program p .

III. HARDWARE-SOFTWARE CONTRACTS

The purpose of a contract is to split the responsibilities for preventing side-channels between software and hardware.

We first formalize the general notion of contracts and we specify when a hardware platform satisfies a contract. Then we present several fundamental contracts for secure speculation.

A. Formalizing contracts

A *contract* is a labeled, deterministic semantics \rightarrow for the ISA. Given a program p and an initial arch. state σ_0 , the labels on the transitions of the corresponding run $\sigma_0 \xrightarrow{\ell_1} \sigma_1 \xrightarrow{\ell_2} \dots \xrightarrow{\ell_n} \sigma_n$ define the *trace* $\llbracket p \rrbracket(\sigma_0) = \ell_1 \ell_2 \dots \ell_n$.

The traces of a contract $\llbracket p \rrbracket$ capture which arch. states are guaranteed to be indistinguishable to an attacker on a hardware satisfying the contract, which is formalized below.

Definition 1 ($\llbracket \cdot \rrbracket \vdash \llbracket \cdot \rrbracket$). A hardware semantics $\llbracket \cdot \rrbracket$ *satisfies a contract* $\llbracket \cdot \rrbracket$ if, for all programs p and all initial arch. states σ, σ' , if $\llbracket p \rrbracket(\sigma) = \llbracket p \rrbracket(\sigma')$, then $\llbracket p \rrbracket(\sigma) = \llbracket p \rrbracket(\sigma')$.

Different contracts correspond to different divisions of security obligations between software and hardware: secrets at the program level must not affect contract observations, because then they can become visible to the adversary. Hence, contracts exposing more observations correspond to hardware with weaker security guarantees, whereas contracts exposing fewer observations correspond to hardware with stronger security guarantees. A degenerate case is a contract with no observations, which is satisfied by an ideal side-channel resilient platform that securely executes every program.

B. Contracts for secure speculation

We now define four fundamental contracts that characterize the security guarantees offered by mechanisms for secure speculation. We derive our contracts as the combination of two kinds of building blocks.

1) *Building blocks for contracts*: The first building block are *observer modes*, which govern what information a contract exposes. We define them via labels on the contract semantics.

- The *constant-time* observer mode (**ct** for short) is commonly used when reasoning about side-channels in cryptosystems. It uses labels **pc** ℓ , **load** n , and **store** n to expose the value ℓ of the program counter and the locations n of load and store operations. The observer mode can be augmented with support for variable-latency instructions by additionally exposing the operands of those instructions as observations, which we forgo for simplicity.

- The *architectural* observer mode (**arch** for short) additionally exposes the *values* v that is loaded from memory locations n via the label **load** $n = v$. As registers are set to zero in the initial arch. state, **arch**-traces also determine the values of registers during execution.

The second building block are *execution modes* that characterize which paths need to be explored to collect observations. For processors with speculative execution, depending on the presence and effectiveness of hardware-level countermeasures, one needs to go beyond those covered by the arch. semantics.

- In the *sequential* execution mode (**seq** for short), programs are executed sequentially and in-order following the arch. semantics.

- In the *always-mispredict* execution mode (**spec** for short), programs are executed sequentially, but incorrect branches are also executed for a bounded number of steps before backtracking. This execution mode is based on [12] and can be used to explore the effects of speculatively executed instructions at the ISA level.

2) *Contract* $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}}$: This contract exposes the program counter and the locations of memory accesses on sequential, non-speculative paths; see Figure 2. $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}}$ is a fundamental baseline that is often implicitly assumed in practice, and that has also been formalized in [18], [20].

In Section VI-A we show that $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}}$ is satisfied by a simple in-order processor without speculation. However, modern out-of-order processors do *not* satisfy $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}}$, as shown below.

Example 1. Consider the vanilla Spectre v1 snippet from Figure 1a, compiled to μ ASM:

```

1  x ← y < size_A
2  beqz x, 1 //checking y < size_A
3  load  z, A + y //accessing A[y]
4  z ← z*64
5  load  w, B+z //accessing B[A[y]*64]

```

Consider arch. states σ and σ' that agree on the observations on trace **pc** 3 · **load** (A+y) · **load** (B+x) (and hence on the content of array **A** within bounds), but for which $\sigma(\mathbf{A} + \mathbf{y}) = 0$ and $\sigma'(\mathbf{A} + \mathbf{y}) = 1$ for some $\mathbf{y} > \text{size_A}$. On processors with speculation, an adversary with cache access can distinguish σ and σ' , as demonstrated by the Spectre attack [21].

$$\begin{array}{c}
\text{LOAD} \\
\hline
p(a(\mathbf{pc})) = \mathbf{load} \ x, e \quad \langle m, a \rangle \rightarrow \langle m', a' \rangle \\
\hline
\langle m, a \rangle \xrightarrow{\text{load } \langle e \rangle(a) \xrightarrow{\text{seq}}_{\text{ct}}} \langle m', a' \rangle \\
\\
\text{STORE} \\
\hline
p(a(\mathbf{pc})) = \mathbf{store} \ x, e \quad n = \langle e \rangle(a) \quad \langle m, a \rangle \rightarrow \langle m', a' \rangle \\
\hline
\langle m, a \rangle \xrightarrow{\text{store } n \xrightarrow{\text{seq}}_{\text{ct}}} \langle m', a' \rangle \\
\\
\text{BEQZ-SAT} \\
\hline
p(a(\mathbf{pc})) = \mathbf{beqz} \ x, \ell \quad \langle m, a \rangle \rightarrow \langle m', a' \rangle \\
\hline
\langle m, a \rangle \xrightarrow{\text{pc } a'(\mathbf{pc}) \xrightarrow{\text{seq}}_{\text{ct}}} \langle m', a' \rangle
\end{array}$$

Fig. 2: $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}}$ contract for a program p - selected rules (here $\langle e \rangle(a)$ is the result of expression e given assignment a). The contract is obtained by augmenting the arch. semantics with observations $\text{load } n$, $\text{store } n$, and $\text{pc } \ell$ exposing the addresses of loads, stores, and the program counter, respectively.

Perhaps surprisingly, processors deploying recent proposals for secure speculation still violate $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}}$, see § VI.

3) *Contract* $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{spec}}$: This contract additionally exposes the program counter and the locations of all memory accesses on speculatively executed paths. It is based on the speculative semantics from [12] and formalized in Figure 3.

In Section VI, we show that speculative out-of-order processors (with and without mechanisms for secure speculation) satisfy $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{spec}}$.

Consider again Example 1: by exposing observations on mispredicted paths, $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{spec}}$ makes the states σ, σ' distinguishable at the contract level, effectively delegating the responsibility of ensuring that $\sigma(\mathbf{A} + \mathbf{y})$ and $\sigma'(\mathbf{A} + \mathbf{y})$ do not carry secret information to the software side.

4) *Contract* $\llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}}$: This contract exposes the program counter, the location of all loads and stores, and the values of all data loaded from memory on standard, i.e. non-speculative, program paths. The contract is obtained by modifying the LOAD rule from Figure 2 as follows:

$$\begin{array}{c}
\text{LOAD} \\
\hline
p(a(\mathbf{pc})) = \mathbf{load} \ x, e \quad \langle m, a \rangle \rightarrow \langle m', a' \rangle \\
\hline
\langle m, a \rangle \xrightarrow{\text{load } \langle e \rangle(a) = m(\langle e \rangle(a)) \xrightarrow{\text{seq}}_{\text{arch}}} \langle m', a' \rangle
\end{array}$$

As we assume that register values are zeroed in the initial state, the $\llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}}$ trace effectively exposes the contents of registers during execution. While this does not seem to guarantee any kind of security, $\llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}}$ *does* guarantee the confidentiality of data that is *only transiently* loaded, thus effectively preventing speculative disclosure gadgets. In that sense, the contract $\llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}}$ is a simple and clean formulation of the idea behind *transient noninterference* [5], making it comparable to the guarantees offered by other contracts, and providing an actionable interface to software.

5) *Special contracts*: We informally present a number of contracts that illustrate our framework’s expressiveness:

- $\llbracket \cdot \rrbracket_{\perp}$ is the contract that does not expose any observa-

tions and corresponds to a hypothetical side-channel resilient processor that can securely execute every program.

- $\llbracket \cdot \rrbracket_{\text{ct-pc}}^{\text{seq-spec}}$ exposes program counter and addresses of loads during sequential execution, and only the program counter during speculative execution. That is, it may intuitively be understood as $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}} + \llbracket \cdot \rrbracket_{\text{pc}}^{\text{spec}}$.
- $\llbracket \cdot \rrbracket_{\text{arch}}^{\text{spec}}$ exposes the values of data loaded from memory also during speculatively executed instructions. It corresponds to a processor that does not offer any confidentiality guarantees for any accessed data.
- $\llbracket \cdot \rrbracket_{\perp}$ exposes all arch. state. It could correspond to a processor vulnerable to all Meltdown-type attacks (see §VII).

C. A lattice of contracts

Finally, we compare contracts in terms of the security guarantees they offer to software. Intuitively, a contract is stronger than another, if it guarantees to leak less information to a microarchitectural adversary.

Definition 2 ($\llbracket \cdot \rrbracket_1 \sqsupseteq \llbracket \cdot \rrbracket_2$). A contract $\llbracket \cdot \rrbracket_1$ is *stronger* than a contract $\llbracket \cdot \rrbracket_2$ if $\llbracket p \rrbracket_2(\sigma) = \llbracket p \rrbracket_2(\sigma') \Rightarrow \llbracket p \rrbracket_1(\sigma) = \llbracket p \rrbracket_1(\sigma')$ for all programs p and all initial arch. states σ, σ' .

Equivalently, $\llbracket \cdot \rrbracket_1 \sqsupseteq \llbracket \cdot \rrbracket_2$ holds whenever two arch. states that can be distinguished by $\llbracket \cdot \rrbracket_1$ ’s traces can also be distinguished by $\llbracket \cdot \rrbracket_2$ ’s traces.

Note that if $\llbracket \cdot \rrbracket_1$ exposes only a subset of the labels of $\llbracket \cdot \rrbracket_2$, then $\llbracket \cdot \rrbracket_1$ is stronger than $\llbracket \cdot \rrbracket_2$ according to Definition 2. For example, the instructions explored by spec are also explored by seq , and the observations of ct are contained in the observations of arch . This enables us to arrange all contracts defined in §III-B in the lattice [22] shown in Figure 4.

Finally, as expected, a hardware platform that satisfies a contract $\llbracket \cdot \rrbracket_1$ also satisfies all weaker contracts $\llbracket \cdot \rrbracket_2$.

Proposition 1. *If $\llbracket \cdot \rrbracket \vdash \llbracket \cdot \rrbracket_1$ and $\llbracket \cdot \rrbracket_2 \sqsubseteq \llbracket \cdot \rrbracket_1$, then $\llbracket \cdot \rrbracket \vdash \llbracket \cdot \rrbracket_2$.*

This implies that processors with stronger contracts $\llbracket \cdot \rrbracket_1$ are backward-compatible in the sense that they can securely execute any side-channel resilient legacy code that was already secure under weaker contracts $\llbracket \cdot \rrbracket_2$.

IV. PROGRAMMING AGAINST CONTRACTS

Contracts are the basis for secure programming. Here, we consider two scenarios that are both instances of secure programming: In the first, which we call “constant-time programming”, the goal is to ensure that a benign program does not leak confidential data to an adversary while computing on this data. In the second, which we call “sandboxing”, the goal is to prevent a potentially malicious program from accessing confidential data.

A. Secure programming

We begin by framing secure programming as an information-flow property. To distinguish confidential from public data, we rely on a policy $\pi: \text{Vals} \rightarrow \{L, H\}$ that labels memory locations as high (H) or low (L), encoding whether locations store confidential data or not. Two arch. states σ, σ'

$$\begin{array}{c}
\text{STEP} \\
\frac{p(\sigma(\mathbf{pc})) \neq \mathbf{beqz} \ x, \ell \quad \sigma \xrightarrow{\tau_{\text{ct}}^{\text{seq}}} \sigma'}{\langle \sigma, \omega + 1 \rangle \cdot s \xrightarrow{\tau_{\text{ct}}^{\text{spec}}} \langle \sigma', \omega \rangle \cdot s} \\
\\
\text{ROLLBACK} \\
\frac{s = \langle \sigma', \omega' \rangle \cdot s'}{\langle \sigma, 0 \rangle \cdot s \xrightarrow{\text{pc } \sigma'(\mathbf{pc}) \text{ spec}_{\text{ct}}} s} \\
\\
\text{BARRIER} \\
\frac{p(\sigma(\mathbf{pc})) = \mathbf{spbarr} \quad \sigma \xrightarrow{\tau_{\text{ct}}^{\text{seq}}} \sigma'}{\langle \sigma, \omega + 1 \rangle \cdot s \xrightarrow{\tau_{\text{ct}}^{\text{spec}}} \langle \sigma', 0 \rangle \cdot s} \\
\\
\text{BRANCH} \\
\frac{p(\sigma(\mathbf{pc})) = \mathbf{beqz} \ x, \ell \quad \ell_{\text{correct}} = \begin{cases} \ell & \text{if } \sigma(x) = 0 \\ \sigma(\mathbf{pc}) + 1 & \text{otherwise} \end{cases} \quad \ell_{\text{mispred}} \in \{\ell, \sigma(\mathbf{pc}) + 1\} \setminus \ell_{\text{correct}} \quad \omega_{\text{mispred}} = \begin{cases} \omega & \text{if } \omega = \infty \\ \omega & \text{otherwise} \end{cases}}{\langle \sigma, \omega + 1 \rangle \cdot s \xrightarrow{\text{pc } \ell_{\text{mispred}} \text{ spec}_{\text{ct}}} \langle \sigma[\mathbf{pc} \mapsto \ell_{\text{mispred}}], \omega_{\text{mispred}} \rangle \cdot \langle \sigma[\mathbf{pc} \mapsto \ell_{\text{correct}}], \omega \rangle \cdot s}
\end{array}$$

Fig. 3: Definition of $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{spec}}$ contract. Configurations are stacks of $\langle \sigma, \omega \rangle$, where $\omega \in \mathbb{N} \cup \{\infty\}$ is the speculative window denoting how many instructions are left to be executed. (initial arch. states σ are treated as $\langle \sigma, \infty \rangle$). At each computation step, the ω at the top of the stack is reduced by 1 (rules STEP and BRANCH). When executing a branch instruction (rule BRANCH), the state $\langle \sigma[\mathbf{pc} \mapsto \ell_{\text{mispred}}], \omega_{\text{mispred}} \rangle$ is pushed on top of the stack, thereby allowing the exploration of the mispredicted branch for ω_{mispred} steps. The correct branch $\langle \sigma[\mathbf{pc} \mapsto \ell_{\text{correct}}], \omega \rangle$ is also recorded on the stack; allowing to later roll back speculatively executed statements. When the ω at the top of the stack reaches 0, we pop it (i.e., we backtrack and discard the changes) and we continue the computation. Speculation barriers trigger a roll back by setting ω to 0 (rule BARRIER).




Fig. 4: Lattice of contracts. An edge from $\llbracket \cdot \rrbracket_1$ to $\llbracket \cdot \rrbracket_2$ means that $\llbracket \cdot \rrbracket_1 \sqsubseteq \llbracket \cdot \rrbracket_2$. The $\llbracket \cdot \rrbracket_{\top}$ contract is the one without observations, and $\llbracket \cdot \rrbracket_{\perp}$ one exposing all the architectural state.

are *low-equivalent*, written $\sigma \simeq_L \sigma'$, iff the values of all low memory locations are the same.

Definition 3 ($p \vdash NI(\pi, \llbracket \cdot \rrbracket)$). Program p is *non-interferent* w.r.t. contract $\llbracket \cdot \rrbracket$ and policy π if for all initial arch. states σ, σ' : $\sigma \simeq_L \sigma' \Rightarrow \llbracket p \rrbracket(\sigma) = \llbracket p \rrbracket(\sigma')$.

That is, a program is non-interferent w.r.t. a contract and a policy, if low-equivalent arch. states are indistinguishable under the contract, i.e., no information about high memory locations leaks into the contract's traces.

Similarly to Def. 3, one can define a notion of non-interference w.r.t. a hardware semantics $\{\cdot\}, \{\cdot\}$, written $p \vdash NI(\pi, \{\cdot\}, \{\cdot\})$, where information about high memory locations cannot flow into hardware observations.

The following proposition, capturing leakage at the hardware level, follows by composition of Definitions 1 and 3:

Proposition 2. *If $p \vdash NI(\pi, \llbracket \cdot \rrbracket)$ and $\{\cdot\} \vdash \llbracket \cdot \rrbracket$, then $p \vdash NI(\pi, \{\cdot\})$.*

B. Sandboxing

The goal of sandboxing is to enable the safe execution of untrusted, potentially malicious code. This is achieved by ensuring that the untrusted code is confined to a set of tightly controlled resources. Here we focus on one important aspect: preventing code from reading outside of its own subset of the address space. To achieve this, just-in-time compilers enforce

access-control policies by inserting checks to ensure that all memory accesses happen within the sandbox's bounds.

We describe sandboxes using policies π , where memory outside of the sandbox is declared high. To account for programs that may escape the sandbox by exploiting speculation across access-control checks, we make the following distinction:

- Traditional sandboxing approaches [23], [24] check/enforce *vanilla sandboxing*: A program p is *vanilla-sandboxed* w.r.t. π if p never accesses high memory locations when executing under the arch. semantics \rightarrow . In our framework, being vanilla-sandboxed is equivalent to $p \vdash NI(\pi, \llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}})$, i.e., being non-interferent w.r.t. $\llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}}$. This follows from $\llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}}$ exposing the value of accessed high memory locations.

- To faithfully reason about sandboxing on out-of-order and speculative processors, one needs to go beyond vanilla sandboxing and make sure that the program does not leak any information that is outside of its sandbox through a covert channel. We say that program is *generally-sandboxed* w.r.t. contract $\llbracket \cdot \rrbracket$, if it is vanilla-sandboxed and in addition non-interferent w.r.t. $\llbracket \cdot \rrbracket$, i.e., $p \vdash NI(\pi, \llbracket \cdot \rrbracket)$. General sandboxing together with Proposition 2 guarantees that no data outside of the sandbox affects what a μ arch. adversary (including the sandboxed program p itself, via probing) can observe on any platform satisfying $\llbracket \cdot \rrbracket$.

Def. 4 enables to bridge the gap between vanilla sandboxing and general sandboxing for a given program.

Definition 4. Program p satisfies *weak speculative non-interference* (wSNI) with respect to $\llbracket \cdot \rrbracket$ if for all initial arch. states σ, σ' : $\llbracket p \rrbracket_{\text{arch}}^{\text{seq}}(\sigma) = \llbracket p \rrbracket_{\text{arch}}^{\text{seq}}(\sigma') \Rightarrow \llbracket p \rrbracket(\sigma) = \llbracket p \rrbracket(\sigma')$.

Weak speculative non-interference is a variant of *speculative non-interference*, the security checked by Spectector [12].

Proposition 3 shows how wSNI bridges the gap between vanilla and general sandboxing.

Proposition 3. *If program p is vanilla-sandboxed w.r.t. π and wSNI w.r.t. $\llbracket \cdot \rrbracket$, then p is generally-sandboxed w.r.t. π and $\llbracket \cdot \rrbracket$.*

Hence, to check whether a program p is generally-

sandboxed w.r.t. $\llbracket \cdot \rrbracket$ and π one can: (1) check/enforce that p is vanilla-sandboxed w.r.t. π , and (2) verify whether p is wSNI.

C. Constant-time programming

Constant-time programming is a coding discipline for the implementation of code like cryptographic algorithms that needs to compute over secret data without leaks. Code without (1) secret-dependent control flow, (2) secret-dependent memory accesses, and (3) secret-dependent inputs to variable-latency instructions is traditionally understood as “constant time”. As discussed before this corresponds to $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}}$, which exposes control flow and memory accesses.³

Again, considering only $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}}$ is insufficient to reason about constant-time on modern processors. For this, we make the following distinction:

- Existing constant-time approaches (type systems [25], static analyses [18], [26], and techniques for secure compilation [27], [28]) check/enforce *vanilla-constant-time*. In our framework, a program p is *vanilla-constant-time* w.r.t. π if $p \vdash NI(\pi, \llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}})$, i.e., p non-interferent w.r.t. $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}}$.
- More generally, a program p is *generally-constant-time* w.r.t. contract $\llbracket \cdot \rrbracket$ iff $p \vdash NI(\pi, \llbracket \cdot \rrbracket)$, i.e., constant timeness coincides with non-interference w.r.t. a contract.

One possibility for checking general-constant-timeness is devising dedicated tools [14]. Alternatively, one can reuse vanilla-constant-time tools [18], [26] and then bridge the gap between vanilla and general-constant-time. To bridge this gap, one can rely on the following generalization of *speculative non-interference* from [12]:

Definition 5 (Speculative non-interference [12]). Program p is *speculatively non-interferent* (SNI) w.r.t. policy π and contract $\llbracket \cdot \rrbracket$ if for all initial arch. states σ, σ' :

$$\sigma \simeq_L \sigma' \wedge \llbracket p \rrbracket_{\text{ct}}^{\text{seq}}(\sigma) = \llbracket p \rrbracket_{\text{ct}}^{\text{seq}}(\sigma') \Rightarrow \llbracket p \rrbracket(\sigma) = \llbracket p \rrbracket(\sigma').$$

Proposition 4 shows how SNI bridges the gap between vanilla and general constant-time.

Proposition 4. *If program p is vanilla-constant-time w.r.t. π and SNI w.r.t. π and $\llbracket \cdot \rrbracket$, then p is generally-constant-time w.r.t. π and $\llbracket \cdot \rrbracket$.*

Thus to check whether a program p is generally-constant-time w.r.t. $\llbracket \cdot \rrbracket$ and π one can (1) check vanilla-constant-time, and (2) verify whether p is SNI w.r.t. $\llbracket \cdot \rrbracket$ and π .

Observe, however, that not all contracts are useful for general-constant-time. Remarkably, the $\llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}}$ contract, which naturally corresponds to the guarantees provided by state-of-the-art HW-level countermeasures like STT [5] and NDA [4] is inherently inadequate for constant-time programming: A program that is non-interferent w.r.t. $\llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}}$ may not access any secret data. However, accessing and computing on secret data is the whole point of constant-time programming.

D. Experiments

In this section, we illustrate how our framework can be used to support secure programming, for both the sandboxing and constant-time scenarios, w.r.t. the contracts from §III.

³As discussed earlier, we forgo variable-latency instructions in μASM .

1 <code>if (y < size_A)</code>	1 <code>x = A[y];</code>
2 <code>x = A[y];</code>	2 <code>if (y < size_A)</code>
3 <code>if (x)</code>	3 <code>if (x)</code>
4 <code>temp &= B[0];</code>	4 <code>temp &= B[0];</code>

(a) Program P_1^f
(b) Program P_2^f

Fig. 5: Variants of Spectre v1 that leak information through the control-flow statement in line 3.

Tooling: To automate our analysis we adapted Spectector [12], which can already check SNI for the $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{spec}}$ contract, to support checking SNI and wSNI w.r.t. all the contracts from §III, i.e., $\llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}}, \llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}}, \llbracket \cdot \rrbracket_{\text{ct}}^{\text{spec}}, \llbracket \cdot \rrbracket_{\text{ct-pc}}^{\text{seq-spec}}$.

Propositions 3–4 present a clear path to check (general) sandboxing/constant-time: (1) use existing tools to verify vanilla sandboxing/constant-time, and (2) verify wSNI/SNI using Spectector.

Experimental setup: We analyze 4 different programs:

- P_1 and P_2 are the Spectre v1 snippet from Figure 1a and its variant from Figure 1b, respectively.
- P_1^f and P_2^f are modifications of P_1 and P_2 that leak information through control-flow statements. The programs are shown in Figure 5.

We compile each program with Clang at `-O2` optimization level. We also compile each program with a countermeasure that automatically injects `lfence` speculation barriers after each branch instruction.⁴ We denote by P^f the program P with `lfences`.

As a result, we have eight small x86 programs. that we analyze with our enhanced version of Spectector.

Sandboxing: We analyze programs P_1, P_1^f, P_1^f, P_1^f w.r.t. the policy π that declares the contents of `A[i]` as *low* for all i that are within the array bounds, and as *high* otherwise.

Our goal is to determine whether these programs satisfy the general-sandboxing property w.r.t. the contracts in §III. We remark that all variants of P_1 are vanilla-sandboxed w.r.t. π : they never access out-of-bound locations under the arch. semantics \rightarrow thanks to the bounds check.

Figure 6 summarizes our findings, which we discuss below:

- For $\llbracket \cdot \rrbracket \in \{\llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}}, \llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}}\}$, the fact that $\llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}}$ and $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}}$ are stronger than $\llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}}$ (see §III-C) directly implies wSNI w.r.t. these contracts for *any* program (denoted by “Y, \sqsupseteq ” in the table). Therefore, programs P_1, P_1^f, P_1^f, P_1^f all satisfy general-sandboxing (see Proposition 3) without further analysis.
 - For $\llbracket \cdot \rrbracket \in \{\llbracket \cdot \rrbracket_{\text{ct}}^{\text{spec}}, \llbracket \cdot \rrbracket_{\text{ct-pc}}^{\text{seq-spec}}\}$, we check whether wSNI holds using Spectector. Table entries “Y, wSNI” denote a successful check, which implies (via Proposition 3) that the program is generally-sandboxed w.r.t. $\llbracket \cdot \rrbracket$. In several cases, denoted by “N”, the wSNI check fails. While this is not generally the case, the counterexamples to wSNI show that the respective programs are indeed not sandboxed w.r.t. $\llbracket \cdot \rrbracket$.
- Program P_1 fails the wSNI check w.r.t. $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{spec}}$, due to the speculative secret-dependent load (line 3 in Fig. 1b), but it

⁴The countermeasure is enabled with the `-x86-speculative-load-hardening -x86-slh-lfence` flags.

	$\llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}}$	$\llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}}$	$\llbracket \cdot \rrbracket_{\text{ct}}^{\text{spec}}$	$\llbracket \cdot \rrbracket_{\text{ct-pc}}^{\text{seq-spec}}$
P_1	Y, \sqsupseteq	Y, \sqsupseteq	N	Y, wSNI
P_1^f	Y, \sqsupseteq	Y, \sqsupseteq	Y, wSNI	Y, wSNI
P_1'	Y, \sqsupseteq	Y, \sqsupseteq	N	N
$P_1^{f'}$	Y, \sqsupseteq	Y, \sqsupseteq	Y, wSNI	Y, wSNI

Fig. 6: Sandboxing analysis w.r.t. different contracts.

	$\llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}}$	$\llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}}$	$\llbracket \cdot \rrbracket_{\text{ct}}^{\text{spec}}$	$\llbracket \cdot \rrbracket_{\text{ct-pc}}^{\text{seq-spec}}$
P_2	Y, \sqsupseteq	N	N	Y, SNI
P_2^f	Y, \sqsupseteq	N	Y, SNI	Y, SNI
P_2'	Y, \sqsupseteq	N	N	N
$P_2^{f'}$	Y, \sqsupseteq	N	Y, SNI	Y, SNI

Fig. 7: Constant-time analysis results w.r.t. different contracts.

satisfies wSNI w.r.t. the stronger contract $\llbracket \cdot \rrbracket_{\text{ct-pc}}^{\text{seq-spec}}$ that ensures confidentiality of secret-dependent speculative loads.

- In contrast, program P_1' violates wSNI due to the speculative branch statement (line 3 in Fig. 5) w.r.t. $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{spec}}$ and $\llbracket \cdot \rrbracket_{\text{ct-pc}}^{\text{seq-spec}}$.
- Finally, programs P_1^f and $P_1^{f'}$, where `lfences` are inserted after the branch, satisfy wSNI w.r.t. $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{spec}}$ and $\llbracket \cdot \rrbracket_{\text{ct-pc}}^{\text{seq-spec}}$.

Constant-time: We analyze programs $P_2, P_2', P_2^f, P_2^{f'}$ w.r.t. the same policy π as before.

This time, our goal is to determine whether these programs are constant-time w.r.t. the contracts in §III. We remark that $P_2, P_2', P_2^f, P_2^{f'}$ are vanilla-constant-time w.r.t. π , while none of these programs is vanilla-sandboxed w.r.t. π .

Figure 7 summarizes our findings, which we discuss below:

- For $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}}$, all programs are constant-time w.r.t. $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}}$ as they are vanilla-constant-time (denoted by “Y, \sqsupseteq ” in the table).
- For $\llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}}$, constant-time is violated for all programs, with and without `lfence`, due to the non-speculative load of a secret into the architectural state.
- For $\llbracket \cdot \rrbracket \in \{\llbracket \cdot \rrbracket_{\text{ct}}^{\text{spec}}, \llbracket \cdot \rrbracket_{\text{ct-pc}}^{\text{seq-spec}}\}$, Table entries “Y, SNI” denote a successful check using Spectector, which implies (via Proposition 4) that the program is constant-time w.r.t. $\llbracket \cdot \rrbracket$. Again, while this is not true in general, the counterexamples to SNI for these particular programs turn out to be proofs that the programs are not constant-time w.r.t. $\llbracket \cdot \rrbracket$.

Program P_2 violates SNI w.r.t. $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{spec}}$ but satisfies it under the stronger contract $\llbracket \cdot \rrbracket_{\text{ct-pc}}^{\text{seq-spec}}$ that does not expose the address of the speculative load (line 3 in Figure 1b). In contrast, P_2' violates SNI against both contracts. Finally, the programs with fences (P_2^f and $P_2^{f'}$) satisfy SNI w.r.t. $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{spec}}$ and $\llbracket \cdot \rrbracket_{\text{ct-pc}}^{\text{seq-spec}}$.

V. MODELING MICROARCHITECTURE AND ADVERSARIES

This section presents a hardware semantics for μ ASM programs. The semantics is based on the semantics from [14], [19] and it models the execution of μ ASM programs by a simple out-of-order processor with a unified cache for data and instructions and a branch predictor for speculative execution over branch instructions. The purpose of this semantics is to allow us to model and reason about hardware-level Spectre

countermeasures; see §VI. To this end, it strives to achieve the following design goals: (1) To faithfully capture the key features of speculative and out-of-order execution, while (2) keeping it simple, and (3) supporting large classes of microarchitectural features like caches and branch predictors. The latter aspect allows us to focus on hardware-level countermeasures in the context of an arbitrary caching algorithms and branch-prediction strategies.

We start by formalizing hardware configurations (Section V-A) that extend arch. states with the state of the μ arch. components, i.e., cache, reorder buffer, and branch predictor.

Next, we formalize the semantics of the pipeline steps (Section V-B). This semantics describes how instructions are fetched, executed, and retired under our semantics as well as how hardware configurations are updated during the execution.

A. Hardware configurations

Each *hardware configuration* $\langle \sigma, \mu \rangle$ consists of its arch. state σ , recording the memory and register assignments, and of its μ arch. state μ , which we formalize next.

The μ arch. state consists of a reorder buffer, which stores the state of in-flight instructions, a cache, a branch predictor, and a scheduler, which orchestrates the pipeline during the computation. Note that, in our model, cache states track which memory blocks are stored in the cache (i.e., they store metadata) but they do *not* store the data itself. While we fix the behavior of the reorder buffer in §V-A1, our semantics is parametric in the models of caches, branch predictors, and the pipeline scheduler; see §V-A2. Theorem statements in §VI (except where explicitly stated) hold for *all* possible choices of cache, predictor, and scheduler in our model.

1) *Reorder buffers:* Reorder buffers store the state of in-flight, i.e., not yet retired, instructions. Initially instructions are *unresolved*, e.g., a load `load x, y + z` that has not yet been performed or an assignment `z ← 2 + k` whose right-hand side has not yet been evaluated. Executing an unresolved instruction can transform it into a *resolved* instruction, where all expressions are replaced with their values. Additionally, to model speculative control flow, reorder buffer entries may be *tagged* with the address of a branch instruction ℓ . We write `pc ← v@ℓ`, whenever the assignment of v to the `pc` is the result of a call to the branch predictor when fetching the branch at address ℓ . Instructions are *untagged*, written $i@ε$, if they are not the result of a prediction.

We model reorder buffers as sequences of *commands* of length at most `w` denoting the buffer’s maximal length:

$$\begin{aligned}
 (\text{Tags}) \quad T &:= \varepsilon \mid \ell \\
 (\text{Commands}) \quad c &:= i@T \\
 (\text{Reorder buffers}) \quad \text{buf} &:= \varepsilon \mid c \cdot \text{buf}
 \end{aligned}$$

A reorder buffer captures the state of execution of in-flight instructions. Consider the buffer $\text{buf} := k \leftarrow 25@ε \cdot \text{load } x, y + z@ε \cdot z \leftarrow 2 + k@ε$. It records that there are three in-flight instructions: one of them ($k \leftarrow 25@ε$) has been resolved and is ready to be retired, while the remaining two are still unresolved. Executing the third command would result in the new buffer $\text{buf}' := k \leftarrow 25@ε \cdot \text{load } x, y + z@ε \cdot z \leftarrow 27@ε$.

Given a buffer buf , its data-independent projection $buf\downarrow$ is obtained by replacing all resolved (respectively unresolved) expressions in instructions with R (respectively UR). For instance, the data-independent projection of the buffer buf from above is $k \leftarrow R@\varepsilon \cdot \mathbf{load} \ x, UR@\varepsilon \cdot z \leftarrow UR@\varepsilon$.

2) *Caches, Branch predictors, and Schedulers*: Rather than providing a fixed model for caches, branch predictors and schedulers, our semantics is parametric in such components. To this end, we only fix the interface to these components, which is given in Figure 8, constraining how the semantics may interact with these components. Each of these components is defined by a set of states, an initial state, and uninterpreted functions modeling their relevant behavior:

- Caches are equipped with a function $access(\ell, cs) \in \{\text{Hit}, \text{Miss}\}$ that captures whether accessing memory address ℓ in cache state cs results in a cache hit (Hit) or miss (Miss), and a function $update(\ell, cs) = cs'$ that updates the state of the cache based on the access to address ℓ . We stress that cache states cs track only the memory addresses of the blocks in the cache, *not* the blocks themselves.

- Branch predictors are equipped with a function $update(bp, \ell, b)$ that updates the state bp of the branch predictor by recording that the branch at program counter ℓ has been resolved to value b , and $predict(bp, \ell)$ that, given a predictor state bp , predicts the outcome of the branch at address ℓ .

- Schedulers determine which pipeline stages to activate next. Following [14], [19], we model this choice using three types of directives: (a) **fetch** is used to fetch and decode the next instruction pointed by the program counter register \mathbf{pc} , (b) **execute** i is used to execute the i -th command in the reorder buffer buf , and (c) **retire** is used to retire (i.e., apply the changes to the memory and register file) the first command in the buffer. Schedulers are equipped with a $next(sc)$ function that produces the next directive given the scheduler's state sc , and a $update(sc, buf)$ function that updates the scheduler's state based on the state of the reorder buffer.

3) *Microarchitectural states*: A $\mu\text{arch. state } \mu$ is a 4-tuple $\langle buf, cs, bp, sc \rangle$ where buf is a reorder buffer, cs is the state of the unified cache (for data and instructions), bp is the branch predictor state, and sc is the scheduler state.

A $\mu\text{arch. state } \mu$ is *initial* if $buf = \varepsilon$ and the $\mu\text{arch. components}$ are in their initial states. Similarly, μ is *final* if $buf = \varepsilon$. Hence, a hardware configuration $\langle \sigma, \mu \rangle$ is initial (respectively final) if σ and μ are so.

For simplicity, we write $\langle m, a, buf, cs, bp, sc \rangle$ to represent the hardware configuration $\langle \langle m, a \rangle, \langle buf, cs, bp, sc \rangle \rangle$.

B. Hardware semantics

We formalize the hardware semantics of a μASM program p using a binary relation $\Rightarrow \subseteq \text{HwStates} \times \text{HwStates}$ that maps hardware states to their successors:

$$\begin{array}{c} \text{STEP} \\ \frac{\langle m, a, buf, cs, bp \rangle \xRightarrow{d} \langle m', a', buf', cs', bp' \rangle}{d = next(sc) \quad sc' = update(sc, buf'\downarrow)} \\ \langle m, a, buf, cs, bp, sc \rangle \Rightarrow \langle m', a', buf', cs', bp', sc' \rangle \end{array}$$

The rule captures one execution step at the $\mu\text{arch. level}$. The scheduler is queried to determine the directive $d = next(sc)$ in-

dicating which pipeline step to execute. Next, the $\mu\text{arch. state}$ is updated by performing one step of the auxiliary relation $\langle m, a, buf, cs, bp \rangle \xRightarrow{d} \langle m', a', buf', cs', bp' \rangle$, which depends on the directive d and is formalized below. Finally, the scheduler state is updated based on the data-independent projection of the reorder buffer, i.e., $sc' = update(sc, buf'\downarrow)$. This formalizes the crucial assumption that the scheduler's decisions may depend upon the dependencies between the instructions in the reorder buffer, but not on the values computed thus far.

For each directive, i.e., **fetch**, **execute** i , and **retire**, we sketch below the rules that govern the definition of the auxiliary relations $\xRightarrow{\text{fetch}}$, $\xRightarrow{\text{execute } i}$, and $\xRightarrow{\text{retire}}$.

1) *Fetch*: Instructions are fetched in-order. Here we present selected rules modeling instruction fetch:

$$\begin{array}{c} \text{FETCH-BRANCH-HIT} \\ \frac{a' = apl(buf, a) \quad |buf| < w \quad a'(\mathbf{pc}) \neq \perp \\ p(a'(\mathbf{pc})) = \mathbf{beqz} \ x, \ell \quad \ell' = predict(bp, a'(\mathbf{pc})) \\ access(cs, a'(\mathbf{pc})) = \text{Hit} \quad update(cs, a'(\mathbf{pc})) = cs'}{\langle m, a, buf, cs, bp \rangle \xRightarrow{\text{fetch}} \langle m, a, buf \cdot \mathbf{pc} \leftarrow \ell' @ a'(\mathbf{pc}), cs', bp \rangle} \end{array}$$

$$\begin{array}{c} \text{FETCH-MISS} \\ \frac{|buf| < w \quad a' = apl(buf, a) \quad a'(\mathbf{pc}) \neq \perp \\ access(cs, a'(\mathbf{pc})) = \text{Miss} \quad update(cs, a'(\mathbf{pc})) = cs'}{\langle m, a, buf, cs, bp \rangle \xRightarrow{\text{fetch}} \langle m, a, buf, cs', bp \rangle} \end{array}$$

In these rules, and in those described later, $apl(buf, a)$ denotes the assignment a' obtained by updating a with the changes performed by the commands in buf . Concretely, $apl(buf, a)$ iteratively applies the pending changes for all commands in buf as follows: (a) Assignments $x \leftarrow e @ T$ set the value of $a'(x)$ to e if the assignment is resolved (i.e., $e \in \text{Vals}$) and to \perp otherwise (denoting unresolved values). (b) Load operations $\mathbf{load} \ x, e @ T$ set the value of $a'(x)$ to \perp (since the load operation has not been performed yet). (c) Whenever buf contains a speculation barrier $\mathbf{spbarr} @ T$, $apl(buf, a) = \lambda x \in \text{Regs. } \perp$. (d) Other instructions are ignored.

The rule **FETCH-BRANCH-HIT** models the fetch of a branch instruction $\mathbf{beqz} \ x, \ell$. Whenever the reorder buffer buf is not full ($|buf| < w$), \mathbf{pc} is defined ($a'(\mathbf{pc}) \neq \perp$), and the instruction is in the cache ($access(cs, a'(\mathbf{pc})) = \text{Hit}$), the branch predictor is queried to obtain the next program counter $\ell' = predict(bp, a'(\mathbf{pc}))$. Next, the cache and the reorder buffer states are updated. The latter is updated by appending the command $\mathbf{pc} \leftarrow \ell' @ a'(\mathbf{pc})$, which records the change to the program counter as well as the label of the branch instruction whose target was predicted. The semantics also contains rules for fetching jumps $\mathbf{jmp} \ e$, which append the command $\mathbf{pc} \leftarrow e @ \varepsilon$ to the buffer, and other instructions i , which append the commands $i @ \varepsilon \cdot \mathbf{pc} \leftarrow a'(\mathbf{pc}) + 1 @ \varepsilon$ to the buffer.

The rule **FETCH-MISS** models a cache miss when loading the next instruction. In this case, the cache is updated while the reorder buffer is not modified. A subsequent **fetch** triggered by the scheduler would result in a cache hit and a corresponding change to the reorder buffer.

2) *Execute*: Commands in-flight are executed out-of-order, where the **execute** i directive triggers the execution of the i -th command in the buffer. Selected rules are given in Figure 9.

Component	States	Initial state	Functions	
Cache	CacheStates	cs_0	$access : Vals \times CacheStates \rightarrow \{\text{Hit}, \text{Miss}\}$	$update : Vals \times CacheStates \rightarrow CacheStates$
Branch predictor	BpStates	bp_0	$predict : BpStates \times Vals \rightarrow Vals$	$update : BpStates \times Vals \times Vals \rightarrow BpStates$
Pipeline scheduler	ScStates	sc_0	$next : ScStates \rightarrow Dir$	$update : ScStates \times Bufs \rightarrow ScStates$

Fig. 8: Signatures of the microarchitectural components

EXECUTE-LOAD-HIT

$$\frac{\text{sbarr} \notin \text{buf} \quad \text{store } x', e' \notin \text{buf} \quad x \neq \text{pc} \quad |buf| = i - 1 \quad a' = \text{apl}(\text{buf}, a) \quad \langle e \rangle(a') \neq \perp \quad \text{access}(cs, \langle e \rangle(a')) = \text{Hit} \quad \text{update}(cs, \langle e \rangle(a')) = cs'}{\langle m, a, \text{buf} \cdot \text{load } x, e@T \cdot \text{buf}', cs, bp \rangle \xrightarrow{\text{execute } i} \langle m, a, \text{buf} \cdot x \leftarrow m(\langle e \rangle(a'))@T \cdot \text{buf}', cs', bp \rangle}$$

EXECUTE-BRANCH-ROLLBACK

$$\frac{\langle a'(x) = 0 \wedge \ell \neq \ell' \rangle \vee \langle a'(x) \notin Vals \setminus \{0, \perp\} \wedge \ell \neq \ell_0 + 1 \rangle \quad |buf| = i - 1 \quad a' = \text{apl}(\text{buf}, a) \quad \text{sbarr} \notin \text{buf} \quad \ell_0 \neq \varepsilon \quad p(\ell_0) = \text{beqz } x, \ell'' \quad \ell' \in \{\ell'', \ell_0 + 1\} \setminus \{\ell\} \quad bp' = \text{update}(bp, \ell_0, \ell')}{\langle m, a, \text{buf} \cdot \text{pc} \leftarrow \ell@l_0 \cdot \text{buf}', cs, bp \rangle \xrightarrow{\text{execute } i} \langle m, a, \text{buf} \cdot \text{pc} \leftarrow \ell'@e, cs, bp' \rangle}$$

Fig. 9: Selected rules for **execute** i

The rule EXECUTE-LOAD-HIT models the successful execution of a load (**load** $x, e@T$) that results in a cache hit. In the rule, $\langle e \rangle(a')$ denotes the result of evaluating e in the context of the assignment a' obtained by applying to a all earlier in-flight commands in buf . Whenever the address is resolved, i.e., $\langle e \rangle(a') \neq \perp$, and accessing the address results in a cache hit ($\text{access}(cs, \langle e \rangle(a')) = \text{Hit}$), the reorder buffer is updated by replacing **load** $x, e@T$ with $x \leftarrow m(\langle e \rangle(a'))@T$, thereby recording that the load operation has been executed and that the value of x is now $m(\langle e \rangle(a'))$. The cache state is also updated to account for the memory access to $\langle e \rangle(a')$.

In contrast, the EXECUTE-BRANCH-ROLLBACK rule models the resolution of a mis-speculated branch instruction that results in rolling back the speculatively executed instructions by dropping their entries from the reorder buffer. Whenever the predicted value ℓ disagrees with the outcome ℓ' of the instruction **beqz** x, ℓ'' at address ℓ_0 , the buffer is updated by (1) recording the new value of **pc** (by replacing $\text{pc} \leftarrow \ell@l_0$ with $\text{pc} \leftarrow \ell'@e$), and (2) squashing all later buffer entries (by discarding the buffer suffix buf'). Moreover, the branch predictor's state is updated by recording that the branch at address ℓ_0 has been resolved to ℓ' .

3) *Retire*: Instructions are retired in-order. This is done by retiring only commands $i@T$ at the head of the reorder buffer where the instruction i has been resolved, and the tag T is ε indicating that there are no unresolved predictions. Selected rules for the **retire** directive are given below:

RETIRE-ASSIGNMENT

$$\frac{\text{buf} = x \leftarrow v@e \cdot \text{buf}' \quad v \in Vals}{\langle m, a, \text{buf}, cs, bp \rangle \xrightarrow{\text{retire}} \langle m, a[x \mapsto v], \text{buf}', cs, bp \rangle}$$

RETIRE-STORE

$$\frac{\text{buf} = \text{store } v, n@e \cdot \text{buf}' \quad v, n \in Vals \quad \text{update}(cs, n) = cs'}{\langle m, a, \text{buf}, cs, bp \rangle \xrightarrow{\text{retire}} \langle m[n \mapsto v], a, \text{buf}', cs', bp \rangle}$$

The rule RETIRE-ASSIGNMENT models the retirement of a command $x \leftarrow v@e$, where the assignment a is permanently

updated by recording that x 's value is now v . In contrast, RETIRE-STORE models the retirement of store commands **store** $v, n@e$. In this case, the memory m is permanently updated by writing the value v to address n and the cache state is updated. Finally, we have rules RETIRE-SKIP and RETIRE-BARRIER modeling the retirement of **skip** and **sbarr** instructions, which are removed from the reorder buffer without modifying the arch. state.

C. Formalizing the adversary model

We conclude by formalizing the adversary model that we use in the security analysis in Section VI.

In our analysis, we consider an adversary \mathcal{A} that can observe almost the entire microarchitectural state. Specifically, it can observe (1) the data-independent projection of the reorder buffer (i.e., which instructions are in-flight, but not to what values they are resolved), (2) the state of cache (which stores only the addresses of the blocks in the cache, not the blocks themselves), branch predictor, and scheduler. We formalize this as $\mathcal{A} = (\langle m, a, \text{buf}, cs, bp, sc \rangle) = \langle \text{buf} \downarrow, cs, bp, sc \rangle$.

VI. MECHANISMS FOR SECURE SPECULATION

In this section, we show how several recent proposals for hardware-level secure speculation can be cast within our framework and we study their security.

We analyze three countermeasures: (1) disabling speculation (**seq** in §VI-A), (2) delaying *all* speculative loads (**loadDelay** in §VI-B), and (3) employing hardware-level taint tracking and selectively delaying tainted instructions (**tt** in §VI-C). For each countermeasure **ctx**, we formalize its semantics using a relation \Rightarrow_{ctx} obtained by modifying the hardware semantics from §V (which induces the corresponding trace semantics \Downarrow_{ctx} in the usual way). Additionally, we characterize their security guarantees by showing which of the contracts from §III they satisfy; see Figure 11 for a summary of the results.

Unless otherwise specified, all theorems hold for any instantiation of cache, branch predictor, and scheduler.

Before analyzing the countermeasure, we observe that *all* possible instances of the hardware semantics satisfy the $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{spec}}$ contract, as stated in Theorem 1.

Theorem 1. $\{\cdot\} \vdash \llbracket \cdot \rrbracket_{\text{ct}}^{\text{spec}}$.

From this, it immediately follows that *all* countermeasures presented below satisfy the $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{spec}}$ contract as well.

A. **seq**: Disabling speculation

A first, drastic countermeasure against speculative execution attacks is disabling speculative and out-of-order execution. To model this, we instantiate the hardware semantics by providing a sequential scheduler that produces directives in a **fetch** – **execute** 1 – **retire** order. The sequential scheduler, formalized in Appendix B, works as follows:

- Whenever the reorder buffer is empty, the scheduler selects the **fetch** directive that adds entries to the buffer.
- If the first entry in the buffer is not resolved, the scheduler selects the **execute** 1 directive. Thus, the instruction is executed and, potentially, resolved.
- If the first entry in the buffer is resolved, the scheduler selects the **retire** directive. Therefore, the instruction is retired and its changes are written into the architectural state.

That is, the sequential scheduler ensures that instructions are executed in an in-order, non-speculative fashion.

As expected, instantiating the hardware semantics with the sequential scheduler (denoted with **seq**) results in strong security guarantees. As stated in Theorem 2, **seq** implements the $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}}$ interface which exposes only the program counter and the location of memory accesses.

Theorem 2. $\{\cdot\}_{\text{seq}} \vdash \llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}}$.

B. **loadDelay**: Delaying all speculative loads

Sakalis et al. [3] propose a family of countermeasures that delay memory loads to avoid leakage. In the following, we analyze the *eager delay of (speculative) loads* countermeasure. This countermeasure consists in delaying **loads** until all sources of mis-speculation have been resolved. We remark that the hardware semantics of Section V supports speculation only over branch instructions. Therefore, we model the **loadDelay** countermeasure by preventing loads whenever there are preceding, unresolved branch instructions in the reorder buffer. Using the terminology of [3], loads are delayed as long as they are under a so-called *control-shadow*.

We formalize the **loadDelay** countermeasure by modifying the STEP rule of the hardware semantics as follows (changes are highlighted in blue):

STEP-OTHERS

$$\frac{\langle m, a, \text{buf}, cs, bp \rangle \xrightarrow{d} \langle m', a', \text{buf}', cs', bp' \rangle \quad d = \text{next}(sc) \quad sc' = \text{update}(sc, \text{buf}' \downarrow) \quad d \in \{\text{fetch}, \text{retire}\} \vee (d = \text{execute } i \wedge \text{buf}'|_i \neq \text{load } x, e)}{\langle m, a, \text{buf}, cs, bp, sc \rangle \Rightarrow_{\text{loadDelay}} \langle m', a', \text{buf}', cs', bp', sc' \rangle}$$

STEP-EAGER-DELAY

$$\frac{\langle m, a, \text{buf}, cs, bp \rangle \xrightarrow{d} \langle m', a', \text{buf}', cs', bp' \rangle \quad d = \text{next}(sc) \quad sc' = \text{update}(sc, \text{buf}' \downarrow) \quad d = \text{execute } i \quad \text{buf}'|_i = \text{load } x, e \quad \forall \text{pc} \leftarrow \ell @ \ell' \in \text{buf}[0..i-1]. \ell' = \varepsilon}{\langle m, a, \text{buf}, cs, bp, sc \rangle \Rightarrow_{\text{loadDelay}} \langle m', a', \text{buf}', cs', bp', sc' \rangle}$$

Fetching, retiring, and executing all instructions that are not **loads** works as before (see STEP-OTHERS rule). However, **load** instructions are executed only if all prior branch instructions are resolved (see STEP-NAIVE-DELAY rule). This is captured by requiring that all branch instructions in the buffer prefix have tag ε , i.e., $\forall \text{pc} \leftarrow \ell @ \ell' \in \text{buf}[0..i-1]. \ell' = \varepsilon$.

Thus, **loads** are delayed until they are guaranteed to be executed, while other instructions may be freely executed speculatively and out-of-order. Hence, no data memory accesses are performed on mis-speculated paths. However, maybe surprisingly, parts of the architectural state can still be leaked on mis-speculated paths as nested conditional branches may modify the instruction cache and the branch predictor state.

As a consequence, **loadDelay** violates the $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}}$ contract capturing the standard constant-time requirements.

Example 2. This program illustrates that $\{\cdot\}_{\text{loadDelay}} \not\vdash \llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}}$:

```

1  x = A[10]
2  y = not (A[20] | 1)
3  if (y) //branch always unsatisfied
4  if (x) //only reachable speculatively
5  skip

```

Consider two configurations σ and σ' such that $\sigma(\text{A}+10) = 0$ and $\sigma'(\text{A}+10) = 1$. Then, $\llbracket p \rrbracket_{\text{ct}}^{\text{seq}}(\sigma) = \llbracket p \rrbracket_{\text{ct}}^{\text{seq}}(\sigma') = \text{load } \text{A}+10 \cdot \text{load } \text{A}+20 \cdot \text{pc} \perp$. However, the hardware can leak information through, e.g., the instruction cache if the branch at line 3 is speculatively taken. Then, the result of branch at line 4, which determines whether or not **skip** at 5 is fetched, leaks whether $\text{A}[10]$ (stored in x) is 0 or not, thereby distinguishing σ and σ' .

To capture the guarantees offered by the eager-delay countermeasure, we can use the $\llbracket \cdot \rrbracket_{\text{ct-pc}}^{\text{seq-spec}}$ contract, which may intuitively be understood as $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}} + \llbracket \cdot \rrbracket_{\text{pc}}^{\text{spec}}$, i.e., control-flow and memory accesses are leaked under sequential execution, and in addition, the program counter is leaked during speculative execution. This new contract is satisfied by the countermeasure, leading to Theorem 3.

Theorem 3. $\{\cdot\}_{\text{loadDelay}} \vdash \llbracket \cdot \rrbracket_{\text{ct-pc}}^{\text{seq-spec}}$.

As the control flow during speculation execution may only depend upon data previously loaded non-speculatively, the security of the countermeasure can also be captured by $\llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}}$.

Theorem 4. $\{\cdot\}_{\text{loadDelay}} \vdash \llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}}$.

C. **tt**: Taint tracking of speculative values

Recent work [4], [5] propose to track transient computations and to selectively delay instructions involving tainted information. While these proposals slightly differ in how instructions

are labelled and on the effects of different labels, they share the same building blocks and provide similar guarantees.

For this reason, we start by presenting an overview of the Speculative Taint Tracking (STT) [5] and Non-speculative Data Access (NDA) [4] countermeasures. Next, we introduce a general extension to the hardware semantics from Section V for supporting taint-tracking schemes. We continue by formalizing a countermeasure inspired by STT and we discuss its security guarantees, and we conclude by discussing NDA.

1) *Overview*: STT [5] and NDA [4] are two recent taint-tracking proposals for secure speculation. These countermeasures extend a processor with hardware-level taint tracking to track whether data has been retrieved by a speculatively executed instruction. The taint-tracking mechanism propagates taint through the computation and whenever operations are no longer transient, the taint is removed. Finally, both NDA and STT selectively delay tainted operations to avoid leaks.

The main difference between the two approaches is that while STT delays the *execution* of tainted transmit instructions (that is, instructions like **loads** that might leak information), NDA adopts a more conservative approach that delays the *propagation* of data from tainted instructions.

2) *Supporting taint tracking*: To support taint tracking, we label entries in the reorder buffer with two labels: S (which stands for “safe”) and U (which stands for “unsafe”). A labeled command is of the form $\langle c@T \rangle_\ell$ where $c@T$ is a reorder buffer entry and $\ell \in \{S, U\}$ is a label. The labels S and U form a lattice with $S \sqsubseteq U$, and thus for all ℓ , $U \sqcup \ell = U$ and $S \sqcap \ell = S$.

Existing proposals differ in (1) how labels are assigned and propagated, and (2) how labels affect the processor’s execution. To accommodate different variants for (1) and (2), we formalize these aspects using two functions:

- The *labeling function* $lbl(buf_{ul}, buf, d)$ computes the new labels associated with the (unlabeled) buffer buf_{ul} given the old labeled buffer buf and the directive d determining the activated pipeline step. This function models how the tracking works, i.e., how labels are assigned to new instructions and how they are propagated.

- The *unlabeling function* $unlbl(buf, d)$ produces an unlabeled buffer buf_{ul} starting from a labeled buffer buf and a directive d . This function models how labels affect the processor’s semantics in terms of changes to the reorder buffer (and these changes might depend on the executed pipeline step modeled by d).

We describe later how these functions can be instantiated to model STT and NDA.

We formalize the **tt** countermeasure by modifying the STEP rule as follows (changes are highlighted in blue):

$$\begin{array}{c}
 \text{STEP} \\
 d = \text{next}(sc) \quad \boxed{buf_{ul} = unlbl(buf, d)} \\
 \langle m, a, \boxed{buf_{ul}}, cs, bp \rangle \xrightarrow{d} \langle m', a', \boxed{buf'_{ul}}, cs', bp' \rangle \\
 sc' = \text{update}(sc, \boxed{buf'_{ul}} \downarrow) \quad \boxed{buf' = lbl(buf'_{ul}, buf, d)} \\
 \hline
 \langle m, a, \boxed{buf}, cs, bp, sc \rangle \Rightarrow_{\text{tt}} \langle m', a', \boxed{buf'}, cs', bp', sc' \rangle
 \end{array}$$

The rule differs from the standard STEP rule in three ways:

- Entries in the reorder buffer are labelled.
- Before activating a step in the pipeline, i.e., before apply-

$$\begin{aligned}
 unlbl(buf, \mathbf{fetch}) &= mask(buf) \\
 unlbl(buf, \mathbf{retire}) &= drop(buf) \\
 unlbl(buf, \mathbf{execute} \ i) &= \begin{cases} mask(buf) & \text{if } transmit(buf|_i) \\ drop(buf) & \text{otherwise} \end{cases} \\
 drop(\varepsilon) &:= \varepsilon \\
 drop(\langle i@T \rangle_\ell \cdot buf) &:= i@T \cdot drop(buf) \\
 mask(\varepsilon) &:= \varepsilon \\
 mask(\langle i@T \rangle_\ell \cdot buf) &:= \begin{cases} x \leftarrow \perp @T \cdot mask(buf) & \text{if } \ell = U \wedge \\ & i = x \leftarrow e \\ i@T \cdot mask(buf) & \text{otherwise} \end{cases}
 \end{aligned}$$

Fig. 10: Unlabeling function $unlbl(buf, d)$ for STT

ing one step of \xrightarrow{d} , we use the unlabeling function to derive an unlabeled buffer $buf_{ul} = unlbl(buf, d)$ representing how labels affect the reorder buffer entries.

- The buffer produced by the application of \xrightarrow{d} is labeled by invoking the labeling function $buf' = lbl(buf_{ul}, buf, d)$. Therefore, the labels in buf' are updated to track the information flows through the computation.

3) *Speculative taint tracking*: Here we present how to model a countermeasure inspired by STT [5]. As mentioned above, STT tracks whether data depends on speculatively accessed data and delays the execution of transient transmit instructions. These features are reflected in our model:

- In μASM , there are three kinds of *transmit instructions*: loads **load** x, e , stores **store** x, e , and assignments to the program counter **pc** $\leftarrow e$. We write $transmit(i@T)$ whenever the instruction i is a transmit instruction.

- To delay *only* transmit instructions, the unlabeling function, defined in Figure 10, replaces unsafe assignments $x \leftarrow e$ with $x \leftarrow \perp$ for **fetch** and **execute** i directives when the i -th entry in the buffer is a transmit instruction. This ensures that transmit instructions are not executed whenever they depend on unsafe data, which are now mapped to \perp . In contrast, the unlabeling function simply strips the taint-tracking labels for **retire** and **execute** i directives whenever the i -th entry is not a transmit instruction; thereby allowing the hardware to freely execute non-transmit instructions.

- The labeling function, formalized in Appendix C, specifies how newly fetched instructions are labeled as well as how labels are updated during computation, and it works as follows:
 - Newly fetched **load** x, e instructions are labelled as safe if there is no unresolved branch instruction in the buffer, and they are labelled unsafe otherwise. In contrast, newly fetched assignments $x \leftarrow e$ are labelled as unsafe if they depend on unsafe data (i.e., if one of the registers y occurring in e is labelled as unsafe), and they are labelled as safe otherwise. All other newly fetched instructions are labelled as safe.

- Whenever we retire or execute non-branch instructions, labels are preserved.

- When we execute and resolve a branch instruction (thereby eliminating one of the sources of speculation), there are two

cases. If an earlier branch instruction has not been resolved yet, we preserve all labels since all the later instructions are still transient. In contrast, if all earlier branch instructions have been resolved, then we label as safe all following instructions until the next unresolved branch since all these instructions are non-transient. Moreover, we update the labels of the remaining entries in the reorder buffer to account for the non-transient instructions.

Overall, the labeling function ensures that reorder buffer entries that depend on transiently retrieved data are labelled as unsafe at every point of the computation.

Concretely, **tt** delays all transmit instructions that depend on transiently retrieved (i.e., unsafe) data. However, **tt** does not delay transient loads that depend on safe data, as acknowledged also in [5]. This means that parts of the architectural state can be leaked using speculatively executed instructions.

As shown in Example 3, **tt** violates the $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{seq}}$ contract.

Example 3. Consider the Spectre v1 variant from Figure 1b, compiled to μASM :

```

1 load z, A + y //accessing A[y]
2 x ← y < size_A
3 beqz x, ⊥ //checking y < size_A
4 z ← z*64
5 load w, B+z //accessing B[A[y]*64]

```

Consider two configurations σ and σ' that agree on the values of A , B , y , and size_A and for which $\sigma(y) > \sigma(\text{size_A})$, i.e., the array A is speculatively accessed out of bounds. Furthermore, assume that $\sigma(A + y) = 0$ and $\sigma'(A + y) = 1$. Then, $\llbracket p \rrbracket_{\text{ct}}^{\text{seq}}(\sigma) = \llbracket p \rrbracket_{\text{ct}}^{\text{seq}}(\sigma') = \text{load } A+y \cdot \text{pc } \perp$. However, the hardware semantics can potentially leak information through the data cache if the hardware speculatively executes the load on line 5. Indeed, the load on line 1 is labeled as **S** since it is *not* transient. Therefore, the load operation on line 5, which depends on the result of 1, is not delayed (even though operations relying on its result would be delayed since 5 is labeled as **U**). Therefore, by probing the state of the cache an attacker can distinguish whether $A[y] = 0$ or $A[y] = 1$, thereby distinguishing σ and σ' .

One way to characterize the guarantees provided by the **tt** countermeasure is with the $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{spec}}$ contract.

Theorem 5. $\{\cdot\}_{\text{tt}} \vdash \llbracket \cdot \rrbracket_{\text{ct}}^{\text{spec}}$.

However, we remark that this contract is already satisfied by the baseline hardware defined in Section V without any countermeasures. A more meaningful characterization of **tt**'s guarantees, stated in Theorem 6, is via the $\llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}}$ contract. Intuitively, **tt** satisfies $\llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}}$ as it prevents the execution of transmit instructions based on unsafe transiently retrieved data.

Theorem 6. $\{\cdot\}_{\text{tt}} \vdash \llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}}$.

Theorem 6 confirms the results of [5] and provides a clean characterization of the *transient noninterference* [5] guarantees in terms of the $\llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}}$ contract.

4) *Non-speculative data access*: Weisse et al. [4] propose NDA, a family of countermeasures for secure speculation that also relies on hardware taint tracking. In a nutshell, NDA de-




Fig. 11: Security guarantees of secure-speculation mechanisms.

lays the propagation of speculatively executed instructions until the corresponding speculation sources have been resolved. NDA comes with two different propagation strategies—*strict* and *permissive* propagation—that can be modeled as follows:

- For both propagation strategies, the unlabeled function simply replaces all unsafe assignments $x \leftarrow e$ with $x \leftarrow \perp$, thereby preventing the propagation of unsafe data. This differs from STT where labels are sometimes stripped to allow the propagation of unsafe data, as long as their propagation does not leak the data.

- The labeling function differs from the one in **tt** in how newly fetched instructions are labeled. For the strict strategy, all newly fetched transient instructions are labelled as unsafe. In contrast, only newly fetched transient **loads** are labeled as unsafe under the permissive strategy.

Despite these changes, NDA provides similar guarantees to **tt**. That is, it satisfies the $\llbracket \cdot \rrbracket_{\text{ct}}^{\text{spec}}$ and $\llbracket \cdot \rrbracket_{\text{arch}}^{\text{seq}}$ contracts.

D. Summary

Figure 11 summarizes the results of this section in the lattice structure established in §III-C. This yields the first rigorous comparison of the security guarantees of mechanisms for secure speculation, and it translates the results from §IV into a principled basis for programming them securely.

VII. DISCUSSION

A. Scope of the model

With our modeling of a generic microarchitecture and corresponding side-channel adversaries (§V), we aim to strike a balance between capturing the central aspects of attacks on speculative and out-of-order processors, while obtaining a general and tractable model.

As a consequence, we simplified many aspects of modern processors. For instance, we model only a simple 3-stage pipeline, single threaded, and with conditional branch prediction as the only source of speculation. Likewise, we consider an adversary that can observe instructions in the reorder buffer and memory blocks in the cache, but not the data they carry.

This modelling is adequate for reasoning about protections against variants of Spectre v1. However, it does not encompass features such as store-to-load forwarding or prediction over memory aliasing, or adversaries that can observe leaks from internal processor buffers, such as those exploited in data-sampling attacks [29], [30].

As a consequence, Theorems 1–6 need not extend to these scenarios. However, our framework for expressing contracts is not limited to this simple model, as we discuss next.

B. Beyond Spectre v1

We now discuss how to extend our framework to other transient execution attacks. For each attack, we discuss how to (1) extend our contracts, and (2) adjust our hardware semantics:

- *Spectre-BTB and Spectre-RSB*: These variants speculate respectively over indirect jumps and return instructions. To support them, the `spec`-contracts can be extended to explore all possible mispredicted paths for a bounded number of steps before rolling back (similarly to the `BRANCH` rule in Figure 3). Moreover, our hardware semantics $\{\cdot\}$ can also easily be extended to handle these new forms of speculation. For instance, speculation over indirect jumps could be modeled similarly to the `FETCH-BRANCH-HIT` rule in §V.

- *Spectre-STL*: This variant speculates over memory aliasing over in-flight store and load operations. Extending our contracts to handle this new kind of speculation requires to modify the `spec`-contracts to model the effects of store-to-load forwarding resulting from memory aliasing predictions. This could be done similarly to Pitchfork [14]. That is, the `spec` semantics could keep track of the issued `store` x, e instructions. Then, whenever a `load` y, e' instruction is executed, one could explore multiple paths representing all possible aliasing predictions for a fixed number of steps and later roll-back. Finally, the $\{\cdot\}$ semantics can be extended to support Spectre-STL similarly to other semantics [14], [15], [31].

- *Meltdown and MDS*: In Spectre-type issues, transient execution is caused by control and data flow mispredictions. In Meltdown-type [2] issues, transient execution is caused by instruction faults or μ code assists (the latter encompasses data sampling attacks [29], [30]). For reasoning about secure programming under Meltdown-vulnerable processors, one would need contracts such $\llbracket \cdot \rrbracket_{\perp}$, which exposes all the memory space⁵. However, there is limited value in deriving very weak contracts, since it effectively makes secure programming impossible.

C. Uses of contracts

The contracts we propose in this paper are designed to adequately capture the security guarantees offered by existing mechanisms for secure speculation, while exposing tractable verification conditions for software. We envision hardware vendors to produce such contracts for their CPUs, to enable users to reason about software security without exposing details of the microarchitecture, and to provide a baseline against which to validate the vendors' security claims.

Moreover, rather than trying to infer contracts for a microarchitecture that has not been designed with security in mind and is ultimately broken, our framework can serve as a basis for a clean-slate approach, where one starts from a desired security contract and aims to design microarchitectures that optimize performance within these constraints.

VIII. RELATED WORK

Speculative execution attacks: These attacks exploit μ arch. side-effects of speculatively executed instructions to

⁵Another example is a contract exposing the whole page of any loaded value, which would correspond to a processor with hardware prefetchers.

leak information. There exist many Spectre [21] variants that differ in the exploited speculation sources [32], [33], [34], the covert channels [35], [36], [37] used, or the target platforms [38]. We refer to [11], [39] for a survey.

Hardware-level countermeasures: Here, we review proposals that we have not formalized in §VI:

- “Redo”-based countermeasures [6], [7], [9] execute speculative memory operations on shadow cache structures. Once a memory operation becomes non-speculative, its effects are replicated on the standard cache hierarchy by re-executing the operation. While these countermeasures satisfy $\llbracket \cdot \rrbracket_{ct}^{spec}$, they likely violate $\llbracket \cdot \rrbracket_{arch}^{seq}$ as they still modify other parts of the μ arch. state such as the reorder buffer.

- In contrast, “Undo”-based countermeasures [10] mitigate Spectre attacks by rolling back the effects of speculatively executed instructions on the cache. Such countermeasures provide security against adversaries that observe the final cache state, but they do not provide guarantees the trace-based attackers we consider in this paper.

- Delay-based mitigations selectively delay the execution of some instructions to prevent speculative leaks. In addition to the `loadDelay` countermeasure studied in §VI-B, Sakalis et al. [3] propose a more permissive scheme, similar to conditional speculation [40], where only loads resulting in cache misses are delayed. These countermeasures, however, would violate the $\llbracket \cdot \rrbracket_{arch}^{seq}$ and $\llbracket \cdot \rrbracket_{ct-pc}^{seq-spec}$ contracts because cache hits would still leak information.

SpecShield [41] proposes two countermeasures: one similar to eager-delay and the other similar to NDA's permissive strategy, with similar guarantees as those of `loadDelay` and `tt`.

Finally, some proposals, like [42], [43], improve efficiency by only delaying instructions that may leak program-level sensitive information. This is achieved by either considering all user-provided data as untrusted [42] or by allowing the specification of program-level policies [43].

Formal microarchitectural models: While several works [44], [45], [46] present formal arch. models for (parts of) the ARMv8-A, RISC-V, MIPS, and x86 ISAs, only recently researchers started to focus on formal models of μ arch. aspects. For instance, Coppelia [47] is a tool to automatically generate software exploits for hardware designs.

The speculative semantics from [12] forms the basis for the $\llbracket \cdot \rrbracket_{ct}^{spec}$ contract that exposes the effects of speculatively executed instructions. In contrast to [12], other semantics [14], [15], [19], [31] more closely resemble the actual μ arch. behavior of out-of-order processors with multiple pipeline stages, rather than concisely capturing the resulting leakage. Specifically, the hardware semantics $\{\cdot\}$ in §V extends [14], [19]'s semantics by making explicit the dependencies with caches, predictors, and pipeline scheduler.

HW-SW contracts for side channels: Recently, researchers [48], [49] have been calling for new HW-SW contracts that expose security-relevant μ arch. details. We answer this call by providing contracts for secure speculation and by showing how they can be leveraged at the software level.

Recent work [50], [51] presents extensions to the RISC-V ISA where data is labeled, e.g., as *Public* or *Secret*; labels

are tracked during the computation; and the microarchitecture ensures that secret data does not leak. This work is orthogonal to ours in that we characterize the security of different HW-level countermeasures for a standard ISA.

IX. CONCLUSIONS

Motivated by a lack of hardware-software contracts that support principled co-design for secure speculation, we presented a framework for specifying such contracts.

On the hardware side, we used our framework to provide the first uniform characterization of guarantees provided by a representative set of mechanisms for secure speculation.

On the software side, we used our framework to characterize secure programming in two scenarios—“constant-time programming” and “sandboxing”—and we show how to automate checks for programs to run securely on top of these mechanisms.

Acknowledgments: This work was supported by a grant from Intel Corporation, Atracción de Talento Investigador grant 2018-T2/TIC-11732A, Juan de la Cierva-Formación grant FJC2018-036513-I, Spanish project RTI2018-102043-B-I00 SCUM, and Madrid regional project S2018/TCS-4339 BLOQUES.

REFERENCES

- [1] 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.ieeecomputersociety.org/10.1109/SP.2019.00002>
- [2] C. Canella, J. V. Bulck, M. Schwarz, M. Lipp, B. von Berg, P. Ortner, F. Piessens, D. Evtushkin, and D. Gruss, “A systematic evaluation of transient execution attacks and defenses,” in *28th USENIX Security Symposium (USENIX Security 19)*. Santa Clara, CA: USENIX Association, Aug. 2019, pp. 249–266. [Online]. Available: <http://www.usenix.org/conference/usenixsecurity19/presentation/canella>
- [3] C. Sakalis, S. Kaxiras, A. Ros, A. Jimborean, and M. Sjölander, “Efficient invisible speculative execution through selective delay and value prediction,” in *Proceedings of the 46th International Symposium on Computer Architecture*, ser. ISCA ’19. New York, NY, USA: ACM, 2019, pp. 723–735. [Online]. Available: <http://doi.acm.org/10.1145/3307650.3322216>
- [4] O. Weisse, I. Neal, K. Loughlin, T. F. Wenisch, and B. Kasikci, “NDA: Preventing speculative execution attacks at their source,” in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, ser. MICRO ’52. ACM, 2019.
- [5] J. Yu, M. Yan, A. Khyzha, A. Morrison, J. Torrellas, and C. W. Fletcher, “Speculative Taint Tracking (STT): A Comprehensive Protection for Speculatively Accessed Data,” in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, ser. MICRO ’52. New York, NY, USA: ACM, 2019, pp. 954–968. [Online]. Available: <http://doi.acm.org/10.1145/3352460.3358274>
- [6] M. Yan, J. Choi, D. Skarlatos, A. Morrison, C. Fletcher, and J. Torrellas, “Invispec: Making speculative execution invisible in the cache hierarchy,” in *Proceedings - 51st Annual IEEE/ACM International Symposium on Microarchitecture, MICRO 2018*, ser. Proceedings of the Annual International Symposium on Microarchitecture, MICRO. IEEE Computer Society, 12 2018, pp. 428–441.
- [7] K. N. Khasawneh, E. M. Koryueh, C. Song, D. Evtushkin, D. Ponomarev, and N. Abu-Ghazaleh, “Safespec: Banishing the spectre of a meltdown with leakage-free speculation,” in *Proceedings of the 56th Annual Design Automation Conference 2019*, ser. DAC ’19. New York, NY, USA: ACM, 2019, pp. 60:1–60:6. [Online]. Available: <http://doi.acm.org/10.1145/3316781.3317903>
- [8] V. Kiriansky, I. A. Lebedev, S. P. Amarasinghe, S. Devadas, and J. S. Emer, “DAWG: A defense against cache timing attacks in speculative execution processors,” in *51st Annual IEEE/ACM International Symposium on Microarchitecture, MICRO 2018, Fukuoka, Japan, October 20-24, 2018*, 2018, pp. 974–987. [Online]. Available: <https://doi.org/10.1109/MICRO.2018.00083>
- [9] S. Anisworth and T. M. Jones, “Muontrap: Preventing cross-domain spectre-like attacks by capturing speculative state,” in *Proceedings of the 47th International Symposium on Computer Architecture*, ser. ISCA ’20, 2020.
- [10] G. Saileshwar and M. K. Qureshi, “Cleanupspec: An “undo” approach to safe speculation,” in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, 2019, pp. 73–86.
- [11] C. Canella, J. Van Bulck, M. Schwarz, M. Lipp, B. von Berg, P. Ortner, F. Piessens, D. Evtushkin, and D. Gruss, “A Systematic Evaluation of Transient Execution Attacks and Defenses,” in *Proceedings of the 28th USENIX Security Symposium*, ser. USENIX Security ’19. USENIX Association, 2019.
- [12] 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.
- [13] G. Barthe, G. Betarte, J. Campo, C. Luna, and D. Pichardie, “System-level non-interference for constant-time cryptography,” in *CCS*. ACM, 2014.
- [14] S. Cauligi, C. Disselkoben, K. v. Gleissenthall, D. Stefan, T. Rezk, and G. Barthe, “Towards constant-time foundations for the new spectre era,” 2019.
- [15] M. Balliu, M. Dam, and R. Guanciale, “Inspectre: Breaking and fixing microarchitectural vulnerabilities by formal analysis,” 2019.
- [16] C. Carruth, “Speculative load hardening,” 2018. [Online]. Available: <http://releases.lvm.org/8.0.0/docs/SpeculativeLoadHardening.html>
- [17] M. Miller, “Mitigating speculative execution side channel hardware vulnerabilities,” <https://blogs.technet.microsoft.com/srd/2018/03/15/mitigating-speculative-execution-side-channel-hardware-vulnerabilities/>, 2018.
- [18] J. B. Almeida, M. Barbosa, G. Barthe, F. Dupressoir, and M. Emmi, “Verifying constant-time implementations,” in *USENIX Security Symposium*. USENIX Association, 2016, pp. 53–70.
- [19] M. Vassena, K. v. Gleissenthall, R. G. Kici, D. Stefan, and R. Jhala, “Automatically eliminating speculative leaks with blade,” 2019.
- [20] G. Barthe, G. Betarte, J. D. Campo, and C. Luna, “System-level non-interference of constant-time cryptography. part I: model,” *J. Autom. Reasoning*, vol. 63, no. 1, pp. 1–51, 2019.
- [21] 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 *Proceedings of the 40th IEEE Symposium on Security and Privacy*, ser. S&P ’19. IEEE, 2019.
- [22] J. Landauer and T. Redmond, “A lattice of information,” in *CSFW*, 1993, pp. 65–70.
- [23] B. Yee, D. Sehr, G. Dardyk, J. B. Chen, R. Muth, T. Ormandy, S. Okasaka, N. Narula, and N. Fullagar, “Native client: A sandbox for portable, untrusted x86 native code,” *Commun. ACM*, vol. 53, no. 1, pp. 91–99, Jan. 2010. [Online]. Available: <http://doi.acm.org/10.1145/1629175.1629203>
- [24] A. Haas, A. Rossberg, D. L. Schuff, B. L. Titzer, M. Holman, D. Gohman, L. Wagner, A. Zakai, and J. Bastien, “Bringing the web up to speed with webassembly,” in *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI 2017. New York, NY, USA: Association for Computing Machinery, 2017, p. 185200. [Online]. Available: <https://doi.org/10.1145/3062341.3062363>
- [25] B. Rodrigues, F. M. Quintão Pereira, and D. F. Aranha, “Sparse representation of implicit flows with applications to side-channel detection,” in *Proceedings of the 25th International Conference on Compiler Construction*, ser. CC 2016. New York, NY, USA: Association for Computing Machinery, 2016, p. 110120. [Online]. Available: <https://doi.org/10.1145/2892208.2892230>
- [26] D. Molnar, M. Piotrowski, D. Schultz, and D. A. Wagner, “The program counter security model: Automatic detection and removal of control-flow side channel attacks,” in *Information Security and Cryptology - ICISC 2005, 8th International Conference, Seoul, Korea, December 1-2, 2005, Revised Selected Papers*, ser. Lecture Notes in Computer Science, D. Won and S. Kim, Eds., vol. 3935. Springer, 2005, pp. 156–168. [Online]. Available: https://doi.org/10.1007/11734727_14
- [27] S. Cauligi, G. Soeller, B. Johannsmeyer, F. Brown, R. S. Wahby, J. Renner, B. Grégoire, G. Barthe, R. Jhala, and D. Stefan, “Fact: A dsl for timing-sensitive computation,” in *Proceedings of the 40th*

- ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI 2019. New York, NY, USA: Association for Computing Machinery, 2019, p. 174189. [Online]. Available: <https://doi.org/10.1145/3314221.3314605>
- [28] G. Barthe, S. Blazy, B. Grégoire, R. Hutin, V. Laporte, D. Pichardie, and A. Trieu, “Formal verification of a constant-time preserving c compiler,” *Proc. ACM Program. Lang.*, vol. 4, no. POPL, Dec. 2019. [Online]. Available: <https://doi.org/10.1145/3371075>
- [29] S. van Schaik, A. Milburn, S. sterlund, P. Frigo, G. Maisuradze, K. Razavi, H. Bos, and C. Giuffrida, “RIDL: Rogue in-flight data load,” in *S&P*, May 2019.
- [30] 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.
- [31] R. McIlroy, J. Sevcik, T. Tebbi, B. L. Titzer, and T. Verwaest, “Spectre is here to stay: An analysis of side-channels and speculative execution,” *CoRR*, vol. abs/1902.05178, 2019.
- [32] G. Maisuradze and C. Rossow, “Ret2Spec: Speculative Execution Using Return Stack Buffers,” in *Proceedings of the 25th ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’18. ACM, 2018.
- [33] E. M. Koruyeh, K. N. Khasawneh, C. Song, and N. Abu-Ghazaleh, “Spectre returns! speculation attacks using the return stack buffer,” in *Proceedings of the 12th USENIX Workshop on Offensive Technologies*, ser. WOOT ’18. USENIX Association, 2018.
- [34] J. Horn, “CVE-2018-3639 - speculative store bypass,” <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2018-3639>, 2018.
- [35] C. Trippel, D. Lustig, and M. Martonosi, “MeltdownPrime and SpectrePrime: Automatically-synthesized attacks exploiting invalidation-based coherence protocols,” *CoRR*, vol. abs/1802.03802, 2018.
- [36] M. Schwarz, M. Schwarzl, M. Lipp, and D. Gruss, “Netspectre: Read arbitrary memory over network,” in *ESORICS*, 2019.
- [37] J. Stecklina and T. Prescher, “LazyFP: Leaking FPU register state using microarchitectural side-channels,” *CoRR*, vol. abs/1806.07480, 2018.
- [38] G. Chen, S. Chen, Y. Xiao, Y. Zhang, Z. Lin, and T. H. Lai, “Stealing intel secrets from SGX enclaves via speculative execution,” in *Proceedings of the 4th IEEE European Symposium on Security and Privacy*, ser. EuroS&P ’19. IEEE, 2019.
- [39] W. Xiong and J. Szefer, “Survey of transient execution attacks,” 2020.
- [40] P. Li, L. Zhao, R. Hou, L. Zhang, and D. Meng, “Conditional speculation: An effective approach to safeguard out-of-order execution against spectre attacks,” in *2019 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, 2019, pp. 264–276.
- [41] K. Barber, A. Bacha, L. Zhou, Y. Zhang, and R. Teodorescu, “Specshield: Shielding speculative data from microarchitectural covert channels,” in *2019 28th International Conference on Parallel Architectures and Compilation Techniques (PACT)*. IEEE, 2019, pp. 151–164.
- [42] M. Taram, A. Venkat, and D. Tullsen, “Context-sensitive fencing: Securing speculative execution via microcode customization,” in *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, ser. ASPLOS 19. ACM, 2019, p. 395410.
- [43] M. Schwarz, M. Lipp, C. Canella, R. Schilling, F. Kargl, and D. Gruss, “Context: A generic approach for mitigating spectre,” in *Proceedings of the 27th Annual Network and Distributed System Security Symposium (NDSS20)*. Internet Society, Reston, VA, 2020.
- [44] A. Armstrong, T. Bauereiss, B. Campbell, A. Reid, K. E. Gray, R. M. Norton, P. Mundkur, M. Wassell, J. French, C. Pulte, S. Flur, I. Stark, N. Krishnaswami, and P. Sewell, “ISA semantics for ARMv8-a, RISC-v, and CHERI-MIPS,” *Proceedings of the ACM on Programming Languages*, vol. 3, no. POPL, 2019.
- [45] U. Degenbaev, “Formal specification of the x86 instruction set architecture,” Ph.D. dissertation, Universität des Saarlandes, 2012.
- [46] S. Goel, W. A. Hunt, and M. Kaufmann, *Engineering a Formal, Executable x86 ISA Simulator for Software Verification*. Springer, 2017.
- [47] R. Zhang, C. Deutschbein, P. Huang, and C. Sturton, “End-to-end automated exploit generation for validating the security of processor designs,” in *Proceedings of the 51st Annual IEEE/ACM International Symposium on Microarchitecture*, ser. MICRO ’18. IEEE/ACM, 2018.
- [48] G. Heiser, “For safety’s sake: We need a new hardware-software contract!” *IEEE Design and Test*, vol. 35, pp. 27–30, 2018.
- [49] Q. Ge, Y. Yarom, and G. Heiser, “No security without time protection: We need a new hardware-software contract,” in *Proceedings of the 9th Asia-Pacific Workshop on Systems*, ser. APSys 18. New York, NY, USA: Association for Computing Machinery, 2018. [Online]. Available: <https://doi.org/10.1145/3265723.3265724>
- [50] J. Yu, L. Hsiung, M. E. Hajj, and C. W. Fletcher, “Data oblivious ISA extensions for side channel-resistant and high performance computing,” in *NDSS*. The Internet Society, 2019.
- [51] D. Zagieboylo, G. E. Suh, and A. C. Myers, “Using information flow to design an ISA that controls timing channels,” in *CSF*. IEEE, 2019, pp. 272–287.

APPENDIX A
ARCHITECTURAL SEMANTICS

The architectural semantics for μ ASM programs is presented in Figure 12.

APPENDIX B
SEQUENTIAL SCHEDULER

Here, we formalize the sequential scheduler from §VI-A. The sequential scheduler Seq is defined as the 4-tuple $\langle ScStates, sc_0, next, update \rangle$ where the components are as follows:

$$\begin{aligned}
 ScStates &:= \{buf \downarrow \mid buf \in Bufs\} \\
 sc_0 &:= \varepsilon \\
 next : ScStates &\rightarrow Dir := \\
 next(\varepsilon) &= \mathbf{fetch} \\
 next(c \cdot buf) &= \begin{cases} \mathbf{execute} \ 1 & \text{if } exec(c) \\ \mathbf{retire} & \text{otherwise} \end{cases} \\
 exec(\mathbf{skip}@T) &= \perp \\
 exec(\mathbf{sbarr}@T) &= \perp \\
 exec(x \leftarrow \mathbf{probe}@T) &= \top \\
 exec(x \leftarrow e@T) &= \begin{cases} \top & \text{if } e = UR \vee T \neq \varepsilon \\ \perp & \text{otherwise} \end{cases} \\
 exec(\mathbf{load} \ x, e@T) &= \top \\
 exec(\mathbf{store} \ x, e@T) &= \begin{cases} \top & \text{if } e = UR \vee x = UR \\ \perp & \text{otherwise} \end{cases} \\
 update : ScStates \times Bufs &\rightarrow ScStates := \\
 update(sc, buf) &= buf
 \end{aligned}$$

APPENDIX C
LABELING FUNCTION FOR \mathbf{tt}

The labeling function for the STT countermeasure is given in Figure 13.

Expression evaluation

$$\llbracket n \rrbracket(a) = n \quad \llbracket x \rrbracket(a) = a(x) \quad \llbracket \ominus e \rrbracket(a) = \ominus \llbracket e \rrbracket(a) \quad \llbracket e_1 \otimes e_2 \rrbracket(a) = \llbracket e_1 \rrbracket(a) \otimes \llbracket e_2 \rrbracket(a)$$

Instruction evaluation

SKIP

$$\frac{p(a(\mathbf{pc})) = \mathbf{skip}}{\langle m, a \rangle \rightarrow \langle m, a[\mathbf{pc} \mapsto a(\mathbf{pc}) + 1] \rangle}$$

BARRIER

$$\frac{p(a(\mathbf{pc})) = \mathbf{sbarr}}{\langle m, a \rangle \rightarrow \langle m, a[\mathbf{pc} \mapsto a(\mathbf{pc}) + 1] \rangle}$$

ASSIGN

$$\frac{p(a(\mathbf{pc})) = x \leftarrow e \quad x \neq \mathbf{pc}}{\langle m, a \rangle \rightarrow \langle m, a[\mathbf{pc} \mapsto a(\mathbf{pc}) + 1, x \mapsto \llbracket e \rrbracket(a)] \rangle}$$

CONDITIONALUPDATE-SAT

$$\frac{p(a(\mathbf{pc})) = x \stackrel{e'}{\leftarrow} e \quad \llbracket e' \rrbracket(a) = 0 \quad x \neq \mathbf{pc}}{\langle m, a \rangle \rightarrow \langle m, a[\mathbf{pc} \mapsto a(\mathbf{pc}) + 1, x \mapsto \llbracket e \rrbracket(a)] \rangle}$$

CONDITIONALUPDATE-UNSAT

$$\frac{p(a(\mathbf{pc})) = x \stackrel{e'}{\leftarrow} e \quad \llbracket e' \rrbracket(a) \neq 0 \quad x \neq \mathbf{pc}}{\langle m, a \rangle \rightarrow \langle m, a[\mathbf{pc} \mapsto a(\mathbf{pc}) + 1] \rangle}$$

TERMINATE

$$\frac{p(a(\mathbf{pc})) = \perp}{\langle m, a \rangle \rightarrow \langle m, a[\mathbf{pc} \mapsto \perp] \rangle}$$

LOAD

$$\frac{p(a(\mathbf{pc})) = \mathbf{load} \ x, e \quad x \neq \mathbf{pc} \quad n = \llbracket e \rrbracket(a)}{\langle m, a \rangle \rightarrow \langle m, a[\mathbf{pc} \mapsto a(\mathbf{pc}) + 1, x \mapsto m(n)] \rangle}$$

STORE

$$\frac{p(a(\mathbf{pc})) = \mathbf{store} \ x, e \quad n = \llbracket e \rrbracket(a)}{\langle m, a \rangle \rightarrow \langle m[n \mapsto a(x)], a[\mathbf{pc} \mapsto a(\mathbf{pc}) + 1] \rangle}$$

BEQZ-SAT

$$\frac{p(a(\mathbf{pc})) = \mathbf{beqz} \ x, \ell \quad a(x) = 0}{\langle m, a \rangle \rightarrow \langle m, a[\mathbf{pc} \mapsto \ell] \rangle}$$

BEQZ-UNSAT

$$\frac{p(a(\mathbf{pc})) = \mathbf{beqz} \ x, \ell \quad a(x) \neq 0}{\langle m, a \rangle \rightarrow \langle m, a[\mathbf{pc} \mapsto a(\mathbf{pc}) + 1] \rangle}$$

JMP

$$\frac{p(a(\mathbf{pc})) = \mathbf{jmp} \ e \quad \ell = \llbracket e \rrbracket(a)}{\langle m, a \rangle \rightarrow \langle m, a[\mathbf{pc} \mapsto \ell] \rangle}$$

Fig. 12: Architectural semantics for a μ ASM program p

$$\begin{aligned}
\text{labels}(\text{buf}) &= \text{derive}(\text{buf}, \lambda x \in \text{Regs. } \mathbb{S}) \\
\text{labels}(\text{buf})(x) &= \text{derive}(\text{buf}, \lambda x \in \text{Regs. } \mathbb{S})(x) \\
\text{labels}(\text{buf})(e) &= \bigsqcup_{x \in \text{vars}(e)} \text{derive}(\text{buf}, \lambda x \in \text{Regs. } \mathbb{S})(x)
\end{aligned}$$

$$\begin{aligned}
\text{derive}(\varepsilon, \Lambda) &= \Lambda \\
\text{derive}(\langle i@T \rangle_1 \cdot \text{buf}, \Lambda) &= \begin{cases} \text{derive}(\text{buf}, \Lambda[x \mapsto 1]) & \text{if } i = \mathbf{load} \ x, e \vee x \leftarrow e \\ \text{derive}(\text{buf}, \Lambda) & \text{otherwise} \end{cases} \\
\text{relbl}(\langle i@T \rangle_1 \cdot \text{buf}, \Lambda) &= \begin{cases} \langle i@T \rangle_U \cdot \text{relbl}(\text{buf}, \Lambda[x \mapsto U]) & \text{if } i = \mathbf{load} \ x, e \\ \langle i@T \rangle_{1'} \cdot \text{relbl}(\text{buf}, \Lambda[x \mapsto 1']) & \text{if } i = x \leftarrow e \wedge 1' = \bigsqcup_{y \in \text{vars}(e)} \Lambda(y) \\ \langle i@T \rangle_1 \cdot \text{relbl}(\text{buf}, \Lambda) & \text{otherwise} \end{cases}
\end{aligned}$$

$$\text{lbl}(\text{buf}_{ul} \cdot \mathbf{pc} \leftarrow e@E, \text{buf}, \mathbf{fetch}) = \text{buf} \cdot \langle \mathbf{pc} \leftarrow e@E \rangle_S$$

$$\text{lbl}(\text{buf}_{ul} \cdot \mathbf{pc} \leftarrow \ell@l_0, \text{buf}, \mathbf{fetch}) = \text{buf} \cdot \langle \mathbf{pc} \leftarrow \ell@l_0 \rangle_S$$

$$\text{lbl}(\text{buf}_{ul} \cdot i@T \cdot \mathbf{pc} \leftarrow \ell@E, \text{buf}, \mathbf{fetch}) = \text{buf} \cdot \langle i@T \rangle_S \cdot \langle \mathbf{pc} \leftarrow \ell@E \rangle_S$$

$$\text{where } i \neq \mathbf{load} \ x, e \wedge i \neq x \leftarrow e$$

$$\text{lbl}(\text{buf}_{ul} \cdot x \leftarrow e@T \cdot \mathbf{pc} \leftarrow \ell@E, \text{buf}, \mathbf{fetch}) = \text{buf} \cdot \langle x \leftarrow e@T \rangle_{\text{labels}(\text{buf})(e)} \cdot \langle \mathbf{pc} \leftarrow \ell@E \rangle_S$$

$$\text{where } x \neq \mathbf{pc}$$

$$\text{lbl}(\text{buf}_{ul} \cdot \mathbf{load} \ x, e@T \cdot \mathbf{pc} \leftarrow \ell@E, \text{buf}, \mathbf{fetch}) = \text{buf} \cdot \langle \mathbf{load} \ x, e@T \rangle_U \cdot \langle \mathbf{pc} \leftarrow \ell@E \rangle_S$$

$$\text{where } \exists i'@T' \in \text{buf}. T' \neq E$$

$$\text{lbl}(\text{buf}_{ul} \cdot \mathbf{load} \ x, e@T \cdot \mathbf{pc} \leftarrow \ell@E, \text{buf}, \mathbf{fetch}) = \text{buf} \cdot \langle \mathbf{load} \ x, e@T \rangle_S \cdot \langle \mathbf{pc} \leftarrow \ell@E \rangle_S$$

$$\text{where } \forall \mathbf{pc} \leftarrow \ell@T' \in \text{buf}. T' = E$$

$$\text{lbl}(\text{buf}_{ul}, \text{buf}, \mathbf{execute} \ i) = \text{buf}[0..i-1] \cdot \langle \text{buf}_{ul}|_i \rangle_1 \cdot \text{buf}[i+1..|\text{buf}|]$$

$$\text{where } |\text{buf}_{ul}| = |\text{buf}| \wedge \text{buf}|_i = \langle c \rangle_1 \wedge \forall \ell, \ell_0 \in \text{Vals}. c \neq \mathbf{pc} \leftarrow \ell@l_0$$

$$\text{lbl}(\text{buf}_{ul}, \text{buf}, \mathbf{execute} \ i) = \text{buf}[0..i-1] \cdot \langle \text{buf}_{ul}|_i \rangle_1 \cdot \text{buf}[i+1..|\text{buf}|]$$

$$\text{where } |\text{buf}_{ul}| = |\text{buf}| \wedge \text{buf}|_i = \langle \mathbf{pc} \leftarrow e@E \rangle_1$$

$$\text{lbl}(\text{buf}_{ul}, \text{buf}, \mathbf{execute} \ i) = \text{buf}[0..i-1] \cdot \langle \text{buf}_{ul}|_i \rangle_S \cdot \langle \text{drop}(\text{buf}[i+1..j]) \rangle_S \cdot \text{relbl}(\text{buf}[j+1..|\text{buf}|], \lambda x \in \text{Regs. } \mathbb{S})$$

$$\text{where } |\text{buf}_{ul}| = |\text{buf}| \wedge \text{buf}|_i = \langle \mathbf{pc} \leftarrow \ell@l_0 \rangle_1 \wedge \ell_0 \neq E \wedge \text{buf}_{ul}|_i = \mathbf{pc} \leftarrow \ell'@E \wedge$$

$$\forall \mathbf{pc} \leftarrow v@T \in \text{buf}[0..i-1]. T = E \wedge$$

$$j = \min(\{i \in \mathbb{N} \mid \text{buf}|_i = \langle \mathbf{pc} \leftarrow \ell@l_0 \rangle_1 \wedge \ell_0 \neq E\} \cup \{|\text{buf}|\})$$

$$\text{lbl}(\text{buf}_{ul}, \text{buf}, \mathbf{execute} \ i) = \text{buf}[0..i-1] \cdot \langle \text{buf}_{ul}|_i \rangle_S \cdot \text{buf}[i+1..|\text{buf}|]$$

$$\text{where } |\text{buf}_{ul}| = |\text{buf}| \wedge \text{buf}|_i = \langle \mathbf{pc} \leftarrow \ell@l_0 \rangle_1 \wedge \ell_0 \neq E \wedge \text{buf}_{ul}|_i = \mathbf{pc} \leftarrow \ell'@E \wedge$$

$$\exists \mathbf{pc} \leftarrow v@T \in \text{buf}[0..i-1]. T \neq E$$

$$\text{lbl}(\text{buf}_{ul}, \text{buf}, \mathbf{execute} \ i) = \text{buf}[0..i-1] \cdot \langle \text{buf}_{ul}|_i \rangle_1 \cdot \text{buf}[i+1..|\text{buf}|]$$

$$\text{where } |\text{buf}_{ul}| = |\text{buf}| \wedge \text{buf}|_i = \langle c \rangle_1 \wedge c = \mathbf{pc} \leftarrow \ell@l_0 \wedge \text{buf}_{ul}|_i = \mathbf{pc} \leftarrow \ell@l_0 \wedge \ell_0 \neq E$$

$$\text{lbl}(\text{buf}_{ul}, \text{buf}, \mathbf{execute} \ i) = \text{buf}[0..i-1] \cdot \langle \text{buf}_{ul}|_i \rangle_S$$

$$\text{where } |\text{buf}_{ul}| = i \wedge \text{buf}|_i = \langle \mathbf{pc} \leftarrow \ell@l_0 \rangle_1 \wedge \ell_0 \neq E \wedge \text{buf}_{ul}|_i = \mathbf{pc} \leftarrow \ell'@E$$

$$\text{lbl}(\text{buf}_{ul}, \langle i@T \rangle_\ell \cdot \text{buf}, \mathbf{retire}) = \text{buf}$$

Fig. 13: Labeling function $\text{lbl}(\text{buf}', \text{buf}, d)$ for STT