

# A Formal Approach to Secure Speculation

Kevin Cheang Cameron Rasmussen Sanjit Seshia  
{kcheang, crasmussen, sseshia}@berkeley.edu  
University of California, Berkeley

Pramod Subramanian  
spramod@cse.iitk.ac.in  
Indian Institute of Technology, Kanpur

**Abstract**—Transient execution attacks like Spectre, Meltdown and Foreshadow have shown that combinations of microarchitectural side-channels can be synergistically exploited to create side-channel leaks that are greater than the sum of their parts. While both hardware and software mitigations have been proposed against these attacks, provable security has remained elusive.

This paper introduces a formal methodology for enabling secure speculative execution on modern processors. We propose a new class of information flow security properties called trace property-dependent observational determinism (TPOD). We use this class to formulate a secure speculation property. Our formulation precisely characterises all transient execution vulnerabilities. We demonstrate its applicability by verifying secure speculation for several illustrative programs.

## I. INTRODUCTION

Recently discovered transient execution attacks like Spectre, Meltdown and Foreshadow [1–13] have shown that side channel vulnerabilities are more exploitable than was previously believed. While caches or branch predictors [14–20] leaking information is not exactly news in 2019, Spectre is interesting because it combines side channels to produce a leak that is “greater than the sum of its parts.” A number of mitigations have been proposed to these vulnerabilities [9–12, 21–25] and many of these have been adopted into widely-used software like the Linux kernel [26, 27], Microsoft Windows [28, 29] and the Microsoft Visual Studio compilers and associated libraries [23]. However, these mitigations are not provably secure and in fact some, e.g. Spectre mitigations in Microsoft Visual Studio, are known to be incomplete [30].

Transient execution attacks exploit microarchitectural side-channels present in modern high-performance processors. High-performance processors contain several microarchitectural optimizations — e.g., branch prediction, data and instruction caching, out-of-order execution, and speculative memory address disambiguation, to name just a few — in order to execute programs more efficiently [31, 32]. Many of these optimizations rely on the technique of speculation [33–36]. The processor uses a prediction structure to guess whether a particular execution is likely to occur before its results are available and speculatively executes as per the prediction. If the prediction turns out to be wrong, architectural state — which consists of register and memory values — is restored to its value before speculative execution started and execution restarts along the correct path. In many cases, it is possible to build predictors that mostly guess correctly, and speculation leads to huge performance and power benefits.

It is important to emphasize that when misspeculation is resolved, only architectural state is restored, while *microarchi-*

*tectural* state, such as cache and branch predictor state, is *not*. Transient execution attacks exploit this fact by mistraining a prediction structure to speculatively execute vulnerable wrong-path instructions and exfiltrate confidential information by examining the microarchitectural side-effects of misspeculation.

The above leads to two obvious templates for preventing these vulnerabilities: (i) do not speculate, or (ii) do not leak information through microarchitectural side channels. Many mitigations do indeed take the first approach by turning off speculation in a targeted manner [9, 10, 21, 24, 26–29]. While most of these mitigations were developed on an ad hoc basis after careful manual analysis of known exploitable vulnerabilities, automated tools for Spectre mitigation have also taken this approach [23]. Unfortunately, the latter have been found to be incomplete [30] while the former do not come with provable security guarantees. The larger point here is that there is no formal methodology for reasoning about the security of mitigations to transient execution vulnerabilities.

Some research has also taken the second approach of attempting to close the exfiltration side-channel. For instance the Dynamically Allocated Way Guard (DAWG) closes the cache side-channel by partitioning between protection domains [37]. However, other side-channels (prefetchers, DRAM row buffers, load store queues, etc.) potentially remain exploitable with these solutions and partitioning comes with a significant performance penalty. Here too, it remains unclear whether partitioning a few exfiltration channels is sufficient to prevent all transient execution vulnerabilities.

Besides the lack of provable security, another problem with current approaches are their large performance penalties. In this context, it is noteworthy that recent versions of the Linux Kernel have turned off certain Spectre mitigations by default because performance slowdowns of up to 50% [38, 39] were observed for certain workloads. We believe these high overheads are largely a result of being unable to reason about security of the mitigations. If we could systematically reason about security, we believe it will be possible to develop more aggressive mitigations that disable speculation in a very targeted manner and have a much lower performance overhead.

All of the above points to the need for verification techniques for secure speculation. This problem is most closely related to the secure information flow problem, which has been studied by a rich body of literature [40–46]. Unfortunately, existing work on secure information flow is not sufficient to precisely capture the class of transient execution vulnerabilities. Specifically, it is important to note that traditional notions of information flow security like non-interference [41]

and observational determinism [42–44] are only satisfied when there is *no information flow* from confidential state to adversary observable state. In the context of Spectre, this would imply no information flow from confidential memory locations to microarchitectural side channels. Unfortunately for most programs of interest, all modern processors leak a lot of information through microarchitectural side channels like caches, prefetchers, DRAM row buffers, etc. Therefore, traditional formulations of secure information flow are always violated for such programs regardless of whether they are vulnerable to transient execution attacks.

The above points to one of the key challenges in the verification of secure speculation: *formulating the right property*. We need a way of precisely capturing the *new* leaks introduced by the interaction of microarchitectural side channels with speculation. A second important challenge is coming up with a general system and adversary model that can be used to reason about all possible transient execution attacks, as opposed to pattern-matching known vulnerabilities. Finally, we need a verification methodology that can be used to prove that specific programs satisfy secure speculation.

In this paper, we address each of the above challenges. We introduce a formal methodology for reasoning about security against transient execution attacks. Our approach is based on the formulation of a new class of information flow security properties called *trace property-dependent observational determinism* (TPOD). These properties, an extension of observational determinism, are defined with respect to a trace property and intuitively TPOD captures the following notion of security: *does violation of the trace property introduce new counterexamples to observational determinism?*

We use TPOD to reason about the security of software-based Spectre mitigations. For this, we present an assembly intermediate representation (AIR) into which machine code can be lifted and introduce speculative operational semantics for this AIR. We introduce a general adversary that captures all transient execution attacks, and define a secure speculation property against this adversary as an instance of TPOD. We verify secure speculation using bounded model checking and induction using the UCLID5 verification tool [47, 48] on a suite of small but illustrative benchmarks, several of which are from the literature on Spectre mitigations [30].

### A. Contributions

This paper’s contributions are the following.

- We introduce a methodology for reasoning about the security of microarchitectural speculation mechanisms. To the best of our knowledge, ours is the first effort that can guarantee provably secure speculation.
- We introduce a new class of information-flow security properties called trace property-dependent observational determinism. This class of properties allows us to reason about information leaks that occur due to interactions between microarchitectural mechanisms.
- We introduce a speculative operational semantics for an assembly intermediate representation, an adversary model

for transient execution attacks over this representation and a secure speculation property. Violations of the property correspond to transient execution vulnerabilities.

- We demonstrate viability of our methodology by proving secure speculation for a suite of small programs.

The rest of this paper is organized as follows. Section II presents an overview of transient execution attacks and associated vulnerabilities. Section III reviews observational determinism and introduces trace property-dependent observational determinism. Section IV describes the assembly intermediate representation, speculative operational semantics for it, the adversary model and the secure speculation property. Sections V and VI present our verification approach and case studies. Section VII reviews related work and section VIII provides some concluding remarks.

## II. OVERVIEW

In this section, we present an overview of transient execution vulnerabilities as exemplified by Spectre and review the verification challenges posed by these vulnerabilities.

### A. Introduction to Transient Execution Attacks

Transient execution attacks involve two components: an untrusted component (the attacker) who interacts with a trusted component (the victim) over some communication interface. The attacker exfiltrates confidential information from the victim by exploiting microarchitectural artifacts of misspeculation in high-performance processors. As shown in Figure 1, a transient execution attack has four stages. We explain these four stages using the code snippet shown in Figure 3(a), which is vulnerable to Spectre variant 1.

*(S1) Prepare:* In the first stage, the attacker prepares the exfiltration side channel and mistrains the branch predictor in order to bring the system into a vulnerable state. A commonly used exfiltration side channel is the cache. In this context, preparation refers to priming the cache by executing load/store instructions with effective addresses that are stored in the same cache sets as the victim data. One way to mistrain the branch predictor is to repeatedly execute the victim code with carefully chosen input arguments so that predictor learns that a particular branch should always be taken (or not taken).

*(S2) Invocation:* In the second stage, the attacker invokes victim code with carefully chosen input arguments to trigger

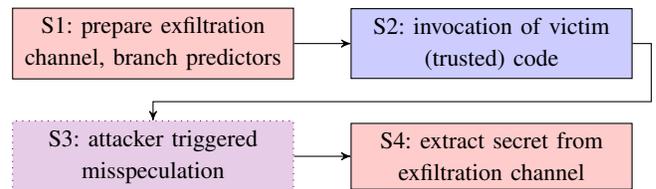


Figure 1: Four Stages of a Speculative Execution Attack. Execution of untrusted code is shown in red, while trusted code is in blue. We show the attacker-triggered misspeculation in the trusted code in the violet dotted box.

misspeculation. This invocation occurs over some communication interface between the untrusted and trusted code. One such interface is through system calls and returns; here the attacker is an untrusted user-mode process while the victim is the operating system kernel. Another example in the context of browser-based sandboxing, e.g. Native Client [49], would be function calls and returns. The attacker is untrusted code running within the sandbox while the victim is the NaCl’s trusted API. Many other vulnerable interfaces exist: hypercalls, enclave entry, software interrupts, etc. The victim may not even be explicitly invoked: an implicit invocation by mistraining the branch predictor is possible! For simplicity, this paper focuses on a function call/return interface but our techniques are easily generalized to other interfaces.

(S3) *Exploitable Misspeculation*: The victim code now executes. At some point it will misspeculate in an attacker-controlled manner resulting in the execution of “wrong path” instructions. These wrong path instructions update speculative architectural state – register and memory values – and microarchitectural state including caches, branch predictors and prefetchers. Eventually the wrong path is resolved and its instructions are flushed. Speculative updates to architectural state (registers and memory) are flushed, but microarchitectural state (e.g., cache updates) is not restored.

While many past attacks have exploited microarchitectural side channels to extract confidential data [14–20], the difference with transient execution vulnerabilities is that the latter only manifest due to misspeculation in the processor. *Even programs whose architectural (non-speculative) execution is carefully designed to not have any side-channel leaks could be vulnerable to transient execution attacks.*

(S4) *Exfiltration*: Finally, control returns to the attacker who examines microarchitectural side-channel state to exfiltrate confidential data from the victim. In cache-based attacks, this involves *probing* the cache in order to infer secrets.

### B. Spectre Variants and Associated Verification Challenges

We now describe the Spectre variant 1 vulnerability and a few modifications to it as exemplars of transient execution attacks. We use this discussion to motivate the research challenges posed by transient execution attacks. While we focus on Spectre variant 1 for ease of exposition, the research questions raised here apply to all other transient execution attacks.

We discuss the four code snippets shown in Figure 3. In each of the snippets, the vulnerable victim function is `foo`. This function is trusted but is invoked by an untrusted attacker with an arbitrary attacker chosen argument `i`. `foo` has access to two arrays: `a1` and `a2`. Note that any architectural execution of these functions should never see accesses to `a1[i]` for  $i \geq N$ . Therefore, one might expect that no information could possibly leak about these values in the array through any side-channel. As we will see, the Spectre attack shows how these values can be inferred by a clever attacker.

1) *Spectre Variant 1*: Figure 3a shows a snippet of code that demonstrates vulnerability to the Spectre variant 1 attack [2,

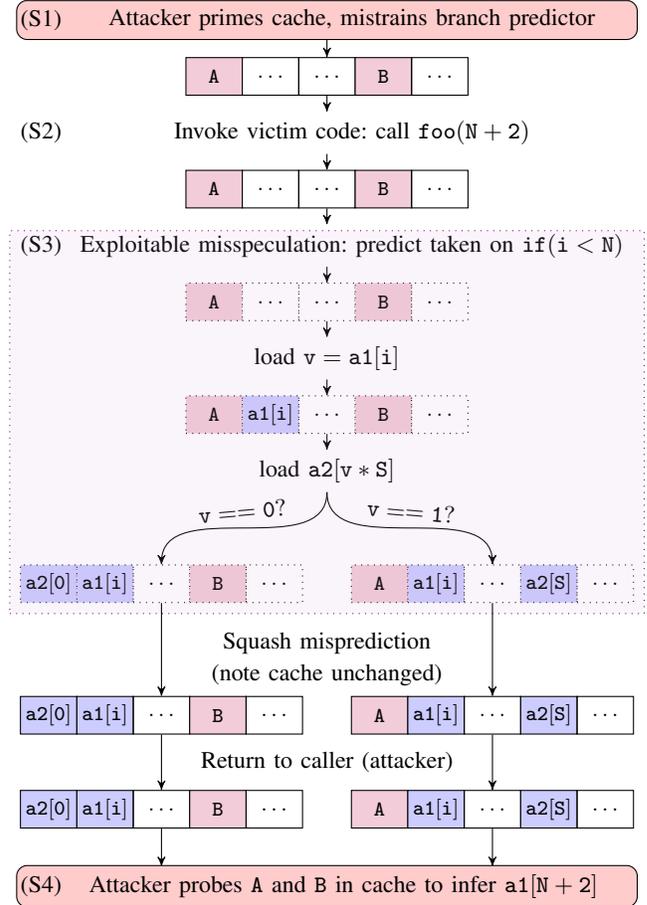


Figure 2: Cache state evolution in Spectre variant 1. The rectangular boxes show the addresses that are cached. Untrusted accesses are red while accesses by trusted code are blue. For simplicity, we show the attack on a direct-mapped cache. The extension to set associative caches is straightforward.

3]. To help explain the vulnerability, we show how cache state evolves during each stage of the attack in Figure 2.

S1 First, the attacker sets up (“primes”) the cache by bringing two addresses A and B into the cache. These addresses are carefully chosen so as to reside in the same cache set as the subsequently-accessed addresses `a2[0]` and `a2[S]` respectively. Next, the attack mistrains the predictor to speculate that the branch on line 4 will be taken. Now, the attack is ready to be launched.

S2 The attacker invokes `foo` with an argument  $i = N + 2$ .

S3 The argument  $i = N + 2$  along with branch predictor mistraining in S1 triggers a misspeculation on line 4. This results in `a1[N+2]` and `a2[a1[N+2]*S]` being speculatively brought into the cache. Eventually, the processor realizes that the branch prediction was incorrect and “undoes” modifications to architectural state, but cache state is *not* restored.

S4 In the final stage, the attacker exploits the fact that the address brought into the cache on line 6 depends on the

```

1 uint8_t a1[M];
2 uint8_t a2[P];
3 uint8_t foo(unsigned i) {
4     if (i < N) {
5         uint8_t v = a1[i];
6         return a2[v*S];
7     }
8     return 0;
9 }

```

(a) Spectre v1 vulnerability.

```

1 uint8_t a1[M];
2 uint8_t a2[P];
3 uint8_t foo(unsigned i) {
4     if (i < N) {
5         __mm_lfence();
6         return a2[a1[i]*S];
7     }
8     return 0;
9 }

```

(b) Fix for Spectre v1.

```

1 uint8_t a1[M];
2 uint8_t a2[P];
3 uint8_t foo(unsigned i) {
4     if (i < N) {
5         uint8_t v = a1[0];
6         return a2[v*S]+i;
7     }
8     return 0;
9 }

```

(c) Conditionally vulnerable variant.

```

1 uint8_t a1[M], a2[P];
2 uint8_t foo(unsigned i) {
3     if (i < N) {
4         uint8_t v = a1[i];
5         __mm_lfence();
6         return a2[v*S];
7     }
8     return 0;
9 }

```

(d) Another fix for Spectre v1.

Figure 3: Illustrative examples for verification of secure speculation. In all code snippets assume that  $M > N$  and that argument  $i$  is an untrusted (low-security) input to the trusted (high-security) function  $f_{\square\square}$ .

value (not address) of  $a1[N+2]$ . The attacker determines this address by loading A and B. One of these will miss in the cache and this timing channel allows the attacker to infer the value of  $a1[N+2]$ .

2) *Fixes to Spectre Variant 1*: As the leaks in Spectre are due to interactions between the branch predictor and the cache, a straightforward fix is to prevent speculation. We can make the code in Figure 3a secure by inserting a *load fence* [9, 22] as shown in Figures 3b and 3d. Figure 3b is easy to understand: the load fence on line 5 ensures that no memory accesses are made until the processor is sure that branch will be taken.

Figure 3d is slightly more involved. The load fence executes after the first load and before the second load. At first glance, it may appear to be insecure, because  $a1[i]$  can still be brought into the cache speculatively. However  $i$  is attacker-chosen while the base address of  $a1$  can also be inferred by the attacker. Therefore, bringing  $a1[i]$  into the cache leaks no additional information. Figure 3d is secure.

3) *Conditional Vulnerability*: Figure 3c presents an interesting variation of Figure 3a. In this case, the first memory load always accesses  $a1[0]$ . Since this value is leaked through the cache (when  $i < N$ ) even without misspeculation, it would seem that this code is not vulnerable to transient execution attacks. However, if  $N = 0$ , then  $a1[0]$  should not be accessed. But the attacker can mistrain the branch predictor to predict that the branch on line 4 is taken and then infer the value of  $a1[0]$ . This code exhibits transient execution vulnerabilities when  $N = 0$  but not when  $N > 0$ !

4) *Verification Challenges*: In Figure 3a, information about  $a1[i]$  leaked when  $i = N + 2$ . For this value of  $i$ ,  $f_{\square\square}$  should not have performed any memory/cache accesses. This points to one challenge in verifying secure speculation:

the verification model needs to capture *interactions* between *microarchitectural side-channels* to detect leaks.

Another challenge is demonstrated by Figures 3b, 3c and 3d. Identifying the vulnerability requires precise semantic analysis of program behavior. Simply matching vulnerable code patterns (e.g., branches followed by dependent loads) results in both false positives and negatives.

Finally, it is important to note that the secure versions  $f_{\square\square}$  in Figures 3b and 3d do not satisfy traditional notions of information flow security [40]: noninterference [41] or observational determinism [42–44] because there *is* information flow from  $a1$  to the cache side-channel even if the function is executed on a processor without a branch predictor.

Current architectural efforts at secure speculation [37] use a “big-hammer” approach of eliminating all cache side-channel leaks – speculative or not – via cache partitioning. While this does ensure security against transient execution attacks when exfiltration is via the cache side-channel, we argue that the approach is sub-optimal for the following reasons. First, there are many side-channels in modern processors besides caches such as prefetchers [50, 51], TLBs [52] and DRAM [53]. These side channels may be vulnerable to transient execution as they are not partitioned. Second, protection mechanisms that rule out all side-channel leaks via partitioning have a higher performance penalty than targeted defences. Given that other side channels are reasonably well-understood and many mitigation/verification techniques exist [54–64], there is reason to believe that a targeted solution will suffice. Finally and perhaps most importantly, this approach gives us no insight into the novel class of transient execution vulnerabilities that Spectre has brought to light.

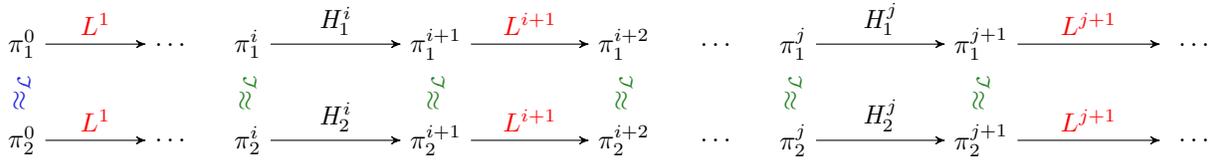


Figure 4: Illustrating observational determinism: low instructions are labelled  $L$ , while high instructions are labelled  $H_1$  and  $H_2$ , proof obligations are shown in green and assumptions are shown in blue.

### III. SPECIFICATION USING TRACE PROPERTY-DEPENDENT OBSERVATIONAL DETERMINISM

To address the challenges raised in § II-B4, this paper formulates a secure speculation property that precisely captures transient execution vulnerabilities. Toward this end, in this section we first review observational determinism [42–44], a class of security properties that can capture certain notions of confidentiality. We then motivate and describe trace property-dependent observational determinism: a novel class of information flow properties that includes secure speculation.

#### A. Preliminaries

We model system behavior using traces which are a sequence of system states. The definition of system states is left abstract for now. We refer to traces using  $\pi$ ,  $\pi_1$ ,  $\pi_2$ , etc. and states by  $s$ ,  $s_0$ ,  $s_1$ , etc. The notation  $\pi^i$  refers the  $i$ th element of the trace; e.g., if  $\pi = \langle s_0, s_1, s_2, s_3, s_4, \dots \rangle$ , then  $\pi^3 = s_3$ .

We consider concurrent systems consisting of two components: an untrusted low-security component and a trusted high-security component. These components interact via some interface (e.g., system calls and returns) which prompt transitions from the low component to the high component or vice versa. A typical confidentiality requirement is that the low component must not be able to distinguish between secret states of the high component.

1) *Low-Equivalence of States*: The above notion of indistinguishability is expressed via low-equivalence of states. We say that two states  $s$  and  $s'$  are low-equivalent if they are indistinguishable to the low component. This is denoted by  $s \approx_{\mathcal{L}} s'$ . Like system states, the definition of low-equivalence is left abstract for now. Low-equivalence is extended to traces in the obvious way. Two traces are low-equivalent if all their states are low-equivalent:  $\pi_1 \approx_{\mathcal{L}} \pi_2$  if  $\forall i. \pi_1^i \approx_{\mathcal{L}} \pi_2^i$ .

2) *Modeling Computation*: The system computes by identifying an operation to execute and transitioning to the next state based on its transition relation  $\rightsquigarrow$ . If system state  $s_i$  can transition to state  $s_j$ , then  $(s_i, s_j) \in \rightsquigarrow$ , which we write as  $s_i \rightsquigarrow s_j$ . As with states, we leave  $\rightsquigarrow$  abstract for now.

The operation executed by the low component in a particular state  $s$  is denoted by  $op_{\mathcal{L}}(s)$ ;  $op_{\mathcal{L}}(s)$  is  $\perp$  if the low component is not being executed in state  $s$ . Similarly, the operation executed by the high component in the state  $s$  is denoted by  $op_{\mathcal{H}}(s)$ . We will overload notation and refer to  $op_{\mathcal{L}}(\pi)$  and  $op_{\mathcal{H}}(\pi)$  to denote the trace of operations executed by the low and high components respectively in  $\pi$ .

#### B. Observational Determinism

A system satisfies observational determinism if for every pair of traces of the system such that: (i) the two traces' initial states are low-equivalent, and (ii) the low operations executed at every step of the two traces are identical, then the two traces are also low-equivalent. Equation 1 shows this definition.

$$\forall \pi_1, \pi_2. \quad (\pi_1^0 \approx_{\mathcal{L}} \pi_2^0 \wedge op_{\mathcal{L}}(\pi_1) = op_{\mathcal{L}}(\pi_2)) \implies (\pi_1 \approx_{\mathcal{L}} \pi_2) \quad (1)$$

Observational determinism is shown pictorially in Figure 4. The figure shows a pair of traces with their initial states being low-equivalent. In subsequent steps, the low operations are identical and this is denoted by labelling low transitions as  $L$ . However, high operations may differ between the traces, so we label its transitions as  $H_1$  and  $H_2$ . Observational determinism holds if every corresponding pair of states in these two traces are low-equivalent. A violation of observational determinism is some sequence of low operations that can distinguish between some two secret states of the high component.

#### C. Trace Property-Dependent Observational Determinism

In a processor that never misspeculates – either because it does not have a branch predictor or because the branch predictor is perfect – there is no information leakage due to transient execution. Therefore, finding transient execution vulnerabilities is equivalent to finding information leaks that would not have occurred in the absence of misspeculation.

1) *Definition of TPOD*: To formulate the above notion of information leakage, we introduce a class of information flow properties called trace property-dependent observational determinism (TPOD), a hyperproperty over four traces that is defined with respect to a trace property [65]. Let the four traces be  $\pi_1, \pi_2, \pi_3, \pi_4$ , and the trace property  $T$ .

Suppose the following assumptions hold:

- 1) traces  $\pi_1$  and  $\pi_2$  satisfy the trace property  $T$ ,
- 2) traces  $\pi_3$  and  $\pi_4$  do not satisfy the trace property  $T$ ,
- 3) all four traces execute the same low operations,
- 4) traces  $\pi_3$  and  $\pi_4$  execute execute the same high operations as  $\pi_1$  and  $\pi_2$  respectively,
- 5) traces  $\pi_1$  and  $\pi_2$  are low-equivalent and the initial states of  $\pi_3$  and  $\pi_4$  are low-equivalent.

Then, TPOD is satisfied if  $\pi_3$  and  $\pi_4$  are low-equivalent. High operations in  $\pi_1, \pi_3$  and  $\pi_2, \pi_4$  respectively must be identical; they are not necessarily identical in  $\pi_1, \pi_2$  or  $\pi_3, \pi_4$ .

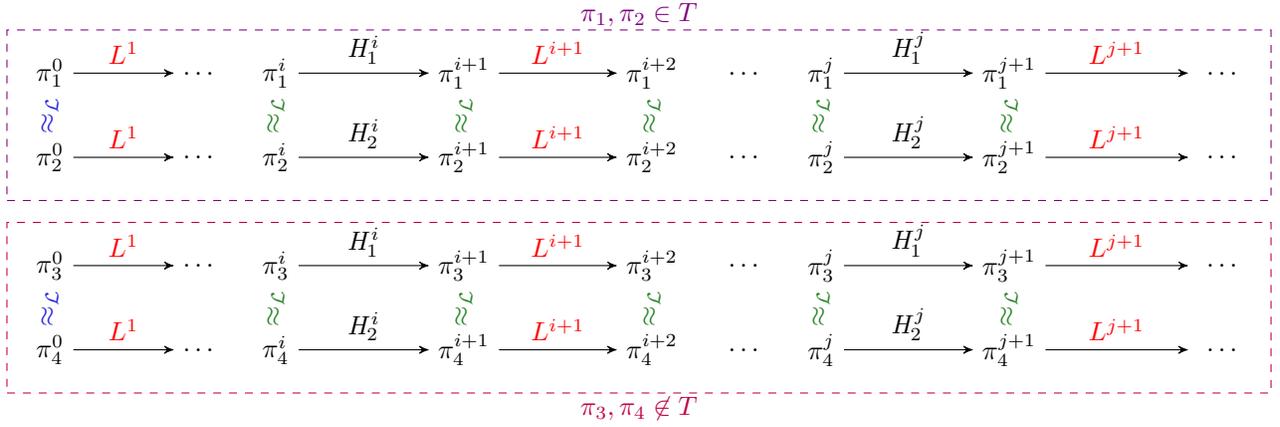


Figure 5: Illustrating trace property-dependent observational determinism.

$\forall \pi_1, \pi_2, \pi_3, \pi_4.$

$$\begin{aligned}
& \pi_1 \in T \wedge \pi_2 \in T \wedge \pi_3 \notin T \wedge \pi_4 \notin T & \implies \\
& op_{\mathcal{L}}(\pi_1) = op_{\mathcal{L}}(\pi_2) = op_{\mathcal{L}}(\pi_3) = op_{\mathcal{L}}(\pi_4) & \implies \\
& op_{\mathcal{H}}(\pi_1) = op_{\mathcal{H}}(\pi_3) \wedge op_{\mathcal{H}}(\pi_2) = op_{\mathcal{H}}(\pi_4) & \implies \\
& \pi_1 \approx_{\mathcal{L}} \pi_2 \wedge \pi_3^0 \approx_{\mathcal{L}} \pi_4^0 & \implies \\
& \pi_3 \approx_{\mathcal{L}} \pi_4 & (2)
\end{aligned}$$

TPOD is shown in Equation 2 and depicted in Figure 5.<sup>1</sup> A violation of TPOD corresponds to a sequence of low operations that were unable to distinguish between high states when the trace property  $T$  was satisfied, but *are* able to distinguish between high states when  $T$  is not satisfied. In other words, violation of the trace property  $T$  introduced a new counterexample to observational determinism.

2) *Refinement and TPOD*: In general, hyperproperties may not be preserved by refinement [43]. However, as we show below TPOD is subset-closed: if any set of traces satisfies TPOD, then every subset of this set also satisfies TPOD.

*Lemma 1*: Trace property-dependent observational determinism is a subset-closed hyperproperty.

Subset-closed hyperproperties are important because they are preserved by refinement [40]. This means that one can prove TPOD on an abstract system, and through iterative refinement show that TPOD holds on a concrete system that is a refinement of the abstract system. Therefore, TPOD can potentially be scalably verified on complex systems.

*Corollary 1*: Trace property-dependent observational determinism is preserved by refinement.

A minor extension to template shown in Equation 2 is to consider an antecedent trace property  $U$  that must be satisfied by all traces. The trace property  $U$  may be used model constraints on valid executions.

$\forall \pi_1, \pi_2, \pi_3, \pi_4.$

$$\begin{aligned}
& \pi_1 \in U \wedge \pi_2 \in U \wedge \pi_3 \in U \wedge \pi_4 \in U & \implies \\
& \pi_1 \in T \wedge \pi_2 \in T \wedge \pi_3 \notin T \wedge \pi_4 \notin T & \implies \\
& op_{\mathcal{L}}(\pi_1) = op_{\mathcal{L}}(\pi_2) = op_{\mathcal{L}}(\pi_3) = op_{\mathcal{L}}(\pi_4) & \implies \\
& op_{\mathcal{H}}(\pi_1) = op_{\mathcal{H}}(\pi_3) \wedge op_{\mathcal{H}}(\pi_2) = op_{\mathcal{H}}(\pi_4) & \implies \\
& \pi_1 \approx_{\mathcal{L}} \pi_2 \wedge \pi_3^0 \approx_{\mathcal{L}} \pi_4^0 & \implies \\
& \pi_3 \approx_{\mathcal{L}} \pi_4 & (3)
\end{aligned}$$

This version of TPOD is shown in Equation 3. This extension is also subset-closed and preserved by refinement

#### IV. SECURE SPECULATION

Having defined TPOD, we now turn to the problem of specifying secure speculation. This requires formalizing models of an adversary and the system for which this property is to hold. Therefore, we first describe an assembly intermediate representation and introduce speculative operational semantics for it. We then introduce our adversary model and present the secure speculation property for this adversary. Violations of the property correspond to transient execution vulnerabilities.

##### A. System Model

Reasoning about secure speculation must be done using assembly language instructions, not in a high-level language because compiler optimizations may introduce branches where none exist in the source program, or may eliminate branches in the source program by turning them into conditional moves. That said, reasoning about a specific instruction set architecture (ISA) is cumbersome and gives little additional insight into the fundamental causes of transient execution vulnerabilities. Therefore, we present an assembly intermediate representation (AIR) that ISAs can be *lifted* into. We model speculation over the AIR by introducing a speculative operational semantics for it.<sup>2</sup>

<sup>2</sup>The AIR itself is based on the binary analysis platform (BAP) intermediate language (IL) [66]; speculative operational semantics for it are novel. “Lifters” from x86 and ARM binaries to BAP can be found at [67].

<sup>1</sup>We follow the convention that the implication operator is right-associative.

$$\begin{aligned}
\langle \text{program} \rangle &::= \langle \text{instr} \rangle^* \\
\langle \text{instr} \rangle &::= \langle \text{reg} \rangle := \langle \text{exp} \rangle \\
&| \langle \text{reg} \rangle := \text{mem}[\langle \text{exp} \rangle] \\
&| \text{mem} := \text{mem}[\langle \text{exp} \rangle \rightarrow \langle \text{exp} \rangle] \\
&| \text{if } \langle \text{exp} \rangle \text{ goto } \langle \text{const} \rangle \\
&| \text{goto } \langle \text{const} \rangle \\
&| \text{specfence} \\
\langle \text{exp} \rangle &::= \langle \text{const} \rangle | \langle \text{reg} \rangle | \diamond_u \langle \text{exp} \rangle | \langle \text{exp} \rangle \diamond_b \langle \text{exp} \rangle
\end{aligned}$$

Figure 6: The Assembly Intermediate Representation (AIR).  $\diamond_u$  and  $\diamond_b$  are typical unary and binary operators respectively.

1) *Assembly Intermediate Representation (AIR)*: The AIR shown in Figure 6. A program is a list of instructions. Instructions are one of the following types:

- updates to registers,
- loads from memory,
- stores to memory,
- conditional and unconditional jumps,
- speculation fences.

The first five types of instructions are standard. We introduce a *speculation fence* instruction which causes the processor to not fetch any more instructions until all outstanding branches are resolved. The load fence instructions in Figures 3b and 3d are modelled as speculation fences because the relevant aspect of these fences for this paper is that they stop speculation.

Note that jump targets must be constants in AIR. This is intentional and precludes the verification of programs using indirect jumps and returns in the current version of our verification tool. We do this to simplify the operational semantics for speculative execution. Modeling speculative execution of indirect jumps requires modeling indirect branch predictors and branch target buffers. Return instructions are predicted using a return address stack. Introducing these structures into our operation semantics is conceptually straightforward but runs into scalability limitations during verification. We plan to extend the operational semantics to include these instructions while addressing scalability in future work.

$$\begin{array}{c}
\text{CONST} \frac{}{\Delta, n \vdash c \Downarrow c} \quad \text{REG} \frac{\Delta[n, r] = v}{\Delta, n \vdash r \Downarrow v} \\
\text{UNOP} \frac{\Delta, n \vdash e \Downarrow v' \quad \diamond_u v' = v}{\Delta, n \vdash \diamond_u e \Downarrow v} \\
\text{BINOP} \frac{\Delta, n \vdash e_1 \Downarrow v_1 \quad \Delta, n \vdash e_2 \Downarrow v_2 \quad v_1 \diamond_b v_2 = v}{\Delta, n \vdash e_1 \diamond_b e_2 \Downarrow v}
\end{array}$$

Figure 7: Semantics of expression evaluation

“Flattening” indirect jumps and returns into a sequence of direct jumps is similar in principle to control-flow integrity (CFI) checks [68, 69]. Since secure programs will likely be implementing CFI anyway, we assert compilers can be modified in straightforward ways to produce code without indirect jumps and returns (with some performance cost).

2) *Operational Semantics for AIR*: In Figures 7 and 8, we introduce operational semantics for **speculative in-order** processors. We model speculation in the branch predictor for direct conditional branches. Other sources of misspeculation such as value prediction and memory address disambiguation are not considered in this model. Extending the semantics to include these is conceptually straightforward. However, this could result in models that are difficult to analyze using automated verification tools because the verification engine would need to explore exponentially more instruction orderings.

**Machine state**  $s$  is the tuple  $\langle \Pi, \Delta, \mu, pc, \omega, \beta, n, \iota \rangle$ .  $\Pi$  is the program memory: a map from program counter values to instructions.  $\Delta$  and  $\mu$  are the state of the registers and data memory respectively while  $pc$  contains the program counter.  $\omega$  is the trace of program and data addresses accessed so far.  $\beta$  is the branch predictor state, which we leave abstract in this paper and  $\iota$  is the instruction that will be executed next.

The main novelty in these semantics is modeling misspeculation.  $n$  is an integer that represents *speculation level*: it is incremented each time we misspeculate on a branch and decremented when a branch is resolved. Speculation level 0 corresponds to architectural (non-speculative) execution.  $\Delta$ ,  $\mu$  and  $pc$  – registers, memory and program counter respectively – are also indexed by the speculation level.  $\Delta[n, r]$  refers to the value of the register  $r$  at speculation level  $n$ .  $\Delta[(n, r) \rightarrow v]$  refers to a register state which is identical to  $\Delta$  except that register  $r$  at speculation level  $n$  has been assigned value  $v$ . We adopt similar notation for  $\mu$  and  $pc$ .

**Expression Semantics** are shown in Figure 7. Expressions are defined over the register state  $\Delta$ . Notation  $\Delta, n \vdash e \Downarrow v$  means that the expression  $e$  evaluates to value  $v$  given register state  $\Delta$  at speculation level  $n$ . These are standard except for the additional wrinkle of the speculation level.

**Statement Semantics** are shown in Figure 8. A transition from the machine state  $s = \langle \Pi, \Delta, \mu, pc, \omega, \beta, n, \iota \rangle$  to the machine state  $s' = \langle \Pi', \Delta', \mu', pc', \omega', \beta', n', \iota' \rangle$  is written as  $\langle \Pi, \Delta, \mu, pc, \omega, \beta, n, \iota \rangle \rightsquigarrow \langle \Pi', \Delta', \mu', pc', \omega', \beta', n', \iota' \rangle$ . We now briefly describe the rules shown in Figure 8.

The REGISTERUPDATE rule models the execution of statements of the form  $r := e$ , expression  $e$  is written to the register  $r$ . This involves: (i) updating the value of register  $r$  at speculation level  $n$  to have the value of the expression  $e$ :  $\Delta' = \Delta[(n, r) \rightarrow v]$ , (ii) incrementing the  $pc$  at speculation level  $n$ :  $pc' = pc[n \rightarrow \rho + 1]$  and (iii) appending  $\langle \rho, \perp \rangle$  to the trace of memory addresses accessed by the program:  $\omega' = \omega.\langle pc, \perp \rangle$ . The  $\perp$  in the second element of the tuple indicates that no data memory access is performed by this instruction. This rule is only executed when a branch is not being resolved:  $\neg \text{resolve}(n, \beta, pc)$  and the next instruction to be executed is  $\iota' = \Pi[\rho']$  where  $\rho' = pc[n]$ .

The LOAD and STORE rules are similar. LOAD updates the register state with value stored at memory location  $a$  at speculation level  $n$ :  $v = \mu[n, a]$  while STORE leaves register state  $\Delta$  unchanged and updates memory address  $a$  at speculation level  $n$ :  $\mu' = \mu[(n, a) \rightarrow v]$ . Both LOAD and STORE append  $\langle \rho, a \rangle$  to the trace of memory addresses

$$\begin{array}{c}
\text{REGISTERUPDATE} \frac{\Delta, n \vdash e \Downarrow v \quad \Delta' = \Delta[(n, r) \rightarrow v] \quad \rho = pc[n] \quad \iota = \Pi[\rho] \quad pc' = pc[n \rightarrow \rho + 1]}{\rho' = pc'[n] \quad \iota' = \Pi[\rho'] \quad \omega' = \omega.\langle \rho, \perp \rangle \quad \neg \text{resolve}(n, \beta, pc)}{\Pi, \Delta, \mu, pc, \omega, \beta, n, r := e \rightsquigarrow \Pi, \Delta', \mu, pc', \omega', \beta, n, \iota'} \\
\\
\text{LOAD} \frac{\Delta, n \vdash e \Downarrow a \quad \mu[n, a] = v \quad \Delta' = \Delta[(n, r) \rightarrow v] \quad \rho = pc[n] \quad \iota = \Pi[\rho] \quad pc' = pc[n \rightarrow \rho + 1]}{\rho' = pc'[n] \quad \iota' = \Pi[\rho'] \quad \omega' = \omega.\langle \rho, a \rangle \quad \neg \text{resolve}(n, \beta, pc)}{\Pi, \Delta, \mu, pc, \omega, \beta, n, r := \text{mem}[e] \rightsquigarrow \Pi, \Delta', \mu, pc', \omega', \beta, n, \iota'} \\
\\
\text{STORE} \frac{\Delta, n \vdash e_1 \Downarrow a \quad \Delta, n \vdash e_2 \Downarrow v \quad \mu' = \mu[(n, a) \rightarrow v] \quad \rho = pc[n] \quad \iota = \Pi[\rho] \quad pc' = pc[n \rightarrow \rho + 1]}{\rho' = pc'[n] \quad \iota' = \Pi[\rho'] \quad \omega' = \omega.\langle \rho, a \rangle \quad \neg \text{resolve}(n, \beta, pc)}{\Pi, \Delta, \mu, pc, \omega, \beta, n, \text{mem} := \text{mem}[e_1 \rightarrow e_2] \rightsquigarrow \Pi, \Delta, \mu', pc', \omega', \beta, n, \iota'} \\
\\
\text{T-PRED} \frac{\Delta, n \vdash e \Downarrow \text{true} \quad \text{mispred}(n, \beta, pc) = \text{false} \quad \rho = pc[n] \quad \iota = \Pi[\rho] \quad pc' = pc[n \rightarrow c]}{\rho' = pc'[n] \quad \iota' = \Pi[\rho'] \quad \omega' = \omega.\langle \rho, \perp \rangle \quad \beta' = \text{update}(n, \rho, \iota, \beta) \quad \neg \text{resolve}(n, \beta, pc)}{\Pi, \Delta, \mu, pc, \omega, \beta, n, \text{if } e \text{ goto } c \rightsquigarrow \Pi, \Delta, \mu, pc', \omega', \beta', n, \iota'} \\
\\
\text{T-MISPRED} \frac{\Delta, n \vdash e \Downarrow \text{true} \quad \text{mispred}(n, \beta, pc) = \text{true} \quad \rho = pc[n] \quad \iota = \Pi[\rho] \quad n' = n + 1}{\forall m, r. \Delta'[m, r] = \text{ITE}(m = n', \Delta(n, r), \Delta(m, r)) \quad \forall m, a. \mu'[m, r] = \text{ITE}(m = n', \mu(n, r), \Delta(m, r))}{pc' = pc[n' \rightarrow \rho + 1, n \rightarrow c] \quad \rho' = pc'[n'] \quad \iota' = \Pi[\rho'] \quad \beta' = \text{update}(n, \rho, \iota, \beta) \quad \neg \text{resolve}(n, \beta, pc)}{\Pi, \Delta, \mu, pc, \omega, \beta, n, \text{if } e \text{ goto } c \rightsquigarrow \Pi, \Delta', \mu', pc', \omega', \beta', n', \iota'} \\
\\
\text{NT-PRED} \frac{\Delta, n \vdash e \Downarrow \text{false} \quad \text{mispred}(n, \beta, pc) = \text{false} \quad \rho = pc[n] \quad \iota = \Pi[\rho] \quad pc' = pc[n \rightarrow \rho + 1]}{\rho' = pc'[n] \quad \iota' = \Pi[\rho'] \quad \omega' = \omega.\langle \rho, \perp \rangle \quad \beta' = \text{update}(n, \rho, \iota, \beta) \quad \neg \text{resolve}(n, \beta, pc)}{\Pi, \Delta, \mu, pc, \omega, \beta, n, \text{if } e \text{ goto } c \rightsquigarrow \Pi, \Delta, \mu, pc', \omega', \beta', n, \iota'} \\
\\
\text{NT-MISPRED} \frac{\Delta, n \vdash e \Downarrow \text{false} \quad \text{mispred}(n, \beta, pc) = \text{true} \quad \rho = pc[n] \quad \iota = \Pi[\rho] \quad n' = n + 1}{\forall m, r. \Delta'[m, r] = \text{ITE}(m = n', \Delta(n, r), \Delta(m, r)) \quad \forall m, a. \mu'[m, r] = \text{ITE}(m = n', \mu(n, r), \Delta(m, r))}{pc' = pc[n' \rightarrow c, n \rightarrow \rho + 1] \quad \rho' = pc'[n'] \quad \iota' = \Pi[\rho'] \quad \beta' = \text{update}(n, \rho, \iota, \beta) \quad \neg \text{resolve}(n, \beta, pc)}{\Pi, \Delta, \mu, pc, \omega, \beta, n, \text{if } e \text{ goto } c \rightsquigarrow \Pi, \Delta', \mu', pc', \omega', \beta', n', \iota'} \\
\\
\text{GOTO} \frac{\rho = pc[n] \quad \iota = \Pi[\rho] \quad pc' = pc[n \rightarrow c] \quad \rho' = pc'[n]}{\iota' = \Pi[\rho'] \quad \omega' = \omega.\langle \rho, \perp \rangle \quad \beta' = \text{update}(n, \rho, \iota, \beta) \quad \neg \text{resolve}(n, \beta, pc)}{\Pi, \Delta, \mu, pc, \omega, \beta, n, \text{goto } c \rightsquigarrow \Pi, \Delta, \mu, pc', \omega', \beta', n, \iota'} \\
\\
\text{SPECFENCE} \frac{n' = 0 \quad \rho' = pc[n'] \quad \iota' = \Pi[\rho'] \quad \neg \text{resolve}(n, \beta, pc)}{\Pi, \Delta, \mu, pc, \omega, \beta, n, \iota \rightsquigarrow \Pi, \Delta, \mu, pc, \omega, \beta, n', \iota'} \\
\\
\text{RESOLVE} \frac{n' = n - 1 \quad \rho' = pc[n'] \quad \iota' = \Pi[\rho'] \quad \beta' = \text{update}(n, \rho, \iota, \beta) \quad \text{resolve}(n, \beta, pc)}{\Pi, \Delta, \mu, pc, \omega, \beta, n, \iota \rightsquigarrow \Pi, \Delta, \mu, pc, \omega, \beta', n', \iota'} \\
\\
\text{HAVOC} \frac{n = 0 \quad \rho \notin \mathcal{T}_\rho \quad pc'[n] = \rho \quad \iota = \Pi[\rho] \quad \forall a. a \notin \mathcal{U}_\mu^{wr} \implies \mu'[0, a] = \mu[0, a]}{\Pi, \Delta, \mu, pc, \omega, \beta, n, \text{havoc } (\Delta, \text{mem}[\mathcal{U}_\mu^{wr}], \beta) \rightsquigarrow \Pi, \Delta', \mu', pc', \omega, \beta', n, \iota}
\end{array}$$

Figure 8: Operational Semantics for Statements in AIR.

signifying accesses to program address  $\rho$  and data address  $a$ . As with REGISTERUPDATE, these rules only apply when a branch is not being resolved in this step:  $\neg \text{resolve}(n, \beta, pc)$ .

The T-PRED rule applies when a conditional jump `if  $e$  goto  $c$`  should be taken and is also predicted taken. In the semantics, we model misspeculation through an uninterpreted function  $\text{mispred}(n, \beta, pc)$  where  $\beta$  is the branch predictor state (left abstract in our model),  $n$  is the speculation level and  $pc$  is a map from speculation levels to program counter values. This rule only applies when  $\text{mispred}$  evaluates false. The rule sets the program counter at speculation level  $n$  to  $c$ :

$pc' = pc[n \rightarrow c]$  and updates the branch predictor state  $\beta'$  using the uninterpreted function  $\text{update}$ . Just like the other rules discussed so far, this applies only when the predicate  $\text{resolve}$  does not hold.

The T-MISPRED rule applies when a conditional jump `if  $e$  goto  $c$`  should be taken but is predicted not taken ( $\text{mispred}$  evaluates to true). This rule changes system state in the following ways. First, the speculation level is incremented:  $n' = n + 1$ . Second, the state of the registers at level  $n$  in  $\Delta$  is now copied over to level  $n'$  in  $\Delta'$  while all other levels are identical between  $\Delta$  and  $\Delta'$ . The memory state  $\mu$  is also

modified in a similar way. The program counter at level  $n$  gets the correct target  $c$ , while the program counter at level  $n'$  gets the mispredicted fall-through target  $\rho+1$ . Execution continues at speculation level  $n'$ .

NT-PRED, NT-MISPRED handle the case when the conditional branch should *not* be taken. These are similar to T-PRED and T-MISPRED. GOTO applies to direct jumps. Note we do not consider misprediction of direct jumps as they have constant targets and will be redirected at decode.

The rule SPECFENCE resolves all outstanding speculative branches by setting the speculation level back to zero. Note that  $pc$ ,  $\Delta$  and  $\mu$  at level zero already have the “correct” values, so nothing further needs to be done.

The rule RESOLVE applies when a mispredicted branch is resolved. Resolution occurs when the uninterpreted predicate  $resolve(n, \beta, pc)$  holds. At the time of resolution, branch predictor state  $\beta'$  is updated using the uninterpreted function  $update$  and the speculation level  $n'$  is decremented. As in SPECFENCE, nothing else need be done as the other state variables have the correct values at the decremented level.

Rule HAVOC will be described in § IV-C1.

### B. Adversary Model

Recall system state  $s$  is the tuple  $\langle \Pi, \Delta, \mu, pc, \omega, \beta, n, \iota \rangle^3$  and evolves according to the transition relation  $\rightsquigarrow$  from Figure 8. As discussed in § III, the system has an untrusted low-security component and a trusted high-security component that execute concurrently. Our verification objective is to prove that confidential states of a specified trusted program are indistinguishable to an arbitrary untrusted program. This verification task requires the definition of: (i) the trusted program to be verified and the family of untrusted adversary programs, (ii) confidential states of the trusted program, (iii) how the adversary *tampers* with system state, and (iv) what parts of state are adversary *observable*.

1) *The Trusted and Untrusted Programs*: We assume that the trusted program resides in the set of instruction memory addresses denoted by  $\mathcal{T}_\rho$ . The trusted program itself is defined by  $\Pi[\rho]$  for each  $\rho \in \mathcal{T}_\rho$ . We are verifying specific programs and not a microarchitecture, so verification has to be redone for modification to a trusted program. Every address  $\rho \notin \mathcal{T}_\rho$  is part of the untrusted component and  $\Pi[\rho]$  is unconstrained for these addresses to model all possible adversarial programs. We assume that untrusted code can invoke trusted code only by jumping to a specific entrypoint address  $\mathcal{EP} \in \mathcal{T}_\rho$ .  $\mathcal{T}_\rho$ ,  $\mathcal{EP}$  and instructions  $\Pi[\rho]$  for all  $\rho \in \mathcal{T}_\rho$  are known to the adversary.

For the motivating examples shown in Figure 3a,  $\mathcal{T}_\rho$  contains all instruction addresses that are part of the function  $\text{f00}$ . The entrypoint  $\mathcal{EP}$  is the address of the first instruction in  $\text{f00}$ . If we are verifying secure speculation for system calls in an operating system kernel,  $\mathcal{T}_\rho$  contains all kernel text addresses and the entrypoint  $\mathcal{EP}$  is the syscall trap address.

Given the above definitions, the low operation executed in a state  $s = \langle \Pi, \Delta, \mu, pc, \omega, \beta, n, \iota \rangle$  is  $op_{\mathcal{L}}(s) \doteq \Pi[pc[0]]$  if

$pc[0] \notin \mathcal{T}_\rho$  and  $\perp$  otherwise. This definition refers to the non-speculative state – we are looking at  $pc[0]$ , not higher speculation levels. The instruction being speculatively executed may be different, and may in fact be from the trusted component. This is important because we use  $op_{\mathcal{L}}$  to constrain adversary actions to be identical across traces, and these constraints can only refer to non-speculative state.

Finally, the trusted program must start off in some well-defined initial state. For instance, global variables may need to be initialized to specific values. We use the predicate  $init_{\mathcal{T}}(s)$  to refer to a valid initial state of the trusted program.

2) *Confidential States*: The secret states that need to be protected from an adversary are the values stored in memory addresses  $a$  that belong to the set  $\mathcal{S}_{\mathcal{T}}$ . For Figure 3,  $\mathcal{S}_{\mathcal{T}}$  contains all addresses that are part of the arrays  $a1$  and  $a2$ .

All other addresses are public state. We will use  $\mathcal{P}_{\mathcal{T}}$  to denote the projection of the values stored at these public addresses:  $\mathcal{P}_{\mathcal{T}}(\mu) \doteq \lambda a. \text{ITE}(a \notin \mathcal{S}_{\mathcal{T}}, \mu[0, a], \perp)$ .

The high instruction executed in a state is denoted  $inst_{\mathcal{T}}(s)$  and has the value  $s.\Pi[s.pc[0]]$  when  $pc[0] \in \mathcal{T}_\rho$  and  $\perp$  otherwise. The high operation executed in state  $s$  is defined as a tuple of the high instruction and the public memory:  $op_{\mathcal{H}}(s) \doteq \langle inst_{\mathcal{T}}(s), \mathcal{P}_{\mathcal{T}}(s.\mu) \rangle$ . We include the values of public memory in this tuple because the high-program may be non-deterministic and we need to constrain the non-determinism to be identical across certain traces.

3) *General Adversary Tampering ( $\mathcal{G}$ )*: The adversary  $\mathcal{G}$  tampers with system state by executing an unbounded number of instructions to modify architectural and microarchitectural state. Adversary tampering is constrained in only two ways.

- 1) **(Conformant Store Addresses)** For every non-speculative state in which an untrusted store is executed, if the target address of the store must be part of the set of addresses the adversary is allowed to write to,  $\mathcal{U}_\mu^{wr}$ . We denote a trace  $\pi$  where every state satisfies the above condition by the predicate  $conformantStoreAddrs(\pi)$ , defined as follows.

$$\begin{aligned} \forall i. \pi^i.n = 0 \wedge \pi^i.pc[0] \notin \mathcal{T}_\rho & \implies \\ \pi^i.\iota = \text{mem} := \text{mem}[e_1 \rightarrow e_2] \wedge \pi^i.\Delta[0, e_1] \Downarrow a & \implies \\ a \in \mathcal{U}_\mu^{wr} & \end{aligned}$$

- 2) **(Conformant Entrypoints)** Non-speculative adversary jumps to trusted code must target the entrypoint  $\mathcal{EP}$ . A trace  $\pi$  where every transition from untrusted to trusted code satisfies this condition is denoted by the predicate  $conformantEntrypoints(\pi)$ . This is defined as follows:

$$\begin{aligned} \forall i, j. i < j \wedge \pi^i.n = \pi^j.n = 0 & \implies \\ (\forall k. i < k < j \implies \pi^k.n \neq 0) & \implies \\ \pi^i.pc[0] \notin \mathcal{T}_\rho \wedge \pi^j.pc[0] \in \mathcal{T}_\rho & \implies \\ \pi^j.pc[0] = \mathcal{EP} & \end{aligned}$$

The above constraints says that if  $\pi^i$  and  $\pi^j$  are non-speculative states, all states between  $\pi^i$  and  $\pi^j$  are

<sup>3</sup>Note: We will use the notation  $s.field$  to refer to elements of the tuple.

speculative, and  $\pi^i$  is part of the untrusted component while  $\pi^j$  is part the trusted component, then  $\pi^j$  must necessarily be at the entrypoint. Note this does not preclude speculative execution of “gadgets” in the trusted code that do not begin at the entrypoint.

The condition conformant store addresses captures the fact that the adversary cannot write to arbitrary memory locations. Conformant entrypoints ensures that execution of the trusted code starts at the entrypoint.

4) *Conformant Traces*: A trace  $\pi$  where: (i)  $\pi^0$  is a non-speculative state and the trusted component has been initialized:  $\pi^0.n = 0 \wedge \text{init}_{\mathcal{T}}(\pi^0)$ , (ii) every state  $\pi^i$  satisfies the conformant stores condition and (iii) every pair of states  $\pi^i$  and  $\pi^j$ , where  $i < j$ , satisfy the conformant entrypoints condition is called a conformant trace, denoted by  $\text{conformant}(\pi)$ .

$$\begin{aligned} \text{conformant}(\pi) \doteq & \pi^0.n = 0 \wedge \text{init}_{\mathcal{T}}(\pi^0) & \wedge \\ & \text{conformantStoreAddrs}(\pi) & \wedge \\ & \text{conformantEntrypoints}(\pi) & (4) \end{aligned}$$

5) *Adversary Observations*: We model an adversary who can observe all architectural state and most microarchitectural state when executing; i.e. when  $n = 0$  and  $pc[0] \notin \mathcal{T}_\rho$ . Specifically, the adversary can observe the following:

- 1) non-speculative register values:  $\Delta[0, r]$  for all  $r$ .
- 2) non-speculative values stored at all memory addresses in the set  $\mathcal{U}_\mu^{rd}$ :  $\mu[0, a]$  for all  $a \in \mathcal{U}_\mu^{rd}$ .
- 3) the trace of instruction and data memory accesses:  $\omega$ .
- 4) the branch predictor state  $\beta$ .

The above implies that two states  $s = \langle \Pi, \Delta, \mu, pc, \omega, \beta, n, \iota \rangle$  and  $s' = \langle \Pi, \Delta', \mu', pc', \omega', \beta', n', \iota' \rangle$  are low-equivalent, denoted  $s \approx_{\mathcal{L}} s'$ , iff  $(n = 0 \wedge pc[0] \notin \mathcal{T}_\rho) \implies (\forall r. \Delta[0, r] = \Delta'[0, r]) \wedge (\forall a. a \in \mathcal{U}_\mu^{rd} \implies \mu[0, a] = \mu'[0, a]) \wedge \omega = \omega' \wedge \beta = \beta'$ .

We do not allow the adversary to observe  $\Delta[n, r]$  and  $\mu[n, a]$  for  $n > 0$  because there is no way to “output” speculative state except through a microarchitectural side-channel. The trace of memory accesses  $\omega$  captures all side-channel leaks via caches, prefetches, DRAM, and other memory-related side-channels. The branch predictor state  $\beta$  models all leaks caused by the branch predictor side-channel.

### C. Formalization of the Security Property

Using the above definitions, we are now ready to formalize the secure speculation property, shown in Equation 5.

$$\begin{aligned} & \forall \pi_1, \pi_2, \pi_3, \pi_4. \\ & \text{conformant}(\pi_1) \wedge \text{conformant}(\pi_2) & \implies \\ & \text{conformant}(\pi_3) \wedge \text{conformant}(\pi_4) & \implies \\ & \forall i. \neg \text{mispred}(\pi_1^i.n, \pi_1^i.\beta, \pi_1^i.pc) & \implies \\ & \forall i. \neg \text{mispred}(\pi_2^i.n, \pi_2^i.\beta, \pi_2^i.pc) & \implies \\ & \exists i. \text{mispred}(\pi_3^i.n, \pi_3^i.\beta, \pi_3^i.pc) & \implies \\ & \exists i. \text{mispred}(\pi_4^i.n, \pi_4^i.\beta, \pi_4^i.pc) & \implies \\ & \text{op}_{\mathcal{L}}(\pi_1) = \text{op}_{\mathcal{L}}(\pi_2) = \text{op}_{\mathcal{L}}(\pi_3) = \text{op}_{\mathcal{L}}(\pi_4) & \implies \\ & \text{op}_{\mathcal{H}}(\pi_1) = \text{op}_{\mathcal{H}}(\pi_3) \wedge \text{op}_{\mathcal{H}}(\pi_2) = \text{op}_{\mathcal{H}}(\pi_4) & \implies \\ & \pi_1 \approx_{\mathcal{L}} \pi_2 \wedge \pi_3^0 \approx_{\mathcal{L}} \pi_4^0 & \implies \\ & \pi_3 \approx_{\mathcal{L}} \pi_4 & (5) \end{aligned}$$

This an instantiation of the TPOD property shown in Equation 3. The trace property  $T$  is satisfied when no misspeculation occurs:  $\pi \in T \iff \forall i. \neg \text{mispred}(\pi^i.n, \pi^i.\beta, \pi^i.pc)$ .<sup>4</sup> The trace property  $U$  requires that all traces be conformant as defined in Equation 4. This ensures we only search for violations among traces representing valid executions of our system/adversary model.

A violation of Equation 5 occurs when there exists a sequence of adversary instructions such that traces  $\pi_1$  and  $\pi_2$  are low-equivalent, but  $\pi_3$  and  $\pi_4$  are not low-equivalent. In other words, we have an information leak that only occurs on a speculative processor; i.e. a transient execution vulnerability.

1) *Adversary Reduction Lemma*: The general adversary’s tampering described in § IV-B3 allows the adversary to execute an unbounded number of arbitrary instructions. While this is fully general, it makes automated reasoning unscalable. To address this problem, we introduce a simpler “havocing adversary”  $\mathcal{H}$  and prove that this adversary is as powerful as the general adversary  $\mathcal{G}$ .

$\mathcal{H}$  executes only one instruction that modifies non-speculative state:  $\text{havoc}(\Delta, \text{mem}[\mathcal{U}_\mu^{wr}], \beta)$ . The semantics of this instruction are shown in Figure 8; it sets the registers, program counter, adversary writeable memory addresses and branch predictor to unconstrained values (i.e. “havocs” them).

*Lemma 2*: Every sequence of  $s_i, \dots, s_j$  with  $\text{op}_{\mathcal{L}}(s_j) \neq \perp$  and  $s_j.n = 0$  for every  $i \leq j \leq k$  can be simulated by a single  $\text{havoc}(\Delta, \text{mem}[\mathcal{U}_\mu^{wr}], \beta)$  instruction.

The adversary reduction lemma lets us replace all sequences of non-speculative instructions executed by the adversary with  $\text{havoc}$ ’s and helps scale verification. It is important to note that we cannot replace instruction sequences which contain speculative instructions because these may contain exploitable transient execution gadgets.

2) *Discussion and Limitations*: An important implication of the secure speculation property is that if a program satisfies Equation 5, then all observational determinism properties where low-equivalence is defined over  $\omega$ ,  $\mu$  and  $\beta$  that hold for non-speculative execution of the program also hold for

<sup>4</sup>Or equivalently in linear temporal logic:  $\pi \models \square \neg \text{mispred}$ .

speculative executions. For instance, a tool like CacheAudit [56, 57] can be used to verify that the cache accesses of a program are independent of some secret. Note that even though a program’s non-speculative execution may not leak information through cache (this is what CacheAudit verifies) that *does not mean* that its speculation execution will have the same properties. This is because CacheAudit does not model speculative execution. However, if we do prove Equation 5 for a program, then all properties proven by tools like CacheAudit also apply to the program’s speculative execution.

Our operational semantics are for in-order processors only. Nevertheless, the secure speculation property can be used to analyze out-of-order execution and other speculation (e.g., memory address disambiguation) in a conceptually straightforward way by extending the semantics to model these features.

Specific programs may need additional constraints on the traces to avoid spurious counterexamples, especially if the set of secrets  $\mathcal{S}_{\mathcal{T}}$  is over-specified. For example, in Figure 3(b), a tuple of traces where the  $i < N$  never occurs would cause a violation of Equation 5 if  $\mathcal{S}_{\mathcal{T}}$  also contained the addresses that point to `a1` and `a2`.

## V. VERIFICATION APPROACH

We have implemented an automated verifier to answer the following question: Given a program (e.g., C code) as input, does it satisfy the secure speculation property in Equation 5?

Our approach is fairly standard, based on the method of *self-composition* (see, e.g., [70]). For lack of space, we present only the essential aspects. Given the input program, we translate it into a transition system based on the adversary model and operational semantics presented in the previous section. The secure speculation property is a 4-safety property, meaning that we can turn it into a safety property to be checked on a 4-way self-composition of the transition system. We use a model checker based on satisfiability modulo theories (SMT) solving to check whether the safety property holds for this 4-way self-composition. The model checker uses either bounded model checking (to find violations of the property) or k-induction (to prove the property).

The main new aspect of our verifier is the implementation of the transformation of the program into a transition system. We rely on two tools: the Binary Analysis Platform (BAP) [66] to translate x86 binaries into an intermediate format called BIL, and UCLID5 [47], an SMT-based model checking tool supporting both bounded model checking (BMC) and k-induction. BIL is an assembly-like intermediate language similar to AIR (described in Sec. IV). Overall, our workflow for each input program is as follows:

- 1) Compile C source code containing the victim function into an x86 binary file.
- 2) Translate the x86 binary file using BAP into the BIL intermediate language.
- 3) Translate the BIL into UCLID5 models and check the secure speculation property via self-composition. For each program, we first obtain a counterexample via BMC demonstrating the vulnerability; then, we insert

an lfence at an appropriate point and prove the secure speculation property via k-induction.

We note that this workflow may be abstracted to a more general TPOD property.

The translation from BIL to UCLID5 implements the operational semantics given earlier, with the following key steps:

- 1) Datatypes in the BIL program such as addresses, memories, and words are converted to uninterpreted types for more scalable analysis and to obtain a more portable model that is not only specific to 32-bit or 64-bit architectures.
- 2) Each basic block of the BIL program is considered an atomic step of the transition system in the UCLID5 model after which the safety property is checked on the 4-way self-composition. This suffices as the deviations in behavior between the 4 traces happen at branch points.
- 3) At any speculative transition step, the program can resolve a misspeculation as per the RESOLVE rule.
- 4) All state variables are initialized to symbolic constants with the exception of the memory, where it is initialized to have the same value at every address except the program-specific secret address that stores the secret.

An example of the BIL to UCLID5 translation is provided side by side in the Appendix.

Given the model, the implication chain of the secure speculation property is translated into a number of assumptions and invariants. The invariants which we wish to check are whether the speculative program traces diverge in control flow, branch prediction or memory access observations, but only in the cases that they do not for the non-speculative traces. Proofs by induction require a few additional auxiliary invariants, whereas bounded model checking does not.

## VI. CASE STUDIES

We used our verifier for a proof-of-concept demonstration to detect whether or not an arbitrary snippet of C code is vulnerable to the Spectre class of attacks. As benchmarks, we rely on Paul Kocher’s list of 15 example victim functions vulnerable to the Spectre attack [30] in addition to the examples we presented earlier.

In particular, we show here results on Examples 1, 5, 7, 8, 10, 11, and 15 from Paul Kocher’s list, along with the example from Figure 3 (c), and an example with nested if statements. We chose these based on what we believe are illustrative of a wide range of victim functions that are not easily detectable using the currently existing static analysis tools such as Qspectre [71], which was only able to detect the first two examples in Paul Kocher’s list. We begin with a brief explanation of some of the benchmark examples and then discuss the results from applying bounded model checking and induction with our secure speculation property on our UCLID5 5 models. Fig. 9 lists all benchmarks we discuss here.

*Example 5 (Figure 9b):* This example is similar to the first variant but implemented within a for loop. The low-security argument  $x$  may be larger than the array size, which causes

```

1 void victim_function_v01(unsigned x) {
2     if (x < array1_size) {
3         __mm_lfence();
4         temp &= array2[array1[x] * 512];
5     }
6 }

```

(a) Example 1: Original Spectre BCB (bounds check bypass) example.

```

1 void victim_function_v07(unsigned x) {
2     static unsigned last_x = 0;
3     if (x == last_x) {
4         __mm_lfence();
5         temp &= array2[array1[x] * 512];
6     }
7     if (x < array1_size)
8         last_x = x;
9 }

```

(c) Example 7: BCB with unsafe static variable check.

```

1 void victim_function_v10(unsigned x, unsigned k) {
2     if (x < array1_size) {
3         __mm_lfence();
4         if (array1[x] == k)
5             temp &= array2[0];
6     }
7 }

```

(e) Example 10: BCB using an additional attacker controlled input.

```

1 void victim_function_v15(unsigned *x) {
2     if (*x < array1_size) {
3         __mm_lfence();
4         temp &= array2[array1[*x] * 512];
5     }
6 }

```

(g) Example 15: BCB using attacker controlled pointer.

```

1 void victim_function_v05(unsigned x) {
2     size_t i;
3     if (x < array1_size) {
4         for (i = x - 1; i > 0; i--)
5             __mm_lfence();
6         temp &= array2[array1[i] * 512];
7     }
8 }

```

(b) Example 5: BCB with a for loop.

```

1 void victim_function_v08(unsigned x) {
2     result = (x < array1_size);
3     __mm_lfence();
4     temp &= array2[array1[result ? (x + 1) : 0] *
5                 512];
}

```

(d) Example 8: BCB with the ternary conditional operator.

```

1 void victim_function_v11(unsigned x) {
2     if (x < array1_size) {
3         __mm_lfence();
4         temp = memcmp(&temp, array2 + (array1[x]
5                 * 512), 1);
6     }
}

```

(f) Example 11: BCB using the memory comparison function.

```

1 void victim_function_nested_ifs(unsigned x) {
2     unsigned val1, val2;
3     if (x < array1_size) {
4         val1 = array1[x];
5         if (val1 & 1) {
6             __mm_lfence();
7             val2 = array2[0];
8         }
9     }
10 }

```

(h) Example NI: BCB with nested if statements.

Figure 9: Examples that were verified for secure speculation with `__mm_lfence()` implemented. In all code snippets assume that that arguments `x` and `k` are untrusted (low-security) inputs to the trusted (high-security) victim functions.

the attack to occur as in example 1, but if `x` is within bounds of the array, note that condition `i > 0` is also potentially vulnerable to the attack.<sup>5</sup>

*Example 7 (Figure 9c):* This example is interesting because it depends on the value of a static variable updated from a previous call of the function. Thus every call to the function should not make the second array access unless `x` is equal to `last_x`.

*Example 8 (Figure 9d):* The ternary operator is interesting because the program counter is allowed to jump to two different basic blocks for the computation of the second array memory access in the BIL as opposed to one block as in Example 1.

*Example 10 (Figure 9e):* This is the first example where a second load dependent on a secret is not required for a leak. Knowing whether or not `array2[0]` was accessed is enough to leak the secret at `array1[x]`.

*Example 11 (Figure 9f):* This example uses a call to `memcmp` to leak the secret, but because of the single byte access, it gets

<sup>5</sup>In Paul Kocher’s list, the condition is `x >= 0`, but this introduces an infinite loop.

optimized out in the assembly to a single load and store.

*Example 15 (figure 9g):* This example is interesting because it passes a pointer instead of an integer as the attacker controlled input. We assume that the value stored at the pointer is constant across traces to ignore cases where the attacker forces a secret dependent branch during non-speculative execution.

	ex1	ex5	ex7	ex8	ex10	ex11	ex15	Fig. 3c	NI
BMC	6.6	9.0	10.2	5.7	9.6	6.4	5.8	6.6	12.9
Ind	5.0	5.0	5.7	4.6	5.8	5.9	4.8	4.8	5.4

Table I: Runtime (sec.) of each example using 5 steps for bounded model checking to find vulnerabilities and 1 step induction to prove correctness after inserting a memory fence. These experiments were run on a machine with an 2.20GHz Intel(R) Core(TM) i7-2670QM CPU with 5737MiB of RAM.

*Example NI (Figure 9h)* In this example, nested if statements cause the attack to occur without a second address load dependent on a secret. If the programs speculatively choose not to execute the second if statement, but only one program

eventually executes the second if as a result of a resolution, then a leak can occur.

Table I lists the run-time (in seconds) required for each verification task with the memory fences implemented. As can be seen, the verifier is able to prove the correctness of these programs within a few seconds. Although these programs are small, this exercise gives us confidence that the method could be useful on larger programs. Scaling to larger programs will need to adopt a stronger software model checking engine and property specific abstractions. We assert that with these improvements, it will be possible to prove secure speculation for larger programs.

## VII. RELATED WORK

The most closely related work to ours is CheckMate [72] which uses happens-before graphs to analyze transient execution vulnerabilities. The insight in CheckMate is that happens-before graphs encode information about the orders in which instructions can be executed. By searching for patterns in the graph where branches are followed by dependent loads, an architectural model can be analyzed for susceptibility to Spectre/Meltdown. A key difference between CheckMate and our approach is that we are not matching patterns of vulnerable instructions. Our verification is semantic, not pattern-based. In particular, the example showing conditional vulnerability in Figure 3(c) cannot be precisely captured by CheckMate.

Another closely related effort is by McIlroy et al. [73] who introduce a formal model of speculative execution in modern processors and analyze it for transient execution vulnerabilities. Similar to our work, they too introduce speculative operational semantics and their model includes indirect jumps and a timer. However, they do not have a property formulation to capture transient execution vulnerabilities, nor do they present an automated verification approach for finding these vulnerabilities. Their semantics can be readily combined with our adversary model and property formulation to algorithmically search for transient execution vulnerabilities.

The Spectre vulnerability was discovered by Kocher et al. [2, 3] while Meltdown was discovered by Lipp et al. [1]. Their public disclosure has triggered an avalanche of new transient execution vulnerabilities notable among which are Foreshadow [4] which attacked enclave platforms and virtual machine monitors, SpectreRSB [7] and Ret2Spec [6]. A thorough study of transient execution vulnerabilities was done by Canella et al. [5]. Transient execution vulnerabilities build on the rich literature of microarchitectural side-channel attacks [14–20, 50–53]. Verification of mitigations to “traditional” side-channel attacks is a well-studied problem [54–64].

TPOD in general and secure speculation in particular are examples of hyperproperties [40]. They are also closely related to notions of noninterference introduced by Goguen and Meseguer [41], separability proposed by Rushby [74], and observational determinism [42–44]. Our verification method is based on self-composition which has been well-studied; see, for example, Barthe et al. [70]. While we take a straightfor-

ward approach to using self-composition, more sophisticated approaches are also possible in some cases (e.g., [75]).

## VIII. CONCLUSION

This paper presented a formal approach for secure speculative execution on modern processors, a key part of which is a formal specification of secure speculation that abstracts away from the particulars of specific vulnerabilities. Our secure speculation formulation is an instance of trace property-dependent observational determinism, a new class of information flow security properties introduced by this work. We introduced an adversary model and an automated approach to verifying secure speculation and demonstrated the approach on several programs that have been used to illustrate the Spectre class of vulnerabilities. To the best of our knowledge, ours is the first effort to guarantee provably secure speculation. In future work, we plan to evaluate our approach on larger programs and more complex platforms like out-of-order processors.

## REFERENCES

- [1] M. Lipp, M. Schwarz, D. Gruss, T. Prescher, W. Haas, S. Mangard, P. Kocher, D. Genkin, Y. Yarom, and M. Hamburg, “Meltdown,” *ArXiv e-prints*, Jan. 2018.
- [2] P. Kocher, D. Genkin, D. Gruss, W. Haas, M. Hamburg, M. Lipp, S. Mangard, T. Prescher, M. Schwarz, and Y. Yarom, “Spectre attacks: Exploiting speculative execution,” *ArXiv e-prints*, Jan. 2018.
- [3] 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,” *Proceedings of the IEEE Symposium on Security and Privacy*, pp. 19–37, 2019.
- [4] J. V. Bulck, M. Minkin, O. Weisse, D. Genkin, B. Kasikci, F. Piessens, M. Silberstein, T. F. Wenisch, Y. Yarom, and R. Strackx, “Foreshadow: Extracting the keys to the intel SGX kingdom with transient out-of-order execution,” in *27th USENIX Security Symposium (USENIX Security 18)*. Baltimore, MD: USENIX Association, 2018, p. 991–1008.
- [5] 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,” *CoRR*, vol. abs/1811.05441, 2018. [Online]. Available: <http://arxiv.org/abs/1811.05441>
- [6] G. Maisuradze and C. Rossow, “Ret2spec: Speculative execution using return stack buffers,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’18, 2018.
- [7] E. M. Koruyeh, K. N. Khasawneh, C. Song, and N. Abu-Ghazaleh, “Spectre returns! speculation attacks using the return stack buffer,” in *12th USENIX Workshop on Offensive Technologies (WOOT 18)*. Baltimore, MD: USENIX Association, 2018.
- [8] J. Horn, “Read privileged memory with a side-channel,” 2018. [Online]. Available: <https://googleprojectzero.blogspot.com/2018/01/reading-privileged-memory-with-side.html>
- [9] Intel, “Bounds Check Bypass / CVE-2017-5753 / INTEL-SA-00088,” 2018. [Online]. Available: <https://software.intel.com/security-software-guidance/software-guidance/bounds-check-bypass>
- [10] —, “Branch Target Injection / CVE-2017-5715 / INTEL-SA-00088,” 2018. [Online]. Available: <https://software.intel.com/security-software-guidance/software-guidance/branch-target-injection>
- [11] —, “Rogue System Register Read / CVE-2018-3640 / INTEL-SA-00115,” 2018. [Online]. Available: <https://software.intel.com/security-software-guidance/software-guidance/rogue-system-register-read>
- [12] —, “Speculative Store Bypass / CVE-2018-3639 / INTEL-SA-00115,” 2018. [Online]. Available: <https://software.intel.com/security-software-guidance/software-guidance/speculative-store-bypass>
- [13] —, “L1 Terminal Fault / CVE-2018-3615 , CVE-2018-3620,CVE-2018-3646 / INTEL-SA-00161,” 2018. [Online]. Available: <https://software.intel.com/security-software-guidance/software-guidance/l1-terminal-fault>
- [14] C. Percival, “Cache missing for fun and profit,” 2005.

- [15] O. Açıçmez, Ç. K. Koç, and J.-P. Seifert, "Predicting secret keys via branch prediction," in *Cryptographers Track at the RSA Conference*. Springer, 2007, pp. 225–242.
- [16] Y. Oren, V. P. Kemerlis, S. Sethumadhavan, and A. D. Keromytis, "The Spy in the Sandbox - Practical Cache Attacks in Javascript," *CoRR*, vol. abs/1502.07373, 2015. [Online]. Available: <http://arxiv.org/abs/1502.07373>
- [17] F. Liu, Y. Yarom, Q. Ge, G. Heiser, and R. B. Lee, "Last-Level Cache Side-Channel Attacks Are Practical," in *Proceedings of the 2015 IEEE Symposium on Security and Privacy*. IEEE Computer Society, 2015, pp. 605–622. [Online]. Available: <http://dx.doi.org/10.1109/SP.2015.43>
- [18] G. Irazoqui, T. Eisenbarth, and B. Sunar, "SSA: A Shared Cache Attack That Works across Cores and Defies VM Sandboxing – and Its Application to AES," in *IEEE Symposium on Security and Privacy*, May 2015, pp. 591–604.
- [19] M. Schwarz, S. Weiser, D. Gruss, C. Maurice, and S. Mangard, "Malware Guard Extension: Using SGX to Conceal Cache Attacks," *CoRR*, vol. abs/1702.08719, 2017. [Online]. Available: <http://arxiv.org/abs/1702.08719>
- [20] S. Lee, M. Shih, P. Gera, T. Kim, H. Kim, and M. Peinado, "Inferring Fine-grained Control Flow Inside SGX Enclaves with Branch Shadowing," *CoRR*, vol. abs/1611.06952, 2016. [Online]. Available: <http://arxiv.org/abs/1611.06952>
- [21] O. Oleksenko, B. Trach, T. Reiher, M. Silberstein, and C. Fetzer, "You shall not bypass: Employing data dependencies to prevent bounds check bypass," *arXiv preprint arXiv:1805.08506*, 2018.
- [22] Intel, "Deep Dive: Analyzing Potential Bounds Check Bypass Vulnerabilities," 2018. [Online]. Available: <https://software.intel.com/security-software-guidance/insights/deep-dive-analyzing-potential-bounds-check-bypass-vulnerabilities>
- [23] Microsoft, "Spectre mitigations in MSVC," 2018. [Online]. Available: <https://devblogs.microsoft.com/cppblog/spectre-mitigations-in-msvc/>
- [24] P. Turner, "Retpoline: a software construct for preventing branch-target-injection," 2018. [Online]. Available: <https://support.google.com/faqs/answer/7625886>
- [25] Intel, "Deep Dive: Managed Runtime Speculative Execution Side Channel Mitigations," 2018. [Online]. Available: <https://software.intel.com/security-software-guidance/insights/deep-dive-managed-runtime-speculative-execution-side-channel-mitigations>
- [26] —, "Deep Dive: Mitigation Overview for Side Channel Exploits in Linux," 2018. [Online]. Available: <https://software.intel.com/security-software-guidance/insights/deep-dive-mitigation-overview-side-channel-exploits-linux>
- [27] B. Stuart, "Current state of mitigations for spectre within operating systems," in *Proceedings of the 4th Wiesbaden Workshop on Advanced Microkernel Operating Systems*, 2018.
- [28] Microsoft, "ADV180012 — Microsoft Guidance for Speculative Store Bypass," 2018. [Online]. Available: <https://portal.mscc.microsoft.com/en-US/security-guidance/advisory/ADV180012>
- [29] J. Cable, "Update on Spectre and Meltdown security updates for Windows devices," 2018. [Online]. Available: <https://blogs.windows.com/windowsexperience/2018/03/01/update-on-spectre-and-meltdown-security-updates-for-windows-devices/>
- [30] P. Kocher, "Spectre mitigations in microsoft's c/c++ compiler," Feb 2018. [Online]. Available: <https://www.paulkocher.com/doc/MicrosoftCompilerSpectreMitigation.html>
- [31] J. E. Smith and G. S. Sohi, "The microarchitecture of superscalar processors," *Proceedings of the IEEE*, vol. 83, no. 12, pp. 1609–1624, Dec 1995.
- [32] J. P. Shen and M. Lipasti, *Fundamentals of Superscalar Processor Design*. In Press, 2001.
- [33] T.-Y. Yeh and Y. N. Patt, "Two-level adaptive training branch prediction," in *Proceedings of the 24th Annual International Symposium on Microarchitecture*, ser. MICRO 24, 1991, pp. 51–61.
- [34] S. McFarling, "Combining branch predictors," Technical Report TN-36, Digital Western Research Laboratory, Tech. Rep., 1993.
- [35] D. M. Gallagher, W. Y. Chen, S. A. Mahlke, J. C. Gyllenhaal, and W. Hwu, "Dynamic memory disambiguation using the memory conflict buffer," in *Proceedings of the Sixth International Conference on Architectural Support for Programming Languages and Operating Systems*, ser. ASPLOS VI. ACM, 1994, pp. 183–193. [Online]. Available: <http://doi.acm.org/10.1145/195473.195534>
- [36] M. H. Lipasti, C. B. Wilkerson, and J. P. Shen, "Value locality and load value prediction," in *Proceedings of the Seventh International Conference on Architectural Support for Programming Languages and Operating Systems*, 1996.
- [37] V. Kiriansky, I. Lebedev, S. Amarasinghe, S. Devadas, and J. Emer, "Dawg: A defense against cache timing attacks in speculative execution processors," *Cryptology ePrint Archive*, Report 2018/418, 2018, <https://eprint.iacr.org/2018/418>.
- [38] M. Larabel, "Benchmarking The Work-In-Progress Spectre/STIBP Code On The Way For Linux 4.20," 2018. [Online]. Available: <https://www.phoronix.com/scan.php?page=article&item=linux-420wip-stibp&num=1>
- [39] L. Tung, "Linus Torvalds: After big Linux performance hit, Spectre v2 patch needs curbs," 2018. [Online]. Available: <https://www.zdnet.com/article/linus-torvalds-after-big-linux-performance-hit-spectre-v2-patch-needs-curbs/>
- [40] M. R. Clarkson and F. B. Schneider, "Hyperproperties," *Journal of Computer Security*, vol. 18, no. 6, pp. 1157–1210, Sep. 2010.
- [41] J. A. Goguen and J. Meseguer, "Security Policies and Security Models," in *IEEE Symposium on Security and Privacy*, 1982, pp. 11–20.
- [42] J. Mclean, "Proving Noninterference and Functional Correctness Using Traces," *Journal of Computer Security*, vol. 1, pp. 37–58, 1992.
- [43] A. W. Roscoe, "CSP and determinism in security modelling," in *Proceedings of the 1995 IEEE Symposium on Security and Privacy, Oakland, California, USA, May 8-10, 1995*, 1995, pp. 114–127.
- [44] S. Zdancewic and A. C. Myers, "Observational determinism for concurrent program security," in *Proceedings of the 16th IEEE Computer Security Foundations Workshop*. IEEE, 2003, pp. 29–43.
- [45] G. Barthe, P. R. D'Argenio, and T. Rezk, "Secure information flow by self-composition," in *17th IEEE Computer Security Foundations Workshop (CSFW-17)*, 2004, pp. 100–114.
- [46] T. Terauchi and A. Aiken, "Secure Information Flow as a Safety Problem," in *Static Analysis Symposium (SAS '05)*, ser. LNCS 3672, 2005, pp. 352–367.
- [47] S. A. Seshia and P. Subramanyan, "UCLID5: Integrating modeling, verification, synthesis and learning," in *Proceedings of the 16th ACM-IEEE International Conference on Formal Methods and Models for System Design (MEMOCODE)*, October 2018.
- [48] UCLID5 Verification and Synthesis System, "Available at <http://github.com/uclid-org/uclid/>."
- [49] 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," *Communications of the ACM*, vol. 53, no. 1, pp. 91–99, 2010. [Online]. Available: <http://doi.acm.org/10.1145/1629175.1629203>
- [50] D. Gruss, C. Maurice, A. Fogh, M. Lipp, and S. Mangard, "Prefetch side-channel attacks: Bypassing smap and kernel aslr," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '16, 2016, pp. 368–379.
- [51] Y. Shin, H. C. Kim, D. Kwon, J. H. Jeong, and J. Hur, "Unveiling hardware-based data prefetcher, a hidden source of information leakage," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '18, 2018, pp. 131–145.
- [52] B. Gras, K. Razavi, H. Bos, and C. Giuffrida, "Translation leak-aside buffer: Defeating cache side-channel protections with TLB attacks," in *27th USENIX Security Symposium (USENIX Security 18)*, 2018, pp. 955–972.
- [53] P. Pessl, D. Gruss, C. Maurice, M. Schwarz, and S. Mangard, "DRAMA: Exploiting DRAM addressing for cross-cpu attacks," in *25th USENIX Security Symposium (USENIX Security 16)*, 2016, pp. 565–581.
- [54] J. Protzenko, J.-K. Zinzindohoué, A. Rastogi, T. Ramanandaro, P. Wang, S. Zanella-Béguelin, A. Delignat-Lavaud, C. Hrițcu, K. Bhargavan, C. Fourmet *et al.*, "Verified low-level programming embedded in f," *Proceedings of the ACM on Programming Languages*, vol. 1, no. ICFP, p. 17, 2017.
- [55] B. Bond, C. Hawblitzel, M. Kapritsos, K. R. M. Leino, J. R. Lorch, B. Parno, A. Rane, S. Setty, and L. Thompson, "Vale: Verifying high-performance cryptographic assembly code," in *26th USENIX Security Symposium (USENIX Security 17)*, 2017, pp. 917–934.
- [56] G. Doychev, D. Feld, B. Köpf, L. Mauborgne, and J. Reineke, "Cacheaudit: A Tool for the Static Analysis of Cache Side Channels," in *Presented as part of the 22nd USENIX Security Symposium (USENIX Security 13)*. Washington, D.C.: USENIX, 2013, pp. 431–446.
- [57] G. Doychev, B. Köpf, L. Mauborgne, and J. Reineke, "Cacheaudit: A tool for the static analysis of cache side channels," *ACM Transactions on Information and System Security (TISSEC)*, vol. 18, no. 1, p. 4, 2015.
- [58] J. B. Almeida, M. Barbosa, J. S. Pinto, and B. Vieira, "Formal verifica-

- tion of side-channel countermeasures using self-composition,” *Science of Computer Programming*, vol. 78, no. 7, pp. 796–812, 2013.
- [59] G. Barthe, B. Grégoire, and V. Laporte, “Secure compilation of side-channel countermeasures: the case of cryptographic constant-time,” in *2018 IEEE 31st Computer Security Foundations Symposium (CSF)*. IEEE, 2018, pp. 328–343.
- [60] J. B. Almeida, M. Barbosa, G. Barthe, and F. Dupressoir, “Verifiable side-channel security of cryptographic implementations: constant-time mce-cbc,” in *International Conference on Fast Software Encryption*. Springer, 2016, pp. 163–184.
- [61] G. Barthe, G. Betarte, J. Campo, C. Luna, and D. Pichardie, “System-level non-interference for constant-time cryptography,” in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2014, pp. 1267–1279.
- [62] J. B. Almeida, M. Barbosa, G. Barthe, F. Dupressoir, and M. Emmi, “Verifying constant-time implementations,” in *25th USENIX Security Symposium (USENIX Security 16)*, 2016, pp. 53–70.
- [63] H. Eldib, C. Wang, and P. Schaumont, “Formal verification of software countermeasures against side-channel attacks,” *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 24, no. 2, p. 11, 2014.
- [64] —, “SMT-based verification of software countermeasures against side-channel attacks,” in *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, 2014, pp. 62–77.
- [65] B. Alpern and F. B. Schneider, “Defining liveness,” *Information Processing Letters*, vol. 21, no. 4, pp. 181 – 185, 1985.
- [66] D. Brumley, I. Jager, T. Avgerinos, and E. J. Schwartz, “BAP: A binary analysis platform,” in *Proceedings of the 23rd International Conference on Computer Aided Verification*, ser. CAV’11, 2011, pp. 463–469.
- [67] Binary Analysis Platform (BAP) Repository, “Available at <https://github.com/BinaryAnalysisPlatform/bap>,” 2019.
- [68] M. Abadi, M. Budi, U. Erlingsson, and J. Ligatti, “Control-flow integrity,” in *Proceedings of the 12th ACM Conference on Computer and Communications Security*, ser. CCS ’05, 2005, pp. 340–353.
- [69] C. Tice, T. Roeder, P. Collingbourne, S. Checkoway, U. Erlingsson, L. Lozano, and G. Pike, “Enforcing forward-edge control-flow integrity in gcc &#38; llvm,” in *Proceedings of the 23rd USENIX Conference on Security Symposium*, ser. SEC’14, 2014, pp. 941–955.
- [70] G. Barthe, P. R. D’Argenio, and T. Rezk, “Secure information flow by self-composition,” *Mathematical Structures in Computer Science*, vol. 21, no. 6, pp. 1207–1252, 2011.
- [71] Microsoft, “qspectre,” Oct 2018. [Online]. Available: <https://docs.microsoft.com/en-us/cpp/build/reference/qspectre?view=vs-2017>
- [72] C. Trippel, D. Lustig, and M. Martonosi, “Checkmate: Automated synthesis of hardware exploits and security litmus tests,” in *51st Annual IEEE/ACM International Symposium on Microarchitecture, MICRO 2018*, 2018, pp. 947–960.
- [73] 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,” *arXiv preprint arXiv:1902.05178*, 2019.
- [74] J. M. Rushby, “Proof of separability: A verification technique for a class of a security kernels,” in *International Symposium on Programming, 5th Colloquium, Torino, Italy, April 6-8, 1982, Proceedings*, 1982, pp. 352–367.
- [75] M. Sousa and I. Dillig, “Cartesian hoare logic for verifying k-safety properties,” in *ACM SIGPLAN Notices*, vol. 51, no. 6. ACM, 2016, pp. 57–69.

## APPENDIX

```

1 // block 1
2 000001ba: sub victim_function_v01()
3 00000172:
4 ... // function prologue
5 00000178: RAX := mem[0x601050, e1]:u64
6 0000017a: CF := mem[RBP + 0xFFFFFFFFFFFFFFF8, e1]:
   u64 < RAX
7 00000180: when ~CF goto %000001b4
8 00000181: goto %00000182

```

(a) Example 1: Block 1 BIL

```

1 // block 2
2 00000182:
3 00000183: RAX := mem[0x601040, e1]:u64
4 00000184: RDX := mem[0x601038, e1]:u64
5 00000185: RCX := mem[RBP + 0xFFFFFFFFFFFFFFF8, e1]:
   u64
6 00000187: RCX := RCX << 3
7 0000018f: v274 := RCX
8 00000190: RDX := RDX + v274
9 00000197: RDX := mem[RDX, e1]:u64
10 00000199: RDX := RDX << 0xC
11 000001a1: v284 := RDX
12 000001a2: RAX := RAX + v284
13 000001a9: RDX := mem[RAX, e1]:u64
14 000001aa: RAX := mem[0x601058, e1]:u64
15 000001ab: RAX := RAX & RDX
16 000001b2: mem := mem with [0x601058, e1]:u64 <- RAX
17 000001b3: goto %000001b4

```

(c) Example 1: Block 2 BIL

```

1 // block 3
2 000001b4:
3 ... // function epilogue
4 000001b9: return v295

```

(e) Example 1: Block 3 BIL

```

1 procedure do_block1()
2   modifies (...);
3 {
4   ... // function prologue
5   call (RAX) = load_mem(array1_size_addr);
6   CF = common.lessthan(RDI, RAX);
7   call branch(!CF, block3, block2);
8 }

```

(b) Example 1: Block 1 UCLID5 procedure.

```

1 procedure do_block2()
2   modifies (...);
3 {
4   var v274, v284 : word_t;
5   assume (common.read(mem, temp_addr) == common.read
   (common.mem_init, temp_addr));
6   if (lfence && spec_level != common.spec_idx0) {
7     call do_resolve();
8   } else {
9     call (RAX) = load_mem(array2_addr);
10    call (RDX) = load_mem(array1_addr);
11    RCX = RDI;
12    RCX = common.left_shift(common.val0x3, RCX);
13    v274 = RCX;
14    RDX = common.add(RDX, v274);
15    call (RDX) = load_mem(RDX);
16    RDX = common.left_shift(common.val0xC, RDX);
17    v284 = RDX;
18    RAX = common.add(RAX, v284);
19    call (RDX) = load_mem(RAX);
20    call (RAX) = load_mem(temp_addr);
21    RAX = common.land(RAX, RDX);
22    call store_mem(temp_addr, RAX);
23    pc = block3;
24    br_pred_state = common.update_br_pred(
   br_pred_state, true);
25 }
26 }

```

(d) Example 1: Block 2 UCLID5 procedure.

```

1 procedure do_block3()
2   modifies (...);
3 {
4   RAX = common.val0x0;
5   RDX = common.val0x0;
6   ... // function epilogue
7   pc = halt;
8   br_pred_state = common.update_br_pred(
   br_pred_state, true);
9 }

```

(f) Example 1: Block 3 UCLID5 procedure.

Figure 10: Translation from BIL to UCLID5 of the victim function in example 1 from Paul Kocher's list. The left side shows the BIL representation of the x86 binary and the right side shows the corresponding translation from a BIL block to a UCLID5 procedure.