

InSpectre: Breaking and Fixing Microarchitectural Vulnerabilities by Formal Analysis

Roberto Guanciale

robertog@kth.se

KTH Royal Institute of Technology
Stockholm SE-100 44, Sweden

Musard Balliu

musard@kth.se

KTH Royal Institute of Technology
Stockholm SE-100 44, Sweden

Mads Dam

mfd@kth.se

KTH Royal Institute of Technology
Stockholm SE-100 44, Sweden

ABSTRACT

The recent Spectre attacks have demonstrated the fundamental insecurity of current computer microarchitecture. The attacks use features like pipelining, out-of-order and speculation to extract arbitrary information about the memory contents of a process. A comprehensive formal microarchitectural model capable of representing the forms of out-of-order and speculative behavior that can meaningfully be implemented in a high performance pipelined architecture has not yet emerged. Such a model would be very useful, as it would allow the existence and non-existence of vulnerabilities, and soundness of countermeasures to be formally established.

This paper presents such a model targeting single core processors. The model is intentionally very general and provides an infrastructure to define models of real CPUs. It incorporates microarchitectural features that underpin all known Spectre vulnerabilities. We use the model to elucidate the security of existing and new vulnerabilities, as well as to formally analyze the effectiveness of proposed countermeasures. Specifically, we discover three new (potential) vulnerabilities, including a new variant of Spectre v4, a vulnerability on speculative fetching, and a vulnerability on out-of-order execution, and analyze the effectiveness of existing countermeasures including constant time and serializing instructions.

1 INTRODUCTION

The wealth of vulnerabilities that have followed on from Spectre and Meltdown [31, 35] have provided ample evidence of the fundamental insecurity of current computer microarchitecture. The extensive use of instruction level parallelism in the form of out-of-order (OoO) and speculative execution has produced designs with side channels that can be exploited by attackers to learn sensitive information about the memory contents of a process. One witness of the subtlety of the issues is the more than 50 years passed since pipelining, caching, and OoO execution, cf. IBM S/360, was first introduced.

Another witness is the fact that two years after the discovery of Spectre, a comprehensive understanding of the security implications of pipeline related microarchitecture features has yet to emerge. One result is the ongoing arms race between researchers discovering new Spectre-related vulnerabilities [9], and CPU vendors providing patches followed by informal arguments [5]. The

security and effectiveness of the currently proposed countermeasures is unknown, and there are continuously new vulnerabilities appearing that exploit specific microarchitecture features.

It is important to note that side channels and functional correctness are to a large extent orthogonal. The latter is usually proved by reducing pipelined behaviour to sequential behaviour through some form of refinement-based argument. The past decades have seen a significant body of work in this area, cf. [1, 8, 38, 47], addressing rich sets of features of concrete pipeline designs such as OoO, speculation, and self-modifying code. Functional correctness, however, focuses on programs' input-output behaviour and fails to adequately capture the differential aspects of speculation and instruction reordering that are at the root of Spectre-like vulnerabilities. For a systematic study of the latter we argue that new tools that are not necessarily tied to any specific pipeline architecture are needed.

Along this line, several recent works [12, 14, 22, 40] have started to propose formal microarchitectural models using information flow analysis to identify information leaks arising from speculative execution in a principled manner. These models capture specific speculation features, e.g. branch prediction, and variants of Spectre, in particular variant 1, and design analyses that detect known attacks [12, 22, 53]. While these approaches illustrate the usefulness of formal models in analyzing microarchitecture leaks, features lying at the heart of modern CPUs such as OoO execution and many forms of speculation remain largely unexplored, implying that new vulnerabilities may still exist.

Contributions. This work presents InSpectre, the first comprehensive model capable of capturing OoO execution and all forms of speculation that can be meaningfully implemented in the context of a high performance pipeline. The model is intentionally very general and provides an infrastructure to define models of real CPUs (Section 3), which can be used to analyze effectiveness of countermeasures for a given processor.

Our first contribution is a novel semantics supporting microarchitectural features such as OoO execution, non-atomicity of instructions, and various forms of speculation, including branch prediction, jump target prediction, return address prediction, and dependency prediction. Additionally, the semantics supports features such as address aliasing, dynamic references, store forward, and OoO memory commits, which are necessary to model all known variants of Spectre. The semantics implements the stages of an abstract pipeline supporting OoO (Section 4) and speculative execution (Section 5). In line with existing work [12, 22], our security condition formalizes the intuition that optimizations should not introduce additional information leaks (conditional noninterference,

Section 2). We use this condition to show that InSpectre can reproduce all four variants of Spectre.

As a second contribution, we use InSpectre to discover three new potential vulnerabilities. The first vulnerability shows that CPUs supporting only OoO may leak sensitive information. We discovered the second vulnerability while attempting to validate a CPU vendor’s claim that microarchitectures like Cortex A53 are immune to Spectre vulnerabilities because they support only speculative fetching [5]. Our model reveals that this may not be the case. The third vulnerability is a variant of Spectre v4 showing that speculation of a dependency, rather than speculation of a non-dependency as in Spectre v4, between a load and a store operation may also leak sensitive information.

Finally, as a third contribution, we leverage InSpectre to analyze the effectiveness of some existing countermeasures. We found that constant-time [7] analysis is unsound for processors supporting only OoO, and propose a provably secure fix that enables constant-time analysis to ensure security for such processors.

2 SECURITY MODEL

Our security model has the following ingredients: (i) an *execution* model which is given by the execution semantics of a program; (ii) an *attacker* model specifying the observations of an attacker; (iii) a *security policy* specifying the parts of the program state that contain sensitive/high information, and the parts that contain public/low information; (iv) a *security condition* capturing a program’s security with respect to an execution model, an attacker model, and a security policy.

First, we consider a general model of attacker that observes the interaction between the CPU and the memory subsystem. This model has been used (e.g., [3]) to capture information leaks via cache-based side channels transparently without an explicit cache model. It can capture trace-driven attackers that can interleave with the victim’s execution and indirectly observe, for instance using Flush+Reload [20], the victim’s cache footprint via latency jitters. The attacker can observe the address of a memory load $dl\ v$ (data load from memory address v), the address of a memory store $ds\ v$ (data store to memory address v), as well as the value of the program counter $il\ v$ (instruction load from memory address v) [42].

We assume a transition relation $\rightarrow \subseteq States \times Obs \times States$ to model the execution semantics of a program as a state transformer producing observations $l \in Obs$. The reflexive and transitive closure of \rightarrow induces a set of executions $\pi \in \Pi$. The function $trace : \Pi \mapsto Obs^*$ extracts the sequence of observations of an execution.

The security policy is defined by an *indistinguishability relation* $\sim \subseteq States \times States$. The relation \sim determines the security of information that is initially stored in a state, modeling the set of initial states that an attacker is not allowed to discriminate. These states represent the initial *uncertainty* of an attacker about sensitive information.

The security condition defines the security of a program on the target execution model (e.g., the speculation model) \rightarrow_t conditionally on the security of the same program on the reference, i.e. sequential, model \rightarrow_r , by requiring that the target model does not leak more information than the reference model for a policy \sim .

Definition 2.1 (Conditional Noninterference). Let \sim be a security policy and \rightarrow_t and \rightarrow_r be transition relations for the target and reference models of a system. The system is conditionally non-interferent if for all $\sigma_1, \sigma_2 \in States$ such that $\sigma_1 \sim \sigma_2$, if for every $\pi_1 = \sigma_1 \rightarrow_r \dots$ there exists $\pi_2 = \sigma_2 \rightarrow_r \dots$ such that $trace(\pi_1) = trace(\pi_2)$ then for every $\rho_1 = \sigma_1 \rightarrow_t \dots$ there exists $\rho_2 = \sigma_2 \rightarrow_t \dots$ such that $trace(\rho_1) = trace(\rho_2)$.

Conditional noninterference captures only the new information leaks that may be introduced by model \rightarrow_t , and ignores any leaks already present in model \rightarrow_r . The target model is constructed in two steps. First, we present an OoO model that extends the sequential model, which is deterministic, by allowing evaluation to proceed out-of-order. Then the OoO model is further extended by adding speculation. At each step the traces of the abstract model are included in the extended model, and a memory consistency result demonstrates that the per location sequence of memory stores is the same for both models. This establishes functional correctness. Conditional noninterference then establishes security of each extension. Each such step strictly increases the set of possible traces by adding nondeterminism. Since refinement is often viewed as essentially elimination of nondeterminism, one can think of the extensions as “inverse refinements”. Since conditional noninterference considers a possibilistic setting, it does not account for information leaks through the number of initial indistinguishable states.

We now elucidate the advantages of conditional noninterference as compared to standard notions of noninterference and declassification. Suppose we define the security condition directly on the target model, in the style of standard noninterference.

Definition 2.2 (Noninterference). Let P be a program with transition relation \rightarrow and \sim_P a security policy. P satisfies noninterference up to \sim_P if for all $\sigma_1, \sigma_2 \in States$ such that $\sigma_1 \sim_P \sigma_2$ and executions $\pi_1 = \sigma_1 \rightarrow \dots$, there exists an execution $\pi_2 = \sigma_2 \rightarrow \dots$ such that $trace(\pi_1) = trace(\pi_2)$.

Noninterference ensures that if the observations do not enable an attacker to refine his knowledge of sensitive information beyond what is allowed by the policy \sim_P , the program can be considered secure. Noninterference can accommodate partial release of sensitive information by refining the definition of the indistinguishability relation \sim_P . In our context, a precise definition of \sim_P can be challenging to define. However, we ultimately aim at showing that the OoO/speculative model does not leak more information than the in-order (sequential) model, thus capturing the intuition that microarchitectural features like OoO and speculation should not introduce additional leaks. Therefore, instead of defining the policy \sim_P explicitly, we split it into two relations \sim (as in Def. 2.1) and \sim_D , where the former models information of the initial state that is known by the attacker, i.e., the public resources, and the latter models information that the attacker is allowed to learn during the execution via observations. Hence, $\sim_P = \sim \cap \sim_D$.

This characterization allows for a simpler formulation of the security condition that is transparent on the definition of \sim_D , as described in Def. 2.1.

3 FORMAL MICROARCHITECTURAL MODEL

We introduce a Machine Independent Language (MIL) which we use to define the semantics of microarchitectural features such as OoO and speculative execution. We use MIL as a form of abstract microcode language: A target language for translating ISA instructions and reasoning about features that may cause vulnerabilities like Spectre. Microinstructions in MIL represent atomic actions that can be executed by the CPU, emulating the pipeline phases in an abstract manner. This model is intentionally very general and provides an infrastructure to define models of real microarchitectures.

We consider a domain of values $v \in V$, a program counter $pc \in \mathcal{PC}$, a finite set of register/flag identifiers $r_0, \dots, r_n, f, z \in \mathcal{R} \subseteq V$, and a finite set of memory addresses $a_0, \dots, a_m \in \mathcal{M} \subseteq V$. The language can be easily extended to support other type of resources, e.g., registers for vector operations. We assume a total order $<$ on a set of names $t_0, t_1, \dots \in N$, which we use to uniquely identify microinstructions. We write $N_1 < N_2$ if for every pair $(t_1, t_2) \in N_1 \times N_2$ it holds that $t_1 < t_2$.

Microinstructions $\iota \in \mathcal{I}$ are conditional atomic single assignments. A microinstruction $\iota = t \leftarrow c?o$ is uniquely identified by its name $t \in N$ and consists of a boolean guard c , which determines if the assignment should be executed, and an operation $o \in \mathcal{Op}$. The MIL language has three types of operations:

$$\begin{aligned} e &::= v \mid t \mid e_1 + e_2 \mid e_1 > e_2 \mid \dots \\ o &::= e \mid ld \tau t_a \mid st \tau t_a t_v \end{aligned}$$

An internal operation e is an expression over standard finite arithmetic and can additionally refer to names in N and values in V . A resource load operation $ld \tau t_a$, where $\tau \in \{\mathcal{PC}, \mathcal{R}, \mathcal{M}\}$, loads the value of resource τ addressed by t_a . We support three types of resources: The program counter \mathcal{PC} , registers \mathcal{R} , and memory locations \mathcal{M} . A resource store operation $st \tau t_a t_v$ uses the value of t_v to update the resource τ addressed by t_a .

The free names $fn(\iota)$ of an instruction $\iota = t \leftarrow c?o$ is the set of names occurring in c or o , the bound names, $bn(\iota)$, is the singleton $\{t\}$, and the names $n(\iota)$ is $fn(\iota) \cup bn(\iota)$.

To model the internal state of a CPU pipeline, we can translate an ISA instruction as multiple microinstructions. For an ISA instruction at address $v \in \mathcal{M}$ and a name $t \in N$, the function $translate(v, t)$ returns the MIL translation of the instruction at address v , ensuring that the names of the microinstructions thus generated are greater than t . Because we assume code to not be self-modifying, an instruction can be statically identified by its address in memory. We assume that the translation function satisfies the properties: (i) for all $t_1, t_2 \in translate(v, t)$, if $t_1 \neq t_2$ then $bn(t_1) \cap bn(t_2) = \emptyset$; for all $\iota \in translate(v, t)$, (ii) $fn(\iota) < bn(\iota)$, and (iii) $\{t\} < n(\iota)$.

These properties ensure that names uniquely identify microinstructions, the name parameters of a single instruction form a Directed Acyclic Graph, the translated microinstructions are assigned

names greater than t , and the translation of two different ISA instructions does not have direct inter-instruction dependencies (but may have indirect ones).


3.1 MIL Program Examples

We introduce some illustrative examples of MIL programs, using their graph representation. For clarity, we omit conditions whenever they are true and visualize only the immediate dependencies between graph elements.

Consider an ISA instruction that increments the value of register r_1 by one, i.e., $r_1 := r_1 + 1$. The instruction can be translated in MIL as follows:

$$\left\{ \begin{array}{l} t_1 \leftarrow r_1, t_2 \leftarrow ld \mathcal{R} t_1, t_3 \leftarrow t_2 + 1, t_4 \leftarrow st \mathcal{R} t_1 t_3, \\ t_5 \leftarrow ld \mathcal{PC}, t_6 \leftarrow t_5 + 4, t_7 \leftarrow st \mathcal{PC} t_6 \end{array} \right\}$$

Intuitively, t_1 refers to the identifier of target register r_1 , t_2 loads the current value of register r_1 , t_3 executes the increment, and t_4 stores the result of t_3 in the register store. The translation of an ISA instruction also updates the program counter to enable the execution of the next instruction. In this case, the program counter is increased by 4, unconditionally. Notice that we omit the program counter's address, since there is only one such resource. We can graphically represent this set of microinstructions using the following graph:




Example 1: $r_1 := r_1 + 1$

In the following we adopt syntactic sugar to use expressions, in place of names, for the address and value of load and store operations. This can be eliminated by introducing the proper intermediary internal assignments. This permits to rewrite the previous example as:



The translation of multiple ISA instructions results in disconnected graphs. This reflects the fact that inter-instruction dependencies may not be statically identified due to dynamic references and must be dynamically resolved by the MIL semantics. When translating multiple instructions, we use the following convention for generated names: the name t_{ij} identifies the j -th microinstruction resulting from the translation of the i -th instruction. Our convention induces a total (lexicographical) order over names (i.e., $t_{ij} < t_{i'j'}$ iff $(i < i') \vee (i = i' \wedge j < j')$), which respects the properties of the translation function.

MIL is expressive enough to support conditional instructions like conditional arithmetic and conditional move. Conditional branches can be modeled in MIL via microinstructions that are guarded by complementary conditions. For instance, the `beq` instruction, which jumps to address a if the z flag is set, can be translated as in Example 2.



Example 2: beq a




Figure 1: OoO semantics: Microinstruction lifecycle

4 OUT-OF-ORDER SEMANTICS

This section presents an OoO semantics for MIL programs, which is extended in Section 5 to account for speculation. The in-order semantics in Section 6 is obtained by constraining the OoO semantics to enforce program order evaluation. We prove memory consistency for both the OoO and speculative semantics. The results show that, for each semantics, the sequences of per location memory stores are the same, thus establishing functional correctness of the OoO and speculative semantics.

4.1 States, Transitions, Observations

We formalize the semantics via a transition relation $\sigma \xrightarrow{l} \sigma'$, which maps a state σ to a state σ' , and produces a (possibly empty, represented by a dot (\cdot)) observation $l \in Obs$, eliding the dot when convenient. As in Section 2, $Obs = \{\cdot, dl\ v, ds\ v, il\ v\}$ captures the attacker model.

States σ are tuples (I, s, C, F) where: (i) I is a set of MIL microinstructions, (ii) $s \in Stores = N \rightarrow V$ is a (partial) storage function from names to values recording microinstructions' execution results, (iii) $C \subseteq N$ is a set of names of store operations that have been committed to the memory subsystem, (iv) $F \subseteq N$ is a set of names of program counter store operations that have been processed, causing the ISA instruction at the stored location to be fetched and decoded.

In the following we write $s[t \mapsto v]$ for substitution of value v for name t in store s . We use $f(x)\downarrow$ to represent that the partial function f is defined on x , and $f(x)\uparrow$ if not $f(x)\downarrow$. We write $\text{dom}(f)$ for the domain of a partial function f . We also use $f|_D$ to represent the restriction of function f to domain D . The semantics of expressions is $[e] : Stores \rightarrow V$ and is defined as expected. An expression is undefined if at least one name is undefined in a storage, i.e., $[e](s)\uparrow \Leftrightarrow \text{fn}(e) \not\subseteq \text{dom}(s)$. For $\sigma = (I, s, C, F)$ we use $[e]\sigma$ for $[e]s, \sigma(t)\uparrow$ for $s(t)\uparrow$, and $l \in \sigma$ for $l \in I$.

4.2 Microinstruction Lifecycle

Figure 1 represents the microinstruction lifecycle in the OoO semantics. For a given state (I, s, C, F) , a microinstruction $\iota = t \leftarrow$

$c?o \in I$ can be in one of four different states. A microinstruction in state Decoded (represented by a gray circle) has not been executed, committed or fetched ($s(t)\uparrow, t \notin C, t \notin F$), and its guard is either true ($[c]s$) or undefined ($[c](s)\uparrow$). If the guard is false, i.e., $\neg[c]s$, the instruction is considered as Discarded (not shown). A microinstruction is able to move to state Executed (represented by a simple circle whose content is $s(t)$) if its guard evaluates to true and all dependencies have been executed. Subsequently, an Executed store microinstruction can either be committed to the memory subsystem (Committed: $t \in C$, represented by a bold circle), or, if it is a program counter store, assign the program counter, causing a new ISA instruction to be fetched and decoded (Fetched: $t \in F$, represented by a double circle). Accordingly, the transition of a program counter store to state Fetched leads to the spawn of a collection of newly decoded microinstructions (i.e., the translation of the subsequent ISA instruction) in state Decoded. The labels of the edges in the diagram correspond to the names of the transition rules of Section 4.4.

4.3 Semantics of Single Microinstructions

The semantics is defined in two steps: we first define the semantics of single microinstructions, then introduce the operational semantics of MIL programs. The semantics of a microinstruction $[l] : States \rightarrow (V \times Obs) \cup \{\perp\}$ returns either a value and an observation, or \perp if the microinstruction cannot be executed.

(Internal operations) The semantics of internal operations is straightforward:


$$[t \leftarrow c?e]\sigma = \begin{cases} ([e]\sigma, \cdot) & \text{if } [e](\sigma)\downarrow \\ \perp & \text{otherwise} \end{cases}$$

An internal operation can be executed as soon as its dependencies are available. In Example 1, the semantics of internal operation t_1 is defined for the empty storage \emptyset , since it does not refer to any names. However, the semantics of t_3 is undefined in \emptyset , since it depends on the value of t_2 that is not available in \emptyset .

(Store operations) The semantics of store operations is defined as follows:

$$[t \leftarrow c?st\ \tau\ t_a\ t_v]\sigma = \begin{cases} ([t_v]\sigma, \cdot) & \text{if } [t_v](\sigma)\downarrow \wedge [t_a](\sigma)\downarrow \\ \perp & \text{otherwise} \end{cases}$$

A resource update can be executed as soon as both the address of the resource and the value are available. Observe that this rule models the *internal* execution of a resource update and not its commit to the memory subsystem. These internal updates are not observable by a programmer, therefore there is no restriction on their execution order. As an example, the ISA program $r_1 := 0; r_2 := r_1; r_1 := 1$ can be implemented by the following microinstructions:



The semantics of t_{11} , i.e., $st\ \mathcal{R}\ r_1\ 0$, and t_{31} , i.e., $st\ \mathcal{R}\ r_1\ 1$, is defined in \emptyset , and yields $(0, \cdot)$ and $(1, \cdot)$, respectively. As we will see, the operational semantics is in charge of ordering resource updates to preserve consistency and dependencies.

(Load operations) While the semantics of internal operations and store operations only depends on the execution of their operands, load operations may depend on past store operations. This requires identifying the previous resource update that determines the correct value to be loaded. We use the following definitions to compute the set of store operations that may affect a load operation.

Definition 4.1. Consider a load operation $t \leftarrow c?ld \tau t_a$

- $str\text{-}may(\sigma, t) = \{t' \leftarrow c'?st \tau t'_a t'_v \in \sigma \mid t' < t \wedge ([c']\sigma \vee [c'](\sigma)\uparrow) \wedge ([t'_a]\sigma = [t_a]\sigma \vee \sigma(t'_a)\uparrow \vee \sigma(t_a)\uparrow)\}$ is the set of stores that *may* affect the load address of t in state σ .
- $str\text{-}act(\sigma, t) = \{t' \leftarrow c'?st \tau t'_a t'_v \in str\text{-}may(\sigma, t) \mid \neg \exists t'' \leftarrow c''?st \tau t''_a t''_v \in str\text{-}may(\sigma, t). t'' > t' \wedge [c'']\sigma \wedge [t''_a]\sigma \in \{[t_a]\sigma, [t'_a]\sigma\}\}$ is the set of *active* stores.

The stores that may affect the address of t are the stores that: (i) have not been discarded, namely they can be executed ($[c]\sigma$) or may be executed ($c(\sigma)\uparrow$), and (ii) the store address in t'_a may result in the same address as the load address in t_a , namely either they both evaluate to the same address ($\sigma(t'_a) = \sigma(t_a)$), or the store address is unknown ($\sigma(t'_a)\uparrow$), or the load address is unknown ($\sigma(t_a)\uparrow$).

The active stores of t are the stores that may affect the load address computed by t_a , and, there are no subsequent stores t'' on the same address as the load address in t_a , or on the same address as the store address in t'_a . This set determines the “minimal” set of store operations that may affect a load operation from address t_a .

The definitions of $str\text{-}act(\sigma, t)$ and $str\text{-}may(\sigma, t)$ are naturally extended to stores $t \leftarrow c?st \tau t_a t_v$. These definitions allow us to define the semantics of loads:

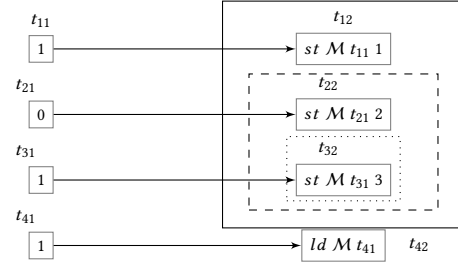
$$[t \leftarrow c?ld \tau t_a]\sigma = \begin{cases} ([t_s]\sigma, l) & \text{if } bn(str\text{-}act(\sigma, t)) = \{t_s\} \wedge \\ & \sigma(t_a)\downarrow \wedge \sigma(t_s)\downarrow \\ \perp & \text{otherwise} \end{cases}$$

where $l = \begin{cases} dl \sigma(t_a) & \text{if } t_s \in C \wedge \tau = M \\ \cdot & \text{otherwise} \end{cases}$

A load operation can be executed if the set of active stores consists of a singleton set with bound name t_s , i.e., the store causing t_a to be assigned is uniquely determined, and both the address t_a of the load and the address t_s of the store can be evaluated in state σ .

Note that the semantics allows forwarding the result of a store to another microinstruction before it is committed to memory. In fact, if the active store is yet to be committed to memory, i.e., $t_s \notin C$, it is possible for the store to forward its data to the load, without causing an interaction with the memory subsystem (i.e., $l = \cdot$). Otherwise, the load yields an observation of a data load from address $\sigma(t_a)$.

Example 3 illustrates the semantics of loads. The program writes 1 into address 1, then writes 2 in 0, overwrites address 1 with 3, and finally loads from address 1. We use active stores to dynamically compute the dependencies of load operations. Let σ_0 be a state containing microinstructions as in the example, and having empty storage. For this state, the active store for the load t_{42} , i.e., $str\text{-}act(\sigma_0, t_{42})$, consists of all stores of the example, as depicted by the solid rectangle. Since none of microinstructions that compute the addresses have been executed, the address t_{41} of the load is



Example 3: $\ast(1):=1$; $\ast(0):=2$; $\ast(1):=3$; $\ast(1)$;

unknown, hence, we cannot exclude any store from affecting the address that will be used by t_{42} . Therefore, the load cannot be executed in σ_0 . This set of active stores will shrink during execution as more information becomes available through the storage.

Let the storage of σ_1 be $\{t_{11} \mapsto 1; t_{31} \mapsto 1\}$, i.e., the result of executing t_{11} and t_{31} . The active stores $str\text{-}act(\sigma_1, t_{42})$ consist of microinstructions depicted by the dashed rectangle. Observe that the store t_{12} is in $str\text{-}may(\sigma_1, t_{42})$, however there exists a subsequent store, namely t_{32} , that overwrites the effects of t_{12} on the same memory address. Therefore, t_{12} is no longer an active store and it can safely be discarded.

Let the storage of σ_2 be $\{t_{11} \mapsto 1; t_{31} \mapsto 1, t_{41} \mapsto 1\}$, i.e., the result of executing t_{11} , t_{31} and t_{41} . The active stores $str\text{-}act(\sigma_2, t_{42})$ now consist of the singleton set $\{t_{32}\}$ as depicted by the dotted rectangle. This is because the address t_{41} of the load can be computed in state σ_2 . Although t_{22} is still in $str\text{-}may(\sigma_2, t_{42})$, there is a subsequent store, t_{32} , that will certainly affect the address of the load. Therefore, t_{22} is no longer an active store.

Finally, let the storage of σ_3 be $\{t_{11} \mapsto 1; t_{31} \mapsto 1, t_{41} \mapsto 1, t_{32} \mapsto 3\}$, i.e., the result of executing t_{11} , t_{31} , t_{41} , and t_{32} . Once $str\text{-}act$ has been reduced to a singleton set ($\{t_{32}\}$), and the active-store has been executed ($\sigma_3(t_{32})\downarrow$), the semantics of the load is defined. This yields the same value as the store in t_{32} . If the store t_{32} has been committed to memory, the execution of the load yields the observation $dl \ 1$.

4.4 Operational Semantics

We can now define the microinstructions' transition relation $\sigma \xrightarrow{l} \sigma'$, implementing the lifecycle of Section 4.2.

(Execute) A microinstruction can be executed if it hasn't already been executed ($s(t)\uparrow$), the guard holds ($[c]s$), and the dependencies have been resolved ($[l](s)\downarrow$):

$$(EXE) \frac{l = t \leftarrow c?o \in I \quad s(t)\uparrow \quad [c]s \quad [l]\sigma = (v, l)}{\sigma = (I, s, C, F) \xrightarrow{l} (I, s[t \mapsto v], C, F)}$$

Observe that if l is a load from the memory subsystem, the rule can produce the observation of a data load.

(Commit) Once a memory store has been executed ($s(t)\downarrow$), it can be committed to memory, yielding an observation. The rule ensures that stores can only be committed once ($t \notin C$) and that stores on the same address are committed in program order, by checking that all past stores are in C , i.e., $bn(str\text{-}may(\sigma, t)) \subseteq C$.

$$\begin{array}{c}
\text{(CMT)} \quad \frac{t \leftarrow c?st \ \mathcal{M} \ t_a \ t_v \in I \quad s(t) \downarrow \quad t \notin C}{bn(str\text{-}may(\sigma, t)) \subseteq C} \\
\sigma = (I, s, C, F) \xrightarrow{ds \ s(t_a)} (I, s, C \cup \{t\}, F)
\end{array}$$

In summary, stores can be executed internally in any order, however, they are committed in order. In Example 3, if σ has storage $s = \{t_{11} \mapsto 1; t_{12} \mapsto 1; t_{31} \mapsto 1; t_{32} \mapsto 3\}$ and commits $C = \emptyset$, then only t_{12} can be committed, since t_{22} has not been executed and $bn(str\text{-}may(\sigma, t_{32})) \not\subseteq C$. Notice that t_{22} is in the may stores since its address has not been resolved. Therefore, t_{32} can be committed only after t_{12} has been committed and t_{21} has been executed. However, the commit of t_{32} does not have to wait for the commit or execution of t_{22} . In fact, if σ' has storage $s' = s \cup \{t_{21} \mapsto 0\}$ then $bn(str\text{-}may(\sigma', t_{32})) = \{t_{12}\}$. That is, the order of store commits is only enforced per location, as expected.

(Fetch-Decode) A program counter store enables the fetching and decoding (i.e., translating) of a new ISA instruction. The rule for fetching is similar to the rule for commit, since instructions are fetched in order. The set F keeps track of program counter updates whose resulting instruction has been fetched and ensures that instructions are not fetched or decoded twice. Fetching the result of a program counter update yields the observation of an instruction load from address a .

$$\begin{array}{c}
\text{(Frc)} \quad \frac{t \leftarrow c?st \ \mathcal{P}C \ t_v \in I \quad s(t) = a \quad t \notin F}{bn(str\text{-}may(\sigma, t)) \subseteq F} \\
\sigma = (I, s, C, F) \xrightarrow{il \ a} (I \cup I', s, C, F \cup \{t\}) \\
\text{where } I' = \text{translate}(a, \max(I))
\end{array}$$

Write $\max(I)$ for the largest name t in I and $\text{translate}(a, \max(I))$ for the translation of the instruction at address a , ensuring that the names of the microinstructions thus generated are greater than $\max(I)$.

(Remarks on OoO semantics) The three rules of the semantics reflect the atomicity of MIL microinstructions: A transition can affect a single microinstruction by either assigning a value to the storage, extending the set of commits, or extending the set of fetches. In the following, we use $\text{step-param}(\sigma, \sigma') = (\alpha, t)$ to identify the rule $\alpha \in \{\text{EXE}, \text{CMT}(a, v), \text{Frc}(I)\}$ that enables $\sigma \rightarrow \sigma'$ and the name t of the affected microinstruction. In case of commits we also extract the modified address a and the saved value v , in case of fetches we extract the newly decoded microinstructions I . The semantics preserves several invariants: Let $(I, s, C, F) = \sigma$ if $\alpha = \text{EXE}$ then $\sigma(t) \uparrow$; if $\alpha = \text{CMT}(a, v)$ then $t \notin C$ and free names (i.e., address and value) of the corresponding microinstruction are defined in s ; if $\alpha = \text{Frc}(I')$ then $t \notin F$; all state components are monotonic.

(Initial state) In order to bootstrap the computation, we assume that the set of microinstructions of the initial state contains one store for each memory address and register, the value of these stores is the initial value of the corresponding resource, and that these stores are in the storage and commits of the initial state.

5 SPECULATIVE SEMANTICS

We now extend the OoO semantics to support speculation. We add two new components to the states: a set of names $P \subseteq n(I)$ whose values have been predicted as result of speculation, and a partial function $\delta : N \rightarrow S$ recording, for each name t , the storage

dependencies at time of execution of the microinstruction identified by t . Therefore, a state in the speculative semantics is a tuple $h = (I, s, C, F, \delta, P)$ where $\sigma = (I, s, C, F)$ is the corresponding state in the OoO semantics. Abusing notation we write (σ, δ, P) to denote a state in the speculative semantics, and use h, h_1, \dots to range over these states. Informally, $\delta(t)$ is a *snapshot* of the storage that affects the value of t due to speculative predictions. As we will see, these snapshots are needed in order to match speculative states with non-speculative states, and to restore the state of the execution in case of misspeculation.

5.1 Managing Microinstruction Dependencies

The execution of a microinstruction may depend on local (intra-) instruction dependencies, the names appearing freely in a microinstruction, as well as cross (inter-) instruction dependencies, caused by memory or register loads.

Definition 5.1. Let $t \leftarrow c?o \in \sigma$. The dependencies of t in σ are

$$\text{deps}(t, \sigma) = \text{fn}(t \leftarrow c?o) \cup \text{depsX}(t, \sigma)$$

where the cross-instruction dependencies are defined as


$$\text{depsX}(t, \sigma) = \begin{cases} \emptyset, & \text{if } t \text{ is not a load} \\ \text{asn}(\sigma, t) \cup \text{srcs}(\sigma, t), & \text{otherwise.} \end{cases}$$

Cross-dependencies are nonempty only for loads and consist of the names of active stores affecting t in state σ , $\text{asn}(\sigma, t) = \text{bn}(str\text{-}act(\sigma, t))$, plus, the names of stores potentially intervening between the earliest active store and t (we call $\text{srcs}(\sigma, t)$ the *potential sources* of t), which are defined as

$$\text{srcs}(\sigma, t) = \bigcup \{ \text{fn}(c'), \{t'_a\} \mid \min(\text{asn}(\sigma, t)) \leq t' < t, t' \leftarrow c'?st \ \tau \ t'_a \ t'_v \in \sigma \}$$

Intuitively, a load depends on the execution of active stores that may affect the address of that load. Moreover, the fact that a name t^* is in the set of active stores asn depends on the addresses and guards of all stores between t^* and t . This is because their values will determine the actual store that affects the address of the load t . Thanks to our ordering relation $<$ between names, we can use the minimum name $\min(\text{asn})$ in asn to compute all stores between any name in asn and t , thus extracting the free names of their guards and addresses.

The following figure illustrates dependencies of the load from Example 3:



If $s = \{t_{11} \mapsto 1; t_{21} \mapsto 0; t_{41} \mapsto 1\}$ then the set of active stores names asn for t_{42} is $\text{bn}(str\text{-}act(\sigma, t_{42})) = \{t_{12}, t_{32}\}$, as depicted by the solid ellipses. In particular, $\min(\text{asn}) = t_{12}$. We consider all stores between t_{12} and the load t_{42} (i.e., t_{12}, t_{22} , and t_{32}), and add to the set of cross-dependencies the names in their guards


and addresses, namely t_{11}, t_{21} and t_{31} , as depicted by the dashed rectangle. Observe that t_{21} is in the set of cross-dependencies, although t_{22} is not an active store. This is because membership of t_{12} in the active stores' set depends on the address t_{21} being set to 0, i.e., $s(t_{21}) = 0$. Therefore, the set of cross-dependencies $depsX(t_{42}, \sigma) = \{t_{12}, t_{32}, t_{11}, t_{21}, t_{31}\}$. Finally, the local dependencies of the load t_{42} consist of its parameter t_{41} (the dotted ellipsis), such that $deps(t_{42}, \sigma) = \{t_{12}, t_{32}, t_{11}, t_{21}, t_{31}, t_{42}\}$.

We verify that the dependencies $deps$ are computed correctly.

Definition 5.2 (t -equivalence). Let σ_1 and σ_2 be states with storage s_1 and s_2 , and t_1 and t_2 be the microinstructions identified by t . Then σ_1 and σ_2 are t -equivalent, $\sigma_1 \sim_t \sigma_2$, if $t_1 = t_2$, $s_1|_{\hat{fn}(t_1)} = s_2|_{\hat{fn}(t_2)}$, and if t 's microinstruction is a load with dependencies $T_i = deps(t, \sigma_i)$ and active stores $SA_i = str-act(\sigma_i|_{T_i}, t)$ for $i \in \{1, 2\}$ then $SA_1 = SA_2$ and $s_1|_{SA_1} = s_2|_{SA_2}$.

Intuitively, t -equivalence states that, if the microinstruction named with t depends (in the sense of $deps$) in both states on the same active stores and these stores assign the same value to t , then the microinstruction has the same dependencies, it is enabled, and it produces the same result in both states.

We use three possible states of the example above to illustrate t -equivalence: σ_1 is a state reachable in the OoO semantics, σ_2 and σ_3 may result from misspeculating the value of t_{31} to be 0 and 5 respectively.



The states σ_1 and σ_2 are not t_{42} -equivalent. In particular, $T_1 = deps(t_{42}, \sigma_1) = \{t_{31}, t_{32}\}$ (notice that t_{12} and t_{22} are not in the dependencies because by we know that $t_{31} \mapsto 1$ and $t_{41} \mapsto 1$), $T_2 = deps(t_{42}, \sigma_2) = \{t_{11}, t_{21}, t_{31}, t_{12}\}$. Notice that $\sigma_1|_{T_1}$ and $\sigma_2|_{T_2}$ contain all the information needed to evaluate the semantics of t_{42} in σ_1 and σ_2 respectively. In this case $SA_1 = str-act(\sigma_1|_{T_1}, t_{42}) = \{t_{32}\}$, and $SA_2 = str-act(\sigma_2|_{T_2}, t_{42}) = \{t_{12}\}$ hence $SA_1 \neq SA_2$: the two states lead the load t_{42} to take the result produced by two different memory stores.

The states σ_2 and σ_3 are t_{42} -equivalent. In fact, $T_3 = deps(t_{42}, \sigma_3) = \{t_{11}, t_{21}, t_{31}, t_{12}\}$ and $SA_3 = str-act(\sigma_3|_{T_3}, t_{42}) = \{t_{12}\}$. Therefore, $SA_2 = SA_3$ and $s_2|_{SA_2} = s_3|_{SA_3}$: The two states lead the load t_{42} to take the result produced by the same memory stores.

LEMMA 5.3. *If $\sigma_1 \sim_t \sigma_2$ and t 's microinstruction in σ_1 is $t = t \leftarrow c?o$, then $deps(t, \sigma_1) = deps(t, \sigma_2)$, $[c]\sigma_1 = [c]\sigma_2$, and if $[i]\sigma_1 = (v_1, l_1)$ and $[i]\sigma_2 = (v_2, l_2)$ then $v_1 = v_2$.*

PROOF. See Appendix A.1. \square

5.2 Microinstruction Lifecycle

Figure 2 depicts the microinstruction lifecycle under speculative execution. Compared to the OoO lifecycle of Section 4.2, states

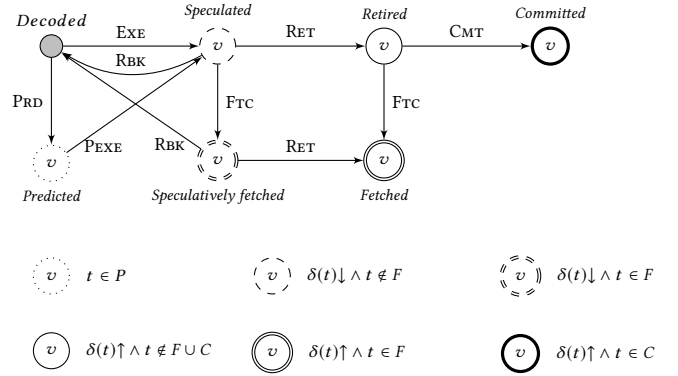


Figure 2: Speculative semantics: Microinstruction lifecycle

Decoded, Predicted, Speculated, and Speculatively Fetched correspond to state Decoded, state Retired corresponds to Executed, otherwise states Fetched and Committed are the same. As depicted in the legend, transitions between states set different properties of a microinstruction's lifecycle, which we will model in the semantics.

State Predicted (dotted circle) models microinstructions that have not yet been executed, but whose result values have been predicted. A Decoded microinstruction can transition to state Predicted by predicting its result value, thus recording that the value was predicted and causing the state of the microinstruction to be defined. A microinstruction that is ready to be executed (in Decoded), possibly relying on predicted values, can be executed and transition to state Speculated (dashed circle), recording its dependencies in the snapshot. Notice that state Speculated models both speculative and non-speculative execution of a microinstruction.

From state Speculated, a microinstruction can: (a) roll back to Decoded (if the predicted values were wrong); (b) speculatively fetch the next ISA instruction to be executed, thus moving to state Speculatively Fetched, doubled dashed circle) and generating newly decoded microinstructions; or (c) retire in state Retired (single circle) if it no longer depends on speculated values.

Microinstructions in state Speculatively Fetched can either be rolled back due to misspeculation, otherwise move to state Fetched (double circle). Finally, in state Retired, as in the OoO case, a PC store microinstruction can be (non-speculatively) fetched and generate newly decoded microinstructions, or, if it is a memory store, it can be committed to the memory subsystem (bold circle).

5.3 Microinstruction Semantics

We now present a speculative semantics, denoted by the transition relation $(\sigma, \delta, P) \rightsquigarrow (\sigma', \delta', P')$, that reflects the microinstructions' lifecycle in Figure 2. We illustrate the rules of our semantics using the graph in Example 4 and the interpretation of states (circles) in Figure 2. Additionally, for two microinstruction identifiers t and t' in speculative state $h = (I, s, C, F, \delta, P)$, we draw an edge from t to t' labeled with v whenever $\delta(t)(t') = v$.

(Predict) The semantics allows to predict the value of an internal operation choosing a value $v \in V$. The rule updates the storage and records the predicted name, while ensuring that the microinstruction has not been executed already.

$$(\text{PRD}) \frac{t \leftarrow c?e \in I \quad s(t) \uparrow \quad \delta' = \delta \cup \{t \mapsto \emptyset\}}{(I, s, C, F, \delta, P) \twoheadrightarrow (I, s[t \mapsto v], C, F, \delta', P \cup \{t\})}$$

We remark that the semantics can predict a value only for an internal operation ($t \leftarrow c?e$) that has not been already executed ($s(t) \uparrow$). As we will see, this choice does not hinder expressiveness while it avoids the complexity in modeling speculative execution of program counter updates and loads. Concretely, the rule assigns an arbitrary value to the name of the predicted microinstruction ($s[t \mapsto v]$) and records that the result is speculated ($\delta \cup \{t \mapsto \emptyset\}$). Observe that the snapshot $\delta'(t)$ is \emptyset because the prediction does not depend on the results of other microinstructions.

Consider state h_0 in Example 4 containing all microinstructions of our running program, which have just been decoded (gray circles). The CPU can predict that the value of arithmetic operation t_2 is 0. Rule PRD updates the storage with $t_2 \mapsto 0$ (dotted circle), the snapshot for t_2 with an empty mapping, and adds t_2 to the prediction set.

(Execute) The rules for execution, commit, and fetch reuse the OoO semantics. First for the case when the instruction has not been predicted already:

$$(\text{EXE}) \frac{\sigma \xrightarrow{I} \sigma' \quad \text{step-param}(\sigma, \sigma') = (\text{EXE}, t)}{(\sigma, \delta, P) \xrightarrow{I} (\sigma', \delta \cup \{t \mapsto s|_{\text{deps}(t, \sigma)}\}, P)}$$

The rule executes a microinstruction t using the OoO semantics and updates the snapshot δ , recording that the execution of t was determined by the value of its dependencies in $\text{deps}(t, \sigma)$ in storage s of state σ . Notice that the premise $\text{step-param}(\sigma, \sigma') = (\text{EXE}, t)$ ensures that microinstruction t has not been predicted. In fact, $\text{step-param}(\sigma, \sigma') = (\text{EXE}, t)$ only if $\sigma(t) \uparrow$, while rule PRD would update the storage with a value for name t , hence $t \notin P$.

Consider now the state h_2 resulting from the execution of t_1 and t_3 in Example 4. In h_2 the CPU can execute the PC update t_6 , updating the storage with $t_6 \mapsto 36$. The rule additionally updates the snapshot for t_6 with the current values of its dependencies, i.e., $\{t_2 \mapsto 0, t_3 \mapsto 32\}$. Since the executed microinstruction t_6 is a store, its dependencies are the free names occurring in the microinstruction. These snapshots are used by rules CMT and RBK to identify mispredictions. Similarly, the rule enables the execution of the memory store t_4 in h_3 , which updates the storage with $t_4 \mapsto 1$ and the snapshot for t_4 with the values of its dependencies $\{t_1 \mapsto 1\}$.

The following rule enables the execution of microinstructions whose result has been previously predicted:

$$(\text{PEXE}) \frac{\sigma = (I, s, C, F) \quad t \in P}{(I, s \setminus \{t\}, C, F) \xrightarrow{I} \sigma' \quad \text{step-param}(\sigma, \sigma') = (\text{EXE}, t)}{(\sigma, \delta, P) \xrightarrow{I} (\sigma', \delta \cup \{t \mapsto s|_{\text{deps}(t, \sigma)}\}, P \setminus \{t\})}$$

The rule removes the value predicted for t from the storage ($s \setminus \{t\}$) to enable the actual execution of t in the OoO semantics. It also removes t from the set of predicted names P and updates the snapshot with the new dependencies of t .

In our example, rule PEXE computes the actual value of t_2 in state h_4 , which was previously mispredicted as 0. The rule corrects the misprediction updating the storage with $t_2 \mapsto 1$ and the snapshot

for t_2 with the values of its dependencies, i.e., $t_1 \mapsto 1$. Notice that in case of a misprediction, the rule does not immediately roll back all other speculated microinstructions that are affected by the mispredicted values, e.g., t_6 .

(Commit) To commit a microinstruction it is sufficient to ensure that there are no dependencies left ($\delta(t) \uparrow$), i.e., the microinstruction has been retired. Since memory commits have observable side effects outside the processor pipeline, only retired memory stores can be sent to the memory subsystem.

$$(\text{CMT}) \frac{\sigma \xrightarrow{I} \sigma' \quad \text{step-param}(\sigma, \sigma') = (\text{CMT}(a, v), t) \quad \delta(t) \uparrow}{(\sigma, \delta, P) \xrightarrow{I} (\sigma', \delta, P)}$$

Consider the state h_4 and the memory store t_4 in our example. Since t_4 has not been retired (i.e., $\delta(t_4) = \{t_1 \mapsto 1\}$) it cannot be committed as $\delta(t_4) \downarrow$. By contrast, the commit of t_4 is allowed in state h_8 where $\delta(t_4) \uparrow$.

(Fetch) Finally, for the case of (speculative or non-speculative) fetching, the snapshot must be updated to record the dependency of the newly added microinstructions:

$$(\text{FRC}) \frac{\sigma \xrightarrow{I} \sigma' \quad \text{step-param}(\sigma, \sigma') = (F(I), t)}{(\sigma, \delta, P) \xrightarrow{I} (\sigma', \delta \cup \{t' \mapsto s|_{\{t\}} \mid t' \in I\}, P)}$$

Following the OoO semantics, if $\text{step-param}(\sigma, \sigma') = (F(I), t)$ then t is a PC update and $s(t)$ is the new value of the PC. For every newly added microinstruction in $t' \in I$, we extend the snapshot δ recording that t' was added as result of updating the PC microinstruction t with the value $s(t)$ (formally, we project the storage s on t , i.e., $s|_{\{t\}}$). The new snapshot may be used later to roll back the newly added microinstructions in I if the value of the PC is misspeculated.

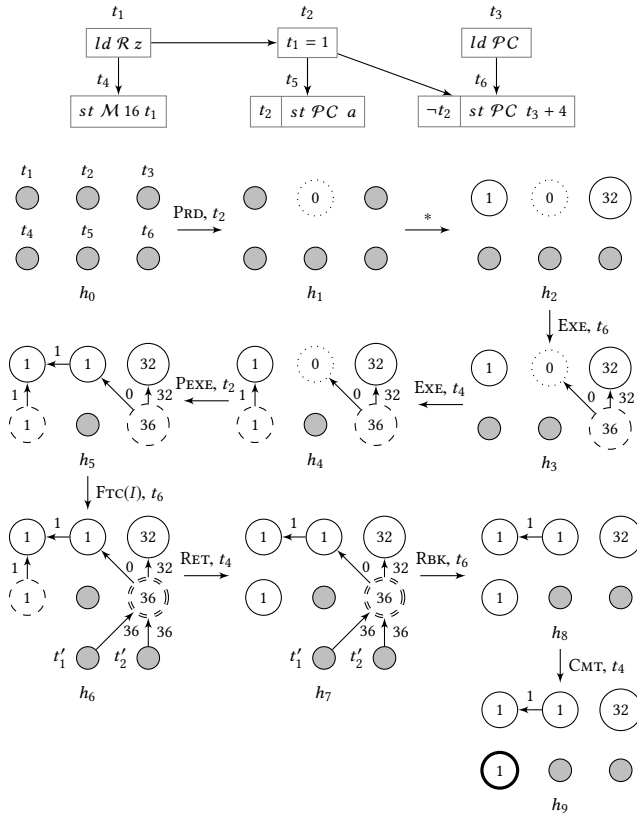
For example, in state h_5 the CPU can speculatively fetch the PC update t_6 , which sets the program counter to 36. Suppose that the newly added microinstructions in I (i.e., the microinstructions resulting from the translation of the ISA instruction at address 36) are t'_1 and t'_2 . Following the OoO semantics, I is added to existing microinstructions in σ' . The rule additionally updates the snapshot for t'_1 and t'_2 recording the PC store that generated the new microinstructions, i.e., $t_6 \mapsto 36$.

(Retire) The following transition rule allows to retire a microinstruction in case of correct speculation:

$$(\text{RET}) \frac{s(t) \downarrow \quad \text{dom}(\delta(t)) \cap \text{dom}(\delta) = \emptyset}{(I, s, C, F) \sim_t (I, \delta(t), C, F) \quad t \notin P} \frac{}{(I, s, C, F, \delta, P) \twoheadrightarrow (I, s, C, F, \delta \setminus \{t\}, P)}$$

The map $\delta(t)$ contains the snapshot of t 's dependencies at time of t 's execution. A microinstruction can be retired only if all its dependencies have been retired ($\text{dom}(\delta(t)) \cap \text{dom}(\delta) = \emptyset$), the microinstruction has been executed (i.e. its value has not been just predicted $s(t) \downarrow \wedge t \notin P$), and the snapshot of t 's dependencies is \sim_t equivalent with the current state, hence the semantics of t has been correctly speculated (see Lemma 5.3). Retiring a microinstruction results in removing the state of its dependencies from δ , as captured by $\delta \setminus \{t\}$.

For instance, in state h_6 the PC store t_6 cannot be retired for two reasons: one of its dependencies has not been retired (i.e., $\delta(t_6) =$



Example 4: Execution trace of speculative semantics.

$\{t_2 \mapsto 0, t_3 \mapsto 32\}$ and $\delta(t_2) = \{t_1 \mapsto 1\}$, hence $\text{dom}(\delta(t_6)) \cap \text{dom}(\delta) = \{t_2\}$, and the snapshot for t_6 differs with respect to the storage (i.e., $\delta(t_6)(t_2) \neq s(t_2)$). Instead, the microinstruction t_4 can be retired because its dependencies (i.e., t_1) have been retired (i.e., $\delta(t_1) \uparrow$) and the snapshot for t_4 (i.e., $t_1 \mapsto 1$) exactly matches the values in the storage. Notice that retiring t_4 would simply remove the mapping for t_4 from δ .

Notice that in case of a load, $(I, s, C, F) \sim_t (I, \delta(t), C, F)$ may hold even if some dependencies of t differ in s and $\delta(t)$. In fact, a load may have been executed as a result of misspeculating the address of a previous store. In this case, \sim_t implies that the misspeculation has not affected the calculation of str-act of the load (i.e., it does not cause a store bypass), hence there is no reason to re-execute the load. This mechanism is demonstrated in examples later in this section.

(Rollback) A microinstruction t can be rolled back when it is found to transitively reference a value that was wrongly speculated. This is determined by comparing t 's dependencies at execution time ($\delta(t)$) with the current storage assignment (s). In case of a discrepancy, if t is not a program counter store, the assignment to t can simply be undone, leaving speculated microinstructions t' that reference t to be rolled back later, if necessary.

$$(\text{RBK}) \frac{t \notin P \quad t \notin F \quad (I, s, C, F) \not\sim_t (I, \delta(t), C, F)}{(I, s, C, F, \delta, P) \twoheadrightarrow (I, s \setminus \{t\}, C, F, \delta \setminus \{t\}, P)}$$

However, if t is a program counter store, the speculative evaluation using rule FRC will have caused a new microinstruction to be speculatively fetched. This fetch needs to be undone. To that end let $t' < t$ (t' refers to t) if $t \in \text{dom}(\delta(t'))$, let $<^+$ be the transitive closure of $<$. As expected $<^+$ is antisymmetric and its reflexive closure is a partial order. Define then the set Δ^+ as $\{t' \mid t' <^+ t\}$; i.e., Δ^+ is the set of names that reference t , not including t itself. Finally, let $\Delta^* = \Delta^+ \cup \{t\}$.

$$(\text{RBK}) \frac{t \notin P \quad t \in F \quad (I, s, C, F) \not\sim_t (I, \delta(t), C, F)}{(I, s, C, F, \delta, P) \twoheadrightarrow (I \setminus \Delta^+, s \setminus \Delta^*, C, F \setminus \Delta^*, \delta \setminus \Delta^*, P \setminus \Delta^*)}$$

For example, in state h_7 the program counter update t_6 can be rolled back because $s(t_2) = 1 \neq 0 = \delta(t_6)(t_2)$. The transition moves the microinstruction t_6 back to the decoded state (i.e., the storage and snapshot h_8 are undefined for t_6) and removes every microinstruction that have been decoded by t_6 (i.e., t'_1 and t'_2).

Notice that rollbacks can be performed out of order and that loads can be retired even in case of mispredictions if their dependencies have been enforced. This permits to model advanced recovery methods used by modern processors, including concurrent and partial recovery in case of multiple mispredictions.

Speculation of load/store dependencies Since the predicted values of internal operations (cf. rule PRD) can affect conditions and targets of program counter stores, the speculative semantics supports speculation of control flow, as well as speculative execution of cross-dependencies resulting from prediction of load/store's addresses. We illustrate these features with an example (Figure 3), which depicts one possible execution of the program in Example 3.

Consider the state h_0 after the CPU has executed and retired first two microinstructions $t_{11}, t_{12}, t_{21}, t_{22}$, and t_{41} , thus resolving the first two stores and the load's address. In state h_0 the CPU can predict the address (i.e., the value of t_{31}) of the third store as 0 and modify the state as in h_1 (rule PRD).

This prediction enables speculative execution of the load t_{42} in state h_1 : the active store's bounded names $\text{bn}(\text{str-act}(\sigma_1, t_{42}))$ consist of the singleton set $\{t_{12}\}$, since $s_1(t_{21}) = s_1(t_{31}) = 0$, while $s_1(t_{41}) = 1$. Hence, we can apply rule EXE to execute t_{42} , thus updating the storage with $t_{42} \mapsto 1$, and recording the snapshot $\{t_{11} \mapsto 1, t_{21} \mapsto 0, t_{31} \mapsto 0, t_{41} \mapsto 1, t_{12} \mapsto 1\}$ for t_{42} . Concretely, t_{42} 's dependencies in state h_1 consists of the local dependencies (i.e., the load's address t_{41}), and the cross dependencies containing t_{12} (i.e., active store it loads the value from), as well as the potential sources of t_{42} , that is, the addresses of all stores between the active store t_{12} and the load t_{42} , namely t_{11}, t_{21} and t_{31} .

At this point, load t_{42} cannot be retired by rule RET in state h_2 since its dependencies, e.g., t_{31} , are yet to be retired. However, we can execute t_{31} by applying rule PEXE . The execution updates the state by removing t_{31} from the prediction set and storing its correct value, as well as extending the snapshot with $t_{31} \mapsto 0$, as depicted in state h_3 .

The execution of t_{31} enables the premises of rule RBK to capture that the dependency misspeculation led to misspeculation of the address of the load t_{42} . Specifically, the set asn at the time of t_{42} 's execution $\text{bn}(\text{str-act}((I_3, \delta_3(t_{42}), C_3, F_3), t_{41})) = \{t_{12}\}$ differs from the active store set $\text{bn}(\text{str-act}(\sigma_3, t_{41})) = \{t_{32}\}$ in the current state.




Figure 3: Speculation of load/store dependencies

Therefore, we roll back the execution removing the mappings for t_{42} from the storage and the snapshot as in h_4 .

Finally, we remark that the speculative execution of loads is rolled back *only if* a misprediction causes a violation of load/store dependencies. For instance, if the value of t_{31} was 5 instead of 1, as depicted in h'_3 , the misprediction of t_{31} 's value as 0 in h_1 does not enable a rollback of the load. This is because the actual value of t_{31} does not change the set of active stores. In fact, the set of active stores at the time of t_{42} 's execution $bn(str-act(I'_3, \delta'_3(t_{42}), C'_3, F'_3), t_{41})) = \{t_{12}\}$ is the same as the active store's set $bn(str-act(\sigma'_3, t_{41})) = \{t_{12}\}$ in the current state.

6 IN-ORDER SEMANTICS

We define the in-order (i.e., sequential) semantics by restricting the scheduling of the OoO semantics and enforcing the execution of microinstructions in program order.

A microinstruction $\iota = t \leftarrow c?o$ is *completed* in state σ (written $C(\sigma, \iota)$) if one of the following conditions hold:

- The instruction's guard evaluates to false in σ , i.e. $\neg[c](\sigma)$.
- The instruction has been executed and is not a memory or a program counter store, i.e., $o \neq st \ \mathcal{M} \ t_a \ t_v \wedge o \neq st \ \mathcal{PC} \ t_v \wedge \sigma(t) \downarrow$.
- The instruction is a committed memory store or a fetched and decoded program counter store, i.e., $t \in C \cup F$

The in-order transition rule allows an evaluation step to proceed only if program-order preceding microinstructions have been completed.

$$\frac{\sigma \xrightarrow{l} \sigma' \quad \text{step-param}(\sigma, \sigma') = (\alpha, t) \quad \forall \iota \in \sigma \text{ if } bn(\iota) < t \text{ then } C(\sigma, \iota)}{\sigma \xrightarrow{l} \sigma'}$$

It is easy to show that the sequential model is deterministic. In fact, the OoO model allows each transition to modify one single name t , while the precondition of the in-order rule forces all previous instructions to be completed, therefore only one transition at a time is enabled.

Definition 6.1. Let $\sigma_1 :: \dots :: \sigma_n$ be the sequence of states of execution π , then $commits(\pi, a)$ is the list of memory commits at address a in π , and is empty if $n < 2$; $v :: commits(\sigma_2 :: \dots :: \sigma_n, a)$ if $\text{step-param}(\sigma_1, \sigma_2) = (CMT(a, v), t)$; and $commits(\sigma_2 :: \dots :: \sigma_n, a)$ otherwise.

We say that two models are memory consistent if writes to the same memory location are seen in the same order.

Definition 6.2. The transition systems \rightarrow_1 and \rightarrow_2 are *memory consistent* if for any program and initial state σ_0 , for all executions $\pi = \sigma_0 \rightarrow_1^* \sigma$, there exists $\pi' = \sigma_0 \rightarrow_2^* \sigma'$ such that for all $a \in \mathcal{M}$ $commits(\pi, a)$ is a prefix of $commits(\pi', a)$.

Intuitively, two models that are memory consistent yield the same sequence of memory updates for each memory address. This ensures that the final result of a program is the same in both models. Notice that since we do not assume any fairness property for the transition systems then an execution $\pi \rightarrow_1$ may indefinitely postpone the commits for a given address. For this reason we only require to find an execution such that $commits(\pi, a)$ is a prefix of $commits(\pi', a)$. We obtain memory consistency of both the OoO and the speculative semantics against the in-order semantics.

THEOREM 6.3. \rightarrow and \rightarrow are memory consistent. \square

PROOF. See Appendix A.2. \square

THEOREM 6.4. $\rightarrow\!\!\rightarrow$ and \rightarrow are memory consistent. \square

PROOF. See Appendix A.3. \square

7 ATTACKS AND COUNTERMEASURES

InSpectre can be used to model and analyze (combinations of) microarchitectural features underpinning Spectre attacks [9, 31, 37], and, importantly, to discover new vulnerabilities and to reason about the security of proposed countermeasures. Observe that these results hold for our generic microarchitectural model, while specific CPUs would require instantiating InSpectre to model their microarchitectural features. We remark that real-world feasibility of our new vulnerabilities falls outside the scope of this work.

Specifically, we use the following recipe: We model a specific prediction strategy in InSpectre and try to prove conditional non-interference for arbitrary programs. Failure to complete the security proof results in new classes of counterexamples as we report below.

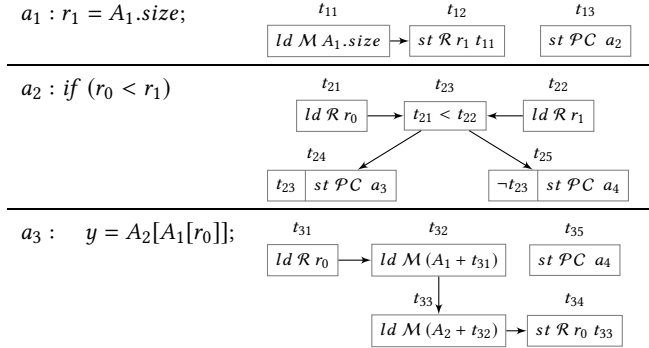
Concretely, prediction strategies and countermeasures are modeled by constraining the nondeterminism in the microinstruction scheduler and in the prediction semantics (see rule PRD). The prediction function $pred_p : \Sigma \rightarrow N \rightarrow 2^V$ captures a prediction strategy p by computing the set of predicted values for a name $t \in N$ and a state $\sigma \in \Sigma$. We assume the transition relation satisfies the following property: If $(\sigma, \delta, P) \xrightarrow{l} (\sigma', \delta', P \cup \{t\})$ then $t \in \text{dom}(pred_p(\sigma))$ and $\sigma'(t) \in pred_p(\sigma)(t)$. This property ensures

that the transition relation chooses predicted values from function $pred_p$.

Following the security model in Section 2, we check conditional noninterference by: (a) using the in-order transition relation \rightarrow as reference model and speculative (OoO) transition relation \twoheadrightarrow (\rightarrow) as target model; (b) providing the security policy \sim for memory and registers. To invalidate conditional noninterference it is sufficient to find two \sim -indistinguishable states that yield the same observations in the reference model and different observations in the target model. We use the classification by Canella et al. [9] to refer to existing attacks.

7.1 Spectre-PHT

Spectre-PHT [31] exploits the prediction mechanism for the outcome of conditional branches. Modern CPUs use *Pattern History Tables* (PHT) to record patterns of past executions of conditional branches, i.e., whether the *true* or the *false* branch was executed, and then use it to predict the outcome of that branch. By poisoning the PHT to execute one direction (say the *true* branch), an attacker can fool the prediction mechanism to execute the *true* branch, even when the actual outcome of the branch is ultimately *false*. The following program (and the corresponding MIL) illustrates information leaks via Spectre-PHT:



Suppose the security policy labels as public the data in arrays A_1 and A_2 , and in register r_0 , and that the attacker controls the value of r_0 . This program is secure at the ISA level as it ensures that r_0 always lies within the bounds of A_1 . However, an attacker can fool the prediction mechanism by first supplying values of r_0 that execute the *true* branch, and then a value that exceeds the size of A_1 . This causes the CPU to perform an out-of-bounds memory access of sensitive data, which is later used as index for a second memory access of A_2 , thus leaving a trace into the cache.

Branch prediction predicts values for MIL instructions that block the evaluation of the guard of a PC store whose target address has been already resolved. For $\sigma = (I, s, C, F, \delta, P)$, we model it as:

$$pred_{br}(\sigma) = \{t' \mapsto v \mid t \leftarrow c?st \mathcal{PC} \ t_a \in I \wedge t' \in fn(c) \wedge s(t_a) \downarrow\}$$

Let σ_0 be the state where only the instruction in a_1 has been translated. Then $pred_{br}(\sigma_0)$ is empty, since σ_0 contains a single unconditional PC update (the guard of t_{13} has no free names). The CPU may apply rules EXE , RET , and FTC on t_{13} without waiting the result of t_{11} . This leads to a new state σ_1 which is obtained by updating the storage with $s_1 = \{t_{13} \mapsto a_2\}$, extending the microinstructions' list with the translation of a_2 , and the snapshot with

$\delta_1 = \{t_i \mapsto t_{13} \mapsto a_2 \text{ for } 1 \leq i \leq 5\}$, while producing the observation $il \ a_2$. In this state $pred_{br}(\sigma_1) = \{t_{23} \mapsto 0, t_{23} \mapsto 1\}$ since the conditions of the two PC stores (i.e., t_{24} and t_{25}) depend on t_{23} which is yet to be resolved. The CPU can now apply rule PRD using the prediction $t_{23} \mapsto 1$, thus guessing that the condition is true. The new state σ_2 contains $s_2 = s_1 \cup \{t_{23} \mapsto 1\}$, $\delta_2 = \delta_1$, and $P_2 = \{t_{23}\}$.

The CPU can follow the speculated branch by applying rules EXE and FTC on t_{24} , which results in state σ_3 with $s_3 = s_2 \cup \{t_{24} \mapsto a_3\}$, $\delta_3 = \delta_2 \cup \{t_{24} \mapsto \{t_{23} \mapsto 1\}, t_{3i} \mapsto \{t_{24} \mapsto a_3\} \text{ for } 1 \leq i \leq 5\}$, and $F_3 = \{t_{13}, t_{24}\}$. Additionally, it produces the observation $il \ a_3$.

Applying rule EXE on t_{31} and t_{32} results in a buffer overread and produces state σ_4 with $s_4 = s_3 \cup \{t_{31} \mapsto r_0, t_{32} \mapsto A_1[r_0]\}$, and observation $dl \ r_0$. Similarly, rule EXE on t_{33} produces the observation $dl \ A_2 + A_1[r_0]$.

Clearly, if $r_0 \geq A_1.size$, the observation reveals memory content outside A_1 , allowing an attacker to learn sensitive data. Observe that this is rejected by the security condition, since such observation is not possible in the sequential semantics.


7.1.1 Countermeasure: Serializing Instructions. Serializing instructions can be modeled by constraining the scheduling of microinstructions. For example, we can model the Intel's *lfence* instruction via a function $lfence(I)$ that extracts all microinstructions resulting from the translation of *lfence*.

Concretely, for $\sigma = (I, s, C, F, \delta, P)$, $t \in lfence(I)$ and $\sigma \twoheadrightarrow \sigma'$, it holds that: (i) if $\sigma(t) \uparrow$ and $\sigma'(t) \downarrow$ then for each $t' \leftarrow c?ld \ M \ t_a \in \sigma$ such that $t' < t \ fn(c) \subseteq \text{dom}(s) \setminus \text{dom}(\delta)$ and $c \Rightarrow (\sigma(t') \downarrow \wedge \delta(t') \uparrow)$; (ii) if $\sigma(t') \uparrow$, $\sigma'(t') \downarrow$, $t' > t$, and $t' \leftarrow c?ld \ M \ t_a \in \sigma$, or $t' \leftarrow c?st \ M \ t_a \ t_v \in \sigma$, or $t' \leftarrow c?st \ R \ t_a \ t_v \in \sigma$, then $\sigma(t) \downarrow \wedge \delta(t) \uparrow$.

Intuitively, the conditions restrict the scheduling of microinstructions to ensure that: (i) whenever a fence is executed, all previous loads have been retired, and (ii) subsequent memory operations or register stores can be executed only if the fence has been retired.

In order to reduce the performance overhead, several works (e.g. [46]) use static analysis to identify necessary serialization points in a program. In the previous example, it is sufficient to place *lfence* after t_{32} and before t_{33} . This does not prevent the initial buffer overread of t_{32} , however, it suspends t_{33} until t_{32} is retired. In case of misprediction, t_{32} and t_{33} will be rolled back, preventing the observation $dl \ A_2 + A_1[r_0]$ which causes the information leak.

7.1.2 Countermeasure: Implicit Serialization. An alternative countermeasure to prevent Spectre-PHT is to use instructions that introduce implicit serialization [18, 41]. For instance, adding the following gadget between instructions a_2 and a_3 in the previous example prevents Spectre-PHT on existing Intel CPUs:




Intuitively, this gadget forces mispredictions to always access $A_1[0]$,

Consider the extension of the previous example with the gadget and suppose $pred_{br}$ mispredicts $t_{23} \mapsto 1$. The instruction in a_3'' introduces a data dependency between t_{11} and t_{32} since $str-act$ of t_{31} includes t_{32}'' until t_{31}'' has been executed; $str-act$ of t_{31}'' includes t_{34}' ; and $str-act$ of t_{32}' includes t_{12} . These names (and intermediate intra-instruction dependencies) are in the free names of some condition of a PC store, hence they cannot be predicted by $pred_{br}$ and their dependencies are enforced by the semantics. In particular, when t_{23} is mispredicted as 1, t_{32}'' is executed after that t_{11} has obtained the value from the memory. This ensures that t_{32}'' sets r_0 to 0 every time a buffer overread occurs. Therefore misspeculations generate the observations $dl A_1 + 0$ and $dl A_2 + A_1[0]$, which do not violate the security condition (since A_1 is labeled as public).

7.1.3 New Vulnerability: Spectre-PHT ICache. When the first Spectre attack was published, some microarchitectures (e.g., Cortex A53) were claimed immune to the attack because of “allowing speculative fetching but not speculative execution” [5]. The informal argument was that mispredictions cannot cause buffer overreads or leave any footprint on the cache in absence of speculative loads. To check this claim, we constrain the semantics to only allow speculation of PC values. Specifically, we require for any transition $(\sigma, \delta, P) \rightarrow (\sigma', \delta', P')$ that executes a microinstruction ($step-param(\sigma, \sigma') = (EXE, t)$) which is either a load ($t \leftarrow c?ld \tau t_a \in \sigma$) or a store ($t \leftarrow c?st \tau t_a t_v \in \sigma$) of a resource other than the program counter ($\tau \neq PC$) to have an empty snapshot on past microinstructions ($dom(\delta) \cap \{t' \mid t' < t\} = \emptyset$).

The analysis of conditional noninterference for this model led to the identification of a class of counterexamples, which we call Spectre-PHT ICache, where branch prediction causes leakage of sensitive data via an ICache disclosure gadget.

Consider a program that jumps to the address pointed to by sec if a user has admin privileges, otherwise it continues to address a_3 .



In the sequential model, an attacker that only observes the instruction cache can see the sequence of observations $il a_1 :: il a_2 :: il a_2$ if $*admin \neq 1$, otherwise the sequence $il a_1 :: il a_2 :: il sec$.

A CPU that supports only speculative fetching may first complete all microinstructions in a_1 , and then predict the result of t_{23} to enable the execution of t_{25} . As a result the PC speculatively fetches the instruction at location sec although $*admin \neq 1$. The transition yields the observation sequence $il a_1 :: il a_2 :: il sec$ which was not possible in the sequential model, thus violating the security condition and leaking the value of sec via the instruction cache.

Intel’s ifence does not stop all microarchitectural operations, like instruction fetching. For this reason ifence may be ineffective against leakage via ICache. In fact, InSpectre reveals that placing a ifence between t_{21} and t_{23} does not prevent the leakage: t_{22}, t_{23}, t_{25}

are neither memory operations nor register stores, hence they can be speculated before the execution of the ifence.


7.2 Spectre-BTB and Spectre-RSB

Two variants of Spectre attacks [9] exploit a CPU’s prediction mechanism for jump targets to leak sensitive data. In particular, Spectre-BTB [31] (Branch Target Buffer) poisons the prediction of indirect jump targets. To model this prediction strategy we assume a function $ijmps(I)$ that extracts all PC stores resulting from the translation of indirect jumps. This can be accomplished by making the translation of these instructions syntactically distinguishable from other control flow updates. As a result, prediction is possible for all indirect jumps whose address is yet to be resolved: Namely, $pred_{BTB}(I, s, C, F, \delta, P) =$

$$\{t_a \mapsto v \mid t \leftarrow c?st \mathcal{PC} t_a \in ijmps(I) \wedge s(t_a) \uparrow\}$$


We do not restrict the possible predicted values v , since an accurate model of jump prediction requires knowing the strategy used by the CPU to update the BTB buffer.

Spectre-RSB [33, 37] poisons the Return Stack Buffer (RSB), which is used to temporally store the N most recent return addresses: $call$ instructions push the return address on the RSB, while ret instructions pop from the RSB to predict the return target. A misprediction can happen if: (i) a return address on the stack has been explicitly overwritten, e.g., when a program handles a software exception using $longjmp$ instructions, or, (ii) returning from a call stack deeper than N , the RSB is empty and the CPU uses the same prediction as for the other indirect jumps. We model $call$ and ret instructions via program counter stores. A $call$ to address b_1 from address a_1 can be modeled as



The $call$ instruction saves (e.g. t_{15}) the return address (e.g. $a_1 + 4$) into the stack, decreases the stack pointer (e.g. t_{13}), and jumps to address b_1 (e.g. t_{16}).

A ret instruction from address a_2 can be modeled as



The instruction loads the return address from the stack (t_{24}), increases the stack pointer (t_{23}), and returns (t_{26}).

We assume functions $calls(I)$ and $rets(I)$ to extract the PC stores that belong to a $call$ and ret respectively. Moreover, if $t \in bn(calls(I))$, we use $ret-ra(I, t)$ to retrieve name of the microinstruction that saves the return address (e.g. t_{15}) of the corresponding $call$. We model return address prediction as

$$pred_{RSB}(I, s, C, F, \delta, P) = \{t_a \mapsto v \mid t \leftarrow c?st \mathcal{PC} t_a \in rets(I) \wedge s(t_a) \uparrow \wedge \exists t' \in bn(calls(I)). t' < t \wedge s(ret-ra(I, t')) = v \wedge RSB-depth(I, t', t) \subseteq \{1 \dots N\}\}$$


Prediction is possible only for ret microinstructions t that have a prior matching $call$ t' , provided that the size of intermediary

stack depth is between 1 and N . We define the latter as the set $RSB\text{-depth}(I, t', t) = \{\#(bn(calls(I)) \cap \{t' \dots t''\}) - \#(bn(rets(I)) \cap \{t' \dots t''\}) \mid t' \leq t'' < t\}$, where $\{t' \dots t''\}$ is an arbitrary continuous sequence of names starting from t' and ending before t'' , and $\#(bn(calls(I)) \cap \{t' \dots t''\})$ and $\#(bn(rets(I)) \cap \{t' \dots t''\})$ count the number of calls and rets in the sequence respectively. The prediction consists in assuming the target address (e.g. t_a) of the ret to be equal to the return address (e.g. v) that has been pushed into the stack by the matching call.

In some microarchitectures (e.g., [?]), RSB prediction falls back to the BTB in case of underflow, i.e., when returning from a function with nested stack deeper than N . This case can be modeled by considering the prediction strategy $pred_{RSB/BTB}$ defined as

$$pred_{RSB} \cup \{t_a \mapsto v \mid t \leftarrow c?st \mathcal{P}C \ t_a \in rets(I) \wedge s(t_a) \uparrow \wedge \exists t' \in bn(calls(I)). t' < t \wedge RSB\text{-depth}(I, t', t) = \{1 \dots N'\} \wedge N' > N\}$$

The following example shows how jump target prediction may violate the security condition. Consider the program $*p = \&f; (*p)()$ that saves the address of a function (i.e., $\&f$) in a function pointer at constant address p and immediately invokes the function. Assuming that these instructions are stored at addresses a_1 and a_2 , their MIL translation is:




Example 5: $*p = \&f; (*p)()$

Because our semantics can predict only internal operations (see rule PRD), the translation function introduces an additional internal operation, i.e., t_{22} which allows predicting the value of the load t_{21} .

Suppose that the function f simply returns and the security policy labels all data, except the program counter, as sensitive. The program is secure (at the ISA level) as it always transfers control to f , producing the sequence of observations $il a_1 :: ds p :: il a_2 :: dl p :: il \&f$ independently of the initial state.

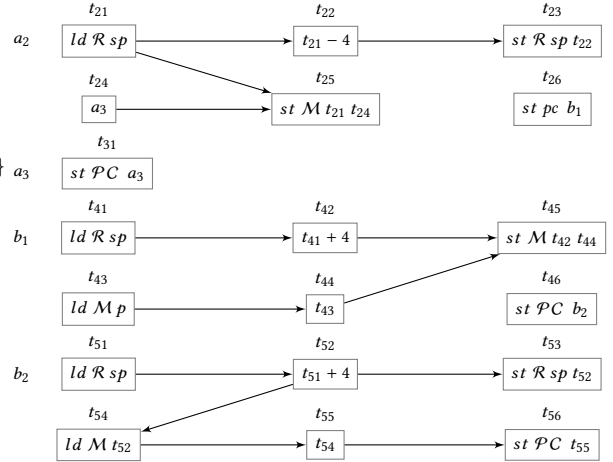
Jump target prediction produces a different behavior. Let σ_0 be the state containing only the translation of the instruction in a_1 . Initially, $pred_{BTB}(\sigma_0)$ is empty since the state contains no PC updates (e.g. t_{12}) that result from translating indirect jumps. The CPU may execute and fetch t_{12} , thus adding t_{21} , t_{22} , and t_{23} to the set of microinstructions I . In the resulting state $pred_{BTB}$ is $\{t_{22} \mapsto v \mid v \in V\}$, since t_{23} models an indirect jump and t_{22} has not been executed. The CPU can therefore predict the value of t_{22} without waiting for the result of the load t_{21} . If the predicted value is the address g of the instruction $r_1 := *(r_2)$ the misprediction can use g as gadget to leak sensitive information.



In fact, the speculative semantics can produce the sequence of observations $il a_1 :: ds p :: il a_2 :: dl p :: il g :: dl v$, where v is the initial value of register r_2 . The last observation of the sequence allows an attacker to learn sensitive data. Observe that this leak is

readily captured by the security condition, since such observation sequence is not possible in the sequential semantics.

7.2.1 Countermeasure: Retpoline. A known countermeasure to Spectre-BTB is the *Retpoline* technique developed by Google [51]. In a nutshell, retpolines are instruction snippets that isolate indirect jumps from speculative execution via call and return instructions. Retpoline has the effect of transforming indirect jumps at address a_2 of Example 5 as:



Instruction at a_2 calls a *trampoline* starting at address b_1 and instruction at a_3 loops indefinitely. The first instruction of the trampoline overwrites the return address on the stack with the value at address p and its second instruction at b_2 returns.

We leverage our model to analyze the effectiveness of Retpoline for indirect jumps. Since address b_1 is known at compile time, t_{26} does not trigger a jump target prediction. While executing the trampoline, the value of t_{55} may be mispredicted, especially if the load from p has not been executed and the store t_{45} is postponed. However, b_2 is a ret, hence the value of t_{55} is predicted via $pred_{RSB}$. Since there is no call between a_1 and b_2 , then prediction can only assign the address a_3 to t_{55} (i.e., $pred_{RSB}|_{t_{55}} \subseteq \{t_{55} \rightarrow a_3\}$). Therefore, the RSB entry generated by a_2 is used and mispredictions are captured with the infinite loop in a_3 . Ultimately, when the value of t_{55} is resolved, the correct return address is used and the control flow is redirected to the value of $*p$, as expected.

7.3 Spectre-STL

Spectre-STL [26] (Store-To-Load) exploits the CPUs mechanism to predict load-to-store data dependencies. A load cannot be executed before executing all the past (in program order) stores that affect the same memory address. However, if the address of a past store has not been resolved, the CPU may execute the load in speculation without waiting for the store, predicting that the target address of the store is different from the load's address. Mispredictions cause store bypasses leading to information leaks and access to stale data. This behavior can be modeled as $pred_{STL}(\sigma, \delta, P) =$

$$\left\{ \begin{array}{l} t' \leftarrow c?ld \ M \ t'_a \in \sigma \wedge \sigma(t'_a) \neq a \\ t_a \mapsto a \mid t \leftarrow c?st \ M \ t_a \ t_v \in str\text{-act}(\sigma, t') \wedge \\ \sigma(t_a) \uparrow \end{array} \right\}$$

A prediction occurs whenever a memory store (t) is waiting an unresolved address ($\sigma(t_a)\uparrow$), while the address ($s(t'_a)$) of a subsequent load (t') has been resolved, and the load may depend on the store ($t \in \text{bn}(\text{str-act}(\sigma, t'))$). Prediction guesses that the store's address (t_a) differs with the load's address.


7.3.1 Hardware countermeasures to Store Bypass. The specification of proposed hardware countermeasures oftentimes comes with no precise semantics and is ambiguous. ARM introduced the Speculative Store Bypass Safe (SSBS) configuration to prevent store bypass vulnerabilities. The specification of SSBS [4] is: *Hardware is not permitted to load ... speculatively, in a manner that could ... give rise to a ... side channel, using an address derived from a register value that has been loaded from memory ... (L) that speculatively reads an entry from earlier in the coherence order from that location being loaded from than the entry generated by the latest store (S) to that location using the same virtual address as L.*

InSpectre provides a ground to formalize the behavior of these hardware mechanisms. We formalize SSBS as follows. Let $\sigma = (I, s, C, F, \delta, P)$ and $t \leftarrow c?ld \tau t_a \in \sigma$. If $\sigma \xrightarrow{t} \sigma'$, $\sigma(t)\uparrow$, and $\sigma'(t)\downarrow$, then for every $t' \in \text{srcs}(t, \sigma)$, if $\sigma(t') \neq \sigma(t_a)$ then $t' \notin P$.

The reason why SSBS prevents Spectre-STL is simple. The rule forbids the execution of a load t if any address used to identify the last store affecting t_a has been predicted to differ from t_a .

7.3.2 New Vulnerability: Spectre-STL-D. Our model reveals that if a microarchitecture mispredicts the existence of a Store-To-Load Dependency (hence Spectre-STL-D), e.g., in order to forward temporary store results, a similar vulnerability may be possible. To model this behavior it is enough to substitute $\sigma(t'_a) \neq a$ with $\sigma(t'_a) = a$ in pred_{STL} . We consider this a new form of Spectre because the implementation of this microarchitectural feature can be substantially different from the one required for Spectre-STL (e.g., Feiste et al. [17] patented a mechanism to implement this feature) and because the vulnerable programs are different.

This feature may cause Spectre-STL-D if a misspeculated dependency is used to perform subsequent memory accesses. Consider the following program:



If the CPU executes and fetches t_{14} , predicts that $t_{12} = b_2$ (i.e., it mispredicts the alias $*b_1 == *b_2$), executes t_{13} , forwards the result of t_{13} to t_{21} , and executes t_{22} before the load t_{11} is retired, then the address accessed by t_{22} depends on t_{sec} . This can produce the secret-dependent sequence of observations $il a_1 :: il a_2 :: dl sec$, while the sequential semantics always produces the secret-independent sequence of observations $il a_1 :: dl b_1 :: ds *b_1 :: il a_2 :: dl b_2 :: dl *b_2$. Notice that SSBS may not be effective against Spectre-STL-D.

7.4 New Vulnerability: Spectre-OoO

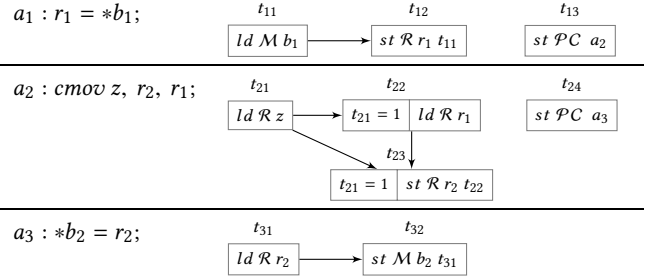
A popular countermeasure to prevent sensitive data from affecting the execution time and caches is “constant time programming”,

also known as “data oblivious algorithms”. This mechanism ensures that branch conditions and memory addresses are independent of sensitive data. The following definition formalizes “ISA constant time” while abstracting from the specific ISA:

Definition 7.1. A program is “ISA constant time” if for every pair of states $\sigma_1 \sim \sigma_2$ and every pair of in-order executions of length n , $\sigma_1 \rightarrow^n \sigma'_1$ and $\sigma_2 \rightarrow^n \sigma'_2$, it is the case that $\sigma'_1 \approx_{ISA} \sigma'_2$, where $(I, s, C, F) \approx_{ISA} (I', s', C', F')$ iff

- (1) $I = I', C = C', F = F'$: the sets of microinstructions, commits and decodes are equal.
- (2) If $t \leftarrow c?ld M t_a \in I$ or $t \leftarrow c?st M t_a t_v \in I$ then $[c]\sigma = [c]\sigma'$, $s(t)\downarrow = s'(t)\downarrow$, (whenever defined, guards evaluate the same, and memory operations execute in lockstep) and $[c]\sigma \Rightarrow (\sigma(t_a) = \sigma'(t_a))$ (the same values are used to address memory)
- (3) If $t \leftarrow c?st PC t_v \in I$ then $[c]\sigma = [c]\sigma'$, $s(t)\downarrow = s'(t)\downarrow$, and $[c]\sigma \Rightarrow (\sigma(t_v) = \sigma'(t_v))$ (the same values are used to update the PC)

The following program (and its MIL translation) exemplifies this policy. It loads register r_1 from address b_1 , copies the value of r_1 in r_2 if the flag z is set, and saves the result into b_2 .



Suppose that flag z contains sensitive information and the attacker observes only the data cache. The “conditional move” instruction in a_2 executes in constant time [27] and is used to re-write branches that may leak information via the execution time or the instruction cache. This allows the program to always access address b_1 and b_2 unconditionally and execute always the same ISA instructions: In the sequential model the program always produces the sequence of observations $dl b_1 :: ds b_2$.

Programs that are ISA constant time could be insecure in presence of speculation, as demonstrated by Spectre-PHT [31]. Perhaps surprisingly, it turns out that ISA constant time is not secure even for the OoO model, in absence of speculation. In fact, our analysis of conditional noninterference for ISA constant time programs in the OoO model led to the identification of a class of vulnerable programs, where secrets influence the existence of data dependency between registers. The above program exemplifies this problem: the data dependency between t_{11} and t_{32} exists only if z is set. Concretely, consider two states σ_0 and σ_1 in which $z = 0$ and $z = 1$, respectively. Then, $\text{str-act}(\sigma_1, t_{31}) = \{t_{23}\}$ and $\text{str-act}(\sigma_1, t_{23}) = \{t_{12}\}$, while $\text{str-act}(\sigma_0, t_{31})$ is the microinstruction representing the initial value of r_2 . Therefore, state σ_0 may produce the observation sequence $ds b_2 :: dl b_1$ only if the flag $z = 0$, thus leaking its value through the data cache.

7.4.1 *MIL Constant Time.* Spectre-OoO motivates the need for a new microarchitecture-aware definition of constant time.

Definition 7.2. A program is “MIL constant time” if for every pair of states $\sigma_1 \sim \sigma_2$ and every pair of in-order executions of length n , $\sigma_1 \rightarrow^n \sigma'_1$ and $\sigma_2 \rightarrow^n \sigma'_2$, it is the case that $\sigma'_1 \approx_{MIL} \sigma'_2$, where $(I, s, C, F) \approx_{MIL} (I', s', C', F')$ iff


- (1) $(I, s, C, F) \approx_{ISA} (I', s', C', F')$
- (2) If $t \leftarrow c?ld \mathcal{R} t_a \in I$ or $t \leftarrow c?st \mathcal{R} t_a t_v \in I$ then $[c]\sigma = [c]\sigma'$, $s(t) \downarrow = s'(t) \downarrow$, and $[c]\sigma \Rightarrow (\sigma(t_a) = \sigma'(t_a))$

Notice that in addition to standard requirements of constant time, MIL constant time requires that starting from two \sim -indistinguishable states the program makes the same accesses to registers. MIL constant time is sufficient to ensure security in the OoO model:

THEOREM 7.3. *If a program P is MIL constant time then P is conditionally noninterferent in the OoO model.*

The theorem enables the enforcement of conditional noninterference for the OoO model by verifying MIL constant time in the sequential model. This strategy has the advantage of performing the verification in the sequential model, which is deterministic, thus making it easier to reuse existing tools for binary code analyses [6].

Finally, we remark that MIL constant time is microarchitecture aware. This means that the same ISA program may or may not satisfy MIL constant time when translated to a given microarchitecture. In fact, the MIL translation of conditional move above is not MIL constant time because of the dependency between the sensitive value in t_{21} and conditional store in t_{23} . However, if a microarchitecture translates the same conditional move as below, the translation is clearly MIL constant time.



8 RELATED WORK

Speculative semantics and foundations Several works have recently addressed the formal foundations of specific forms of speculation to capture Spectre-like vulnerabilities. Cheang et al. [12], Guarnieri et al. [22], and Mcilroy et al. [40] propose semantics that support branch prediction, thus modeling only Spectre v1. Neither work supports speculation of target address, speculation of dependencies, or OoO execution. Disselkoen et al. [14] propose a pomset-based semantics that supports OoO execution and branch prediction. Their model targets a higher abstraction level modeling memory references using logical program variables. Hence, the model cannot support dynamic dependency resolution, dependency prediction, and speculation of target addresses.

Like us, Cauligi et al. [11] propose a model that captures existing variants of Spectre and independently discover a vulnerability similar to our Spectre-STL-D. Remarkably, they demonstrate the feasibility of the attack on Intel Broadwell and Skylake processors. A key difference between the two models is that Cauligi et al. impose sequential order to instruction retire and memory stores. While simplifying the proof of memory consistency and verification, it

does not reflect the inner workings of modern CPUs, which reorder memory stores and implement a relaxed consistency model. These features are required to capture Spectre-OoO in Section 7.4. Moreover, our model provides a clean separation between the general speculative semantics and microarchitecture-specific features, where the latter is obtained by reducing the nondeterminism of the former. This enables a modular analysis of (combinations of) predictive strategies, as in Spectre-PHT ICACHE in Section 7.1.3.

Cache side channels In line with prior works [11, 12, 22], our attacker model abstracts away the mechanism used by an attacker to profile the sequence of a victim’s memory accesses, providing a general account of trace-driven attacks [45]. Complementary works [15, 20, 36, 58] show that cache profiling is becoming increasingly steady and precise. Performance jitters caused by cache usage have been widely exploited to leak sensitive data [2, 21, 32, 39, 43, 44, 59], e.g., in cryptography software. Miller [41], and Fogh and Ertl [18] propose a taxonomy for mitigating speculative execution vulnerabilities. We refer to a recent survey by Canella et al. [19] on cache-based countermeasures.

Spectre vs Meltdown Recent attacks that use microarchitectural effects of speculative execution have been generally distinguished as Spectre and Meltdown attacks [9]. We focus on the former [13, 16, 23, 24, 30, 31, 34, 37, 57], which exploits speculation to cause a victim program to transiently access sensitive memory locations that the attacker is not authorized to read. Meltdown attacks [35] transiently bypass the hardware security mechanisms that enforce memory isolation. Importantly, Meltdown attacks can be easily countered in hardware, while Spectre attacks require hardware-software co-design, which motivates our model. We remark that the vulnerability in Section 7.3 is different from the recent Microarchitectural Data Sampling attacks [10, 48, 52], since it only requires the CPU to predict memory aliases with no need of violating memory protection mechanisms. Microarchitectures supporting this feature have been proposed, e.g., in Feiste et al. [17].

Tool support Several prototypes have been developed to reproduce and detect known Spectre-PHT attacks [12, 22, 53, 54]. Checkmate [50] synthesizes proof-of-concept attacks by using models of speculative and OoO pipelines. Tool support for vulnerabilities beyond Spectre-PHT requires dealing with a large number of possible predictions and instruction interleavings. In fact, current tools mainly focus on Spectre-PHT ignoring OoO execution.

Functional Pipeline Correctness A number of authors, cf. [1, 8, 28, 38, 47], have studied the orthogonal problem of functional correctness in the context of concrete pipeline architectures involving features such as OoO and speculation, usually using a complex refinement argument based on Burch-Dill style flushing [8] in order to align OoO executions with their sequential counterparts. Our correlate is the serialization proofs for OoO and speculation sketched in appendices 6.3 and A.3. It is of interest to mechanize these proofs and to examine if a generic account of serialization using, e.g., InSpectre can help also in the functional verification of concrete pipelines.

Hardware countermeasures While CPU vendors and researchers propose countermeasures, it is hard to validate their effectiveness without a model. InSpectre can help modeling and reasoning about their security guarantees, as in Section 7.1.1. Similarly, InSpectre

can model the hardware configurations and fences designed by Intel [25] to stall (part of) an instruction stream in case of speculation. Several works [29, 49, 55, 56, 60] propose security-aware hardware that prevent Spectre-like attacks. InSpectre can help formalizing these hardware features and analyzing their security.

9 CONCLUDING REMARKS

This paper presented InSpectre, the first comprehensive model capable of capturing out-of-order execution along with the different forms of speculation that could be implemented in a high-performance pipeline. We used InSpectre to model existing vulnerabilities, to discover three new potential vulnerabilities, and to reason about the security of existing countermeasures proposed in the literature. There are a number of interesting directions left open in this work.

Foundations of microarchitecture security We argue that InSpectre pushes the boundary on foundations of microarchitecture security with respect to the current state-of-the-art substantially. Existing models [11, 12, 14, 22, 40] miss features like dynamic inter-instruction dependency (except [11]), instruction non-atomicity, OoO memory commits, and partial misprediction of rollbacks. These features were essential to discover the vulnerabilities, as well as to reason about countermeasures like retpoline or memory fences for data dependency. For instance, InSpectre would not have captured our Spectre-OoO vulnerability if the memory stores and instruction retire are performed in the *sequential* order. Similarly, *static* computation of active stores would not have exposed Store-To-Load variants of Spectre. Moreover, forcing the rollback of all subsequent microinstructions as soon as a value is mispredicted prevents modeling advanced recovery methods used by modern processors, including concurrent and partial recovery in case of multiple mispredictions.

A novel feature of our approach is to decompose instructions into smaller microinstruction-like units. We argue that the modeling of pipelines using ISA level instructions as atomic units is in the long run the wrong approach, not reflecting well the behaviour at the hardware level, and unable to provide the foundation for real pipeline information flow. Non-atomicity is needed to handle, for instance, intra-instruction dependencies and interactions between I/D-caches. Therefore, decomposing instructions into smaller microinstructions, as we do, appears essential.

InSpectre lacks explicit support of Meltdown-like vulnerabilities, multicore and hyperthreading, fences, TLBs, cache eviction policies, and mechanisms used to update branch predictor tables. Our model can already capture many of these features. In the paper we give Intel’s *lfence* as an example. We focus here on core aspects of out-of-order and speculation, but there is nothing inherent in the framework that prevents modeling the above additional features. Also, by providing a general model we cannot currently argue if a concrete architecture is secure. For that we need to specialize the model to a given architecture, by adding detail and eliminating nondeterminism.

Tooling Tooling is needed to explore more systematically the utility of the model for exploit search and countermeasure proof, and the framework needs to be instantiated to different concrete pipeline architectures and be experimentally validated.

One can envisage MIL-based analysis tools like Spectector [22], Pitchfork [11], and oo7 [54]. However, the large nondeterminism introduced by out-of-order and speculation will make such an approach inefficient. We are currently taking a different route by modeling concrete microarchitectures within a theorem prover. This allows verifying conditional noninterference if the microarchitecture is inherently secure. A failing security proof gives a basis for proving countermeasure soundness as in Section 7.4, and the identification of sufficient conditions that can be verified in the (more tractable) sequential model.

REFERENCES

- [1] Mark D. Aagaard, Byron Cook, Nancy A. Day, and Robert B. Jones. 2001. A Framework for Microprocessor Correctness Statements. In *Correct Hardware Design and Verification Methods*. 433–448.
- [2] Onur Aciğmez and Çetin Kaya Koç. 2006. Trace-driven cache attacks on AES (short paper). In *International Conference on Information and Communications Security*. Springer, 112–121.
- [3] José Bacelar Almeida, Manuel Barbosa, Gilles Barthe, François Dupressoir, and Michael Emmi. 2016. Verifying constant-time implementations. In *25th {USENIX} Security Symposium ({USENIX} Security 16)*. 53–70.
- [4] ARM. 2018. SSBS, Speculative Store Bypass Safe. (2018). <https://developer.arm.com/docs/ddi0595/d/aarch64-system-registers/ssbs> Accessed: 2020-01-16.
- [5] ARM. 2019. Cache Speculation Side channels v2.4. (2019).
- [6] Musard Balliu, Mads Dam, and Roberto Guanciale. 2014. Automating Information Flow Analysis of Low Level Code. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3-7, 2014*. 1080–1091.
- [7] Daniel J. Bernstein. 2005. Cache-timing attacks on AES. <http://cr.yp.to/antiforgery/cachetiming-20050414.pdf>. (2005).
- [8] Jerry R. Burch and David L. Dill. 1994. Automatic verification of pipelined microprocessor control. In *Computer Aided Verification*, David L. Dill (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 68–80.
- [9] Claudio Canella, Jo Van Bulck, Michael Schwarz, Moritz Lipp, Benjamin von Berg, Philipp Ortner, Frank Piessens, Dmitry Evtvushkin, and Daniel Gruss. 2019. A Systematic Evaluation of Transient Execution Attacks and Defenses. In *28th USENIX Security Symposium, USENIX Security 2019, Santa Clara, CA, USA, August 14-16, 2019*. 249–266.
- [10] Claudio Canella, Daniel Genkin, Lukas Giner, Daniel Gruss, Moritz Lipp, Marina Minkin, Daniel Moghimi, Frank Piessens, Michael Schwarz, Berk Sunar, Jo Van Bulck, and Yuval Yarom. 2019. Fallout: Leaking Data on Meltdown-resistant CPUs. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM.
- [11] Sunjay Cauligi, Craig Disselkoen, Klaus v Gleissenthall, Deian Stefan, Tamara Rezk, and Gilles Barthe. 2019. Towards Constant-Time Foundations for the New Spectre Era. *arXiv preprint arXiv:1910.01755* (October 2019).
- [12] Kevin Cheang, Cameron Rasmussen, Sanjit Seshia, and Pramod Subramanyan. 2019. A Formal Approach to Secure Speculation. In *CSF 2019*.
- [13] Guoxing Chen, Sanchuan Chen, Yuan Xiao, Yinqian Zhang, Zhiqiang Lin, and Ten H Lai. 2018. Sgxspectre attacks: Stealing intel secrets from sgx enclaves via speculative execution. *arXiv preprint arXiv:1802.09085* (2018).
- [14] Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, and James Riely. 2019. The Code That Never Ran: Modeling Attacks on Speculative Evaluation. In *S&P 2019*.
- [15] Craig Disselkoen, David Kohlbrenner, Leo Porter, and Dean Tullsen. 2017. Prime+ Abort: A Timer-Free High-Precision L3 Cache Attack using Intel {TSX}. In *26th USENIX Security Symposium*. 51–67.
- [16] Dmitry Evtvushkin, Ryan Riley, Nael CSE Abu-Ghazaleh, and Dmitry Ponomarev. et al. 2018. BranchScope: A new side-channel attack on directional branch predictor. In *Proceedings of the Twenty-Third International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM.
- [17] Kurt A. Feiste, John S. Muhich, Larry E. Thatcher, and Steven W. White. 2000. Forwarding store instruction result to load instruction with reduced stall or flushing by effective/real data address bytes matching. (February 2000). <https://patents.google.com/patent/US6021485> US6021485A.
- [18] Anders Fogh and Christopher Ertl. 2018. Wrangling with the Ghost An Inside Story of Mitigating Speculative Execution Side Channel Vulnerabilities. (2018).
- [19] Qian Ge, Yuval Yarom, David Cock, and Gernot Heiser. 2018. A survey of microarchitectural timing attacks and countermeasures on contemporary hardware. *Journal of Cryptographic Engineering* 8, 1 (2018), 1–27.
- [20] Daniel Gruss, Clémentine Maurice, Klaus Wagner, and Stefan Mangard. 2016. Flush+ Flush: a fast and stealthy cache attack. In *International Conference on*

- Detection of Intrusions and Malware, and Vulnerability Assessment*. Springer, 279–299.
- [21] Daniel Gruss, Raphael Spreitzer, and Stefan Mangard. 2015. Cache template attacks: Automating attacks on inclusive last-level caches. In *24th USENIX Security Symposium*. 897–912.
- [22] Marco Guarnieri, Boris Köpf, José F. Morales, Jan Reineke, and Andrés Sánchez. 2020. SPECTECTOR: Principled detection of speculative information flows. In *Proceedings of the 41st IEEE Symposium on Security and Privacy*. IEEE.
- [23] Jann Horn. 2018. speculative execution, variant 4: speculative store bypass. (2018).
- [24] Jann Horn et al. 2018. Reading privileged memory with a side-channel. *Project Zero* 39 (2018).
- [25] Intel. 2018. Speculative Execution Side Channel Mitigations, Revision 3.0. (2018).
- [26] Intel. 2018. Speculative Execution Side Channel Update. (2018).
- [27] Intel. 2019. Guidelines for Mitigating Timing Side Channels Against Cryptographic Implementations. (2019). [https://software.intel.com/security-software-guidance/insights/guidelines-mitigating-timing-side-channel-hardware-vulnerabilities/](https://software.intel.com/security-software-guidance/insights/guidelines-mitigating-timing-side-channel-hardware-vulnerabilities)
- [28] Ranjit Jhala and Kenneth L McMillan. 2001. Microarchitecture verification by compositional model checking. In *International Conference on Computer Aided Verification*. Springer, 396–410.
- [29] V. Kiriansky, I. Lebedev, S. Amarasinghe, S. Devadas, and J. Emer. 2018. DAWG: A Defense Against Cache Timing Attacks in Speculative Execution Processors. In *2018 51st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*.
- [30] Vladimir Kiriansky and Carl Waldspurger. 2018. Speculative buffer overflows: Attacks and defenses. *arXiv preprint arXiv:1807.03757* (2018).
- [31] Paul Kocher, Jann Horn, Anders Fogh, Daniel Genkin, Daniel Gruss, Werner Haas, Mike Hamburg, Moritz Lipp, Stefan Mangard, Thomas Prescher, Michael Schwarz, and Yuval Yarom. 2019. Spectre Attacks: Exploiting Speculative Execution. In *2019 IEEE Symposium on Security and Privacy, SP 2019, San Francisco, CA, USA, May 19–23, 2019*. 1–19.
- [32] Paul C Kocher. 1996. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In *Annual International Cryptology Conference*. Springer, 104–113.
- [33] Esmail Mohammadian Koruyeh, Khaled N. Khasawneh, Chengyu Song, and Nael Abu-Ghazaleh. 2018. Spectre Returns! Speculation Attacks using the Return Stack Buffer. In *12th USENIX Workshop on Offensive Technologies (WOOT 18)*. USENIX Association.
- [34] Esmail Mohammadian Koruyeh, Khaled N Khasawneh, Chengyu Song, and Nael Abu-Ghazaleh. 2018. Spectre returns! speculation attacks using the return stack buffer. In *12th USENIX Workshop on Offensive Technologies (WOOT)*.
- [35] Moritz Lipp, Michael Schwarz, Daniel Gruss, Thomas Prescher, Werner Haas, Anders Fogh, Jann Horn, Stefan Mangard, Paul Kocher, Daniel Genkin, et al. 2018. Meltdown: Reading kernel memory from user space. In *27th USENIX Security Symposium*. 973–990.
- [36] Fangfei Liu, Yuval Yarom, Qian Ge, Gernot Heiser, and Ruby B Lee. 2015. Last-level cache side-channel attacks are practical. In *2015 IEEE Symposium on Security and Privacy*. IEEE, 605–622.
- [37] Giorgi Maisuradze and Christian Rossow. 2018. ret2spec: Speculative execution using return stack buffers. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2109–2122.
- [38] P. Manolios and S. K. Srinivasan. 2005. Refinement maps for efficient verification of processor models. In *Design, Automation and Test in Europe*. 1304–1309 Vol. 2.
- [39] Clémentine Maurice, Manuel Weber, Michael Schwarz, Lukas Giner, Daniel Gruss, Carlo Alberto Boano, Stefan Mangard, and Kay Römer. 2017. Hello from the Other Side: SSH over Robust Cache Covert Channels in the Cloud.. In *NDSS*, Vol. 17. 8–11.
- [40] Ross McIlroy, Jaroslav Sevcik, Tobias Tebbi, Ben L Titzer, and Toon Verwaest. 2019. Spectre is here to stay: An analysis of side-channels and speculative execution. *arXiv preprint arXiv:1902.05178* (2019).
- [41] Matt Miller. 2018. Mitigating speculative execution side channel hardware vulnerabilities. (2018). <https://msrc-blog.microsoft.com/2018/03/15/mitigating-speculative-execution-side-channel-hardware-vulnerabilities/>
- [42] David Molnar, Matt Piotrowski, David Schultz, and David A. Wagner. 2005. The Program Counter Security Model: Automatic Detection and Removal of Control-Flow Side Channel Attacks. In *Information Security and Cryptology - ICISC 2005, 8th International Conference, Seoul, Korea, December 1-2, 2005, Revised Selected Papers*. 156–168.
- [43] Michael Neve and Jean-Pierre Seifert. 2006. Advances on access-driven cache attacks on AES. In *International Workshop on Selected Areas in Cryptography*. Springer, 147–162.
- [44] Dag Arne Osvik, Adi Shamir, and Eran Tromer. 2006. Cache attacks and countermeasures: the case of AES. In *Cryptographers’ track at the RSA conference*. Springer, 1–20.
- [45] Dan Page. 2002. Theoretical use of cache memory as a cryptanalytic side-channel. *IACR Cryptology ePrint Archive* 2002, 169 (2002).
- [46] Andrew Pardoe. 2018. Spectre mitigations in MSVC. (2018).
- [47] Jun Sawada and Warren A Hunt. 2002. Verification of FM9801: An out-of-order microprocessor model with speculative execution, exceptions, and program-modifying capability. *Formal Methods in System Design* 20, 2 (2002), 187–222.
- [48] Michael Schwarz, Moritz Lipp, Daniel Moghimi, Jo Van Bulck, Julian Stecklina, Thomas Prescher, and Daniel Gruss. 2019. ZombieLoad: Cross-Privilege-Boundary Data Sampling. In *CCS*.
- [49] Mohammadkazem Taram, Ashish Venkat, and Dean Tullsen. 2019. Context-sensitive fencing: Securing speculative execution via microcode customization. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM, 395–410.
- [50] Caroline Trippel, Daniel Lustig, and Margaret Martonosi. 2018. CheckMate: Automated Synthesis of Hardware Exploits and Security Litmus Tests. In *51st Annual IEEE/ACM International Symposium on Microarchitecture, MICRO 2018, Fukuoka, Japan, October 20-24, 2018*. 947–960.
- [51] Paul Turner. 2018. Retpoline: a software construct for preventing branch-target-injection. (2018). <https://support.google.com/faqs/answer/7625886>
- [52] Siddhant Aravamudan, Benjamin Berg, Misha Chertov, Daniel Cohen, Pietro Frigo, Giorgi Maisuradze, Kaveh Razavi, Herbert Bos, and Cristiano Giuffrida. 2019. RIDL: Rogue In-flight Data Load. In *S&P*.
- [53] Guanhua Wang, Sudipta Chattopadhyay, Arnab Kumar Biswas, Tulika Mitra, and Abhik Roychoudhury. 2019. KLEESPECTRE: Detecting Information Leakage through Speculative Cache Attacks via Symbolic Execution. *arXiv preprint arXiv:1909.00647* (2019).
- [54] G. Wang, S. Chattopadhyay, I. Gotovchits, T. Mitra, and A. Roychoudhury. 2019. oo7: Low-overhead Defense against Spectre attacks via Program Analysis. *IEEE Transactions on Software Engineering* (2019).
- [55] Ofir Weisse, Ian Neal, Kevin Loughlin, Thomas F Wenisch, and Baris Kasikci. 2019. NDA: Preventing Speculative Execution Attacks at Their Source. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*. ACM, 572–586.
- [56] Jonathan Woodruff, Robert NM Watson, David Chisnall, Simon W Moore, Jonathan Anderson, Brooks Davis, Ben Laurie, Peter G Neumann, Robert Norton, and Michael Roe. 2014. The CHERI capability model: Revisiting RISC in an age of risk. In *2014 ACM/IEEE 41st International Symposium on Computer Architecture (ISCA)*. IEEE, 457–468.
- [57] Yuan Xiao, Yingqian Zhang, and Mircea-Radu Teodorescu. 2020. SPEECHMINER: A Framework for Investigating and Measuring Speculative Execution Vulnerabilities. In *NDSS*.
- [58] Yuval Yarom and Katrina Falkner. 2014. FLUSH+ RELOAD: a high resolution, low noise, L3 cache side-channel attack. In *23rd USENIX Security Symposium*. 719–732.
- [59] Yuval Yarom, Daniel Genkin, and Nadia Heninger. 2017. CacheBleed: a timing attack on OpenSSL constant-time RSA. *Journal of Cryptographic Engineering* 7, 2 (2017), 99–112.
- [60] Drew Zagieboylo, Gookwon Edward Suh, and Andrew C Myers. 2019. *Using Information Flow to Design an ISA that Controls Timing Channels*. Technical Report.

A PROOFS


A.1 Correctness of t -equivalence: Lemma 5.3

If $\sigma_1 \sim_t \sigma_2$ and the t 's microinstruction in σ_1 is $t = t \leftarrow c?o$, then $deps(t, \sigma_1) = deps(t, \sigma_2)$, $[c]\sigma_1 = [c]\sigma_2$, and if $[l]\sigma_1 = (v_1, l_1)$ and $[l]\sigma_2 = (v_2, l_2)$ then $v_1 = v_2$.

PROOF. For non-load operations and guards the proof is trivial, since their semantics only depends on the value of the bound names of the operation and guard. These names are statically identified from c and o and their value is the same in σ_1 and σ_2 by definition of \sim_t . For loads, the proof relies on showing that for every state σ , $str-act(\sigma|_T, t) = str-act(\sigma|_T, t)$, where $T = deps(t, \sigma)$. Let $t' \in str-act(\sigma, t)$. By definition, t' and its bound names are in T , therefore their values are equal in σ and $\sigma|_T$ and hence $t' \in str-may(\sigma|_T, t)$. Also, by definition of $deps$, all names referred to by conditions and addresses of subsequent stores of t' are in T . Therefore if there is no subsequent store that overwrites t' (i.e., $t'' \leftarrow c''?st \tau t''_a t''_v$ such that $[c'']\sigma$ and $(\sigma(t''_a) = \sigma(t) \vee \sigma(t''_a) = \sigma(t_a))$) then there is no store overwriting t' in $\sigma|_T$ and hence $t' \in str-act(\sigma|_T, t)$. \square

A.2 OoO Memory Consistency: Theorem 6.3

To prove that \rightarrow and \rightarrow are memory consistent we demonstrate a reordering lemma, which allows to commute transitions if the $(n + 1)$ -th transition modified name t_2 , n -th transition modified name t_1 , and $t_2 < t_1$.



We use the following notation. Let $\sigma_0 :: \dots :: \sigma_n$ a sequence of states, we define $(I_i, s_i, C_i, F_i) = \sigma_i$ for $i \in \{0 \dots n\}$, $(\alpha_i, t_i) = \text{step-param}(\sigma_{i-1}, \sigma_i)$ and $(I_{i-1} \cup \hat{I}_i, s_{i-1} \cup \hat{s}_i, C_{i-1} \cup \hat{C}_i, F_{i-1} \cup \hat{F}_i) = (I_i, s_i, C_i, F_i)$ for $i \in \{1 \dots n\}$.

We first demonstrate that *str-may* and *str-act* of a microinstruction t do not depend on names bigger than t and that they are monotonic.

LEMMA A.1. *Let σ_0 and σ_1 be two states, if $bn(\hat{I}_1) \geq t$ and $\text{dom}(\hat{s}_1) \geq t$ then $\text{str-may}(\sigma_1, t) = \text{str-may}(\sigma_0, t)$ and $\text{str-act}(\sigma_1, t) = \text{str-act}(\sigma_0, t)$*

PROOF. Let t be a load or store accessing address t_a (the other cases are trivial, since *str-act* is undefined), hence $t_a < t$.

(1) The set of instructions that precedes t is the same in σ_0 and σ_1 . In fact, since $bn(\hat{I}_1) \geq t$ then $\{t' \leftarrow c'?o \in I_0 \cup \hat{I}_1 \mid t' < t\} = \{t' \leftarrow c'?o \in I_0 \mid t' < t\}$.

(2) For every store that precedes t , the evaluation of condition and address is the same in σ_0 and σ_1 . In fact, let $t' = t' \leftarrow c'?st \tau t'_a t'_v$ and $t' < t$ then $n(c' \cup \{t'_a\}) < t$. Therefore, $[c']s_0 \cup \hat{s}_1 = [c']s_0$ and $(s_0 \cup \hat{s}_1)(t'_a) = s_0(t'_a)$.

(3) Similarly, since $t_a < t$ then $(s_0 \cup \hat{s}_1)(t_a) = s_0(t_a)$.

Properties (1, 2, 3) guarantee that $\text{str-may}(\sigma_1, t) = \text{str-may}(\sigma_0, t)$. Similarly, since *str-act* depends on the addresses and conditions of stores in $\text{str-may}(\sigma_1, t)$ and these have names smaller than t then $\text{str-act}(\sigma_1, t) = \text{str-act}(\sigma_0, t)$. \square

LEMMA A.2. *if $\sigma_0 \xrightarrow{l} \sigma_1$ and $t \leftarrow c'?o \in \sigma$ then $\text{str-may}(\sigma_1, t) \subseteq \text{str-may}(\sigma_0, t)$ and $\text{str-act}(\sigma_1, t) \subseteq \text{str-act}(\sigma_0, t)$*

PROOF. The proof is done by case analysis on α_1 . For commits the proof is trivial, since $\hat{I}_1 = \emptyset$ and $\hat{s}_1 = \emptyset$. For fetches, $\hat{s}_1 = \emptyset$ and the transition may decode new stores in \hat{I}_1 . However, these new stores have names greater than $\max(I_0)$, hence their names are greater than t . Therefore the additional stores do not affect *str-may* and *str-act*.

For executions, $\hat{I}_1 = \emptyset$, $\hat{s}_1 = \{t_1 \mapsto v\}$, and $s_0(t_1) \uparrow$ for some v . This store update can make defined the evaluation of the condition or expression of a store. In this case, if a store is in $\text{str-may}(\sigma_1, t)$ it must also be in $\text{str-may}(\sigma_0, t)$. Stores that are in $\text{str-may}(\sigma_0, t)$ but are not in $\text{str-may}(\sigma_1, t)$ have undefined conditions or addresses in σ_0 and false condition or non matching address in σ_1 .

To show that *str-act* does not increase we proceed as follows. Let t' be a store in $\text{str-may}(\sigma_0, t) \setminus \text{str-act}(\sigma_0, t)$. There must be a subsequent overwriting store t'' in $\text{str-may}(\sigma_0, t)$ whose condition holds in σ_0 and address is defined in σ_0 . Such store cannot have t_1 in its free names, hence it is also in $\text{str-may}(\sigma_1, t)$. Therefore the store t'' overwrites t' in σ_1 too. \square

Proof of Theorem 6.3 is done by induction on the length of traces and relies on Lemma A.3 to demonstrate that (α_2, t_2) can be applied in σ_n and Lemma A.4 to show that (α_1, t_1) can be applied in the resulting state σ'_{n+1} to obtain σ_{n+2} .

LEMMA A.3. *If $\sigma_0 \xrightarrow{l_1} \sigma_1 \xrightarrow{l_2} \sigma_2$ and $t_2 < t_1$ then exists l'_2 such that $\sigma_0 \xrightarrow{l'_2} (I_0 \cup \hat{I}_2, s_0 \cup \hat{s}_2, C_0 \cup \hat{C}_2, F_0 \cup \hat{F}_2) = \sigma'$, $\text{step-param}(\sigma_0, \sigma') = (\alpha_2, t_2)$, and if $\alpha_2 = \text{CMT}(a, v)$ then $\alpha_1 \neq \text{CMT}(a, v')$.*

PROOF. We first bound the effects of the transitions that modified t_1 .

(1) If $\sigma_1 \xrightarrow{l_2} \sigma_2$ and $\text{step-param}(\sigma_1, \sigma_2) = (\alpha_2, t_2)$, then exists $l_2 = t_2 \leftarrow c'?o \in I_1$. Since $\text{step-param}(\sigma_0, \sigma_1) = (\alpha_1, t_1)$ then $bn(\hat{I}_1) > t_1 > t_2$. Therefore $t_2 \leftarrow c'?o \in I_0$.

(2) Similarly, $\text{dom}(\hat{s}_1) \subseteq \{t_1\}$, $\hat{C}_1 \subseteq \{t_1\}$, and $\hat{F}_1 \subseteq \{t_1\}$.

The proof continues by case analysis over the transition rule α_2 . (CASE EXE) The hypothesis of the rule ensure that $s_1(t_2) \uparrow$, $[c]s_1$, and $[l_2]\sigma_1 = (v, l_2)$. The conclusion of the rule ensures that $\hat{s}_2 = \{t_2 \mapsto v\}$, $\hat{C}_2 = \emptyset$, $\hat{F}_2 = \emptyset$, and $\hat{I}_2 = \emptyset$.

The proof that t_2 can be executed in σ_0 relies on the fact that all free names of the instruction t_2 must be smaller than t_2 . Property (2), $fn(t_2) < t_2$, and $t_2 < t_1$ ensure that $s(t_2) \uparrow$ and $[c]s$.

The same properties guarantee that $[l_2]\sigma_0 = (v, l_2)$. For internal operations and stores the proof is trivial, since $fn(t_2) < t_2$, and $t_2 < t_1$. The proof for loads uses Lemma A.1 to guarantee that $\text{str-act}(\sigma_0, t_2) = \text{str-act}(\sigma_1, t_2)$.

Hence we can apply rule (exec) to show that exists l'_2 such that $\sigma_0 \xrightarrow{l'_2} (I_0, s_0 \cup \{t_2 \mapsto v\}, C_0, F_0) = \sigma'$. Notice that $l'_2 \neq l_2$. In fact, if t_2 is a load, t_1 is the corresponding active store, and $\alpha_1 = \text{CMT}(a, v)$ then the execution of t_2 needs to access the memory subsystem in σ_1 while it can simply forward the value of t_1 in σ_0 : i.e., an observable load becomes silent.

(CASE CMT(a, v)) In this case $o = st \mathcal{M} t_a t_v$. The hypothesis of the rule ensure that $s_1(t_2) = v$, $t_2 \notin C_1$, $bn(\text{str-may}(\sigma_1, t_2)) \subseteq C_1$, and $s_1(t_a) = a$. The conclusion of the rule ensures that $\hat{s}_2 = \emptyset$, $\hat{C}_2 = \{t_2\}$, $\hat{F}_2 = \emptyset$, $\hat{I}_2 = \emptyset$, and $l_2 = ds a$.

Property (2) and $t_a < t_2 < t_1$ ensure that $s_0(t_2) = v$, $t_2 \notin C_0$, and $s_0(t_a) = a$. Similarly to the case EXE-load, Lemma A.1 guarantees that $bn(\text{str-may}(\sigma_1, t_2)) = bn(\text{str-may}(\sigma_0, t_2))$. Since the *str-may* are smaller than t_2 then $bn(\text{str-may}(\sigma_0, t_2)) \subseteq C_0$.

Hence we can apply rule (commit) to show that $\sigma_0 \xrightarrow{l_2} (I_0, s_0, C_0 \cup \{t_2\}, F_0) = \sigma'$. Finally, to show that $\alpha_1 \neq \text{CMT}(a, v')$ we proceed by contradiction. If $\alpha_1 = \text{CMT}(a, v')$, then $bn(\text{str-may}(\sigma_0, t_1)) \subseteq C_0$. However, $t_2 \in bn(\text{str-may}(\sigma_0, t_1))$, because they write the same address a and $t_2 < t_1$. This contradict that $t_2 \notin C_0$.

(CASE FRC) In this case $o = st \mathcal{PC} t_v$. The hypothesis of the rule ensure that $s_1(t_2) = v$, $t_2 \notin F_1$, and $bn(\text{str-may}(\sigma_1, t_2)) \subseteq F_1$. The conclusion of the rule ensures that $\hat{s}_2 = \emptyset$, $\hat{C}_2 = \emptyset$, $\hat{F}_2 = \{t_2\}$, $\hat{I}_2 = \text{translate}(v, \max(I_1))$, and $l_2 = il a$.

Property (2) and $fn(c) \cup \{t_a\} < t_2 < t_1$ ensure that $s_0(t_2) = v$ and $t_2 \notin F_0$. Similarly to the case commit, Lemma A.1 guarantee that $bn(\text{str-may}(\sigma_1, t_2)) = bn(\text{str-may}(\sigma_0, t_2))$. Since the *str-may* are smaller than t_2 then $bn(\text{str-may}(\sigma_0, t_2)) \subseteq F_0$.

To complete the proof we must show that $\hat{I}_1 = \emptyset$. We proceed by contradiction: if $\hat{I}_1 \neq \emptyset$ then $\alpha_1 = \text{FRC}(\hat{I}_1)$, hence this transition

fetched t_1 and $bn(str\text{-}may(\sigma_0, t_1)) \subseteq F_0$. However, $t_2 < t_1$ and both update the program counter, therefore $t_2 \in bn(str\text{-}may(\sigma_0, t_1))$. This contradicts $t_2 \notin F_0$.

Finally, we can apply rule Frc to show that $\sigma_0 \xrightarrow{l_2} (I_0 \cup \text{translate}(v, \max(I_0), s_0, C_0, F_0 \cup \{t_2\})) = (I_0 \cup \text{translate}(v, \max(I_1), s_0, C_0, F_0 \cup \{t_2\})) = \sigma'$. \square

LEMMA A.4. *If $\sigma_0 \xrightarrow{l_1} \sigma_1 \xrightarrow{l_2} \sigma_2$ and $t_2 < t_1$ then there exists σ', l'_1 , and l'_2 such that $\sigma_0 \xrightarrow{l'_2} \sigma' \xrightarrow{l'_1} \sigma_2$, $step\text{-}param(\sigma_0, \sigma') = (\alpha_2, t_2)$, $step\text{-}param(\sigma', \sigma_2) = (\alpha_1, t_1)$, and if $\alpha_2 = \text{CMT}(a, v)$ then $l_1 \neq \text{CMT}(a, v')$.*

PROOF. The existence of σ' and l'_2 is given by Lemma A.3. For transition $\sigma' \xrightarrow{l'_1} \sigma_2$ we first bound the effects of $\sigma_0 \xrightarrow{l_1} \sigma_1$.

(1) If $\sigma_0 \xrightarrow{l_1} \sigma_1$ and $step\text{-}param(\sigma_0, \sigma_1) = (\alpha_1, t_1)$, then there exists $t_1 = t_1 \leftarrow c?o \in I_0$. Therefore $t_1 \leftarrow c?o \in I_0 \cup \hat{I}_2 = I'$.

(2) $\text{dom}(\hat{s}_2) \subseteq \{t_2\}$, $\hat{C}_2 \subseteq \{t_2\}$, and $\hat{F}_2 \subseteq \{t_2\}$.

We continue the proof by case analysis on α_2 .

(Case EXE) The hypothesis of the rule ensure that $s_0(t_1) \uparrow$, $[c]s_0$, and $[l_1]\sigma_0 = (v, l_1)$. The conclusion of the rule ensures that $\hat{s}_1 = \{t_1 \mapsto v\}$, $\hat{C}_1 = \emptyset$, $\hat{F}_1 = \emptyset$, and $\hat{I}_1 = \emptyset$.

(3) Property (2) and $t_2 < t_1$ ensure that $s'(t_1) \uparrow$. Moreover, if $\hat{s}_2 = \emptyset$ then $[c]s' = [c]s_0$. Otherwise, $\hat{s}_2 = \{t_2 \mapsto v_2\}$ for some v_2 . In this case, $\alpha_2 = \text{EXE}$ and $s_0(t_2) \uparrow$. Since $[c](s_0) \downarrow$ then $t_2 \notin n(c)$, hence $[c]s' = [c]s_0$.

The same argument is used to guarantee that $[l_1]\sigma' = (v, l'_1)$. For internal operations and stores the proof follows the same approach of (3). The proof for loads uses Lemma A.2. Notice that for internal operations and stores $l'_1 = l_1 = \cdot$. For loads, either $l'_1 = l_1 = dl\ a$ or $l'_1 = dl\ a$ and $l'_1 = \cdot$. The latter happens when $step\text{-}param(\sigma_0, \sigma') = (\text{CMT}(a, v), t_2)$. In this case, we have reordered a memory commit before a load thus making it not possible to forward the value of the store to the load and requiring a new memory interaction. This shows that the observations of the OoO model are a subset of the observations of in-order model, since the OoO model can execute loads before the corresponding stores are committed.

Finally, we can apply rule EXE to show that $\sigma' \xrightarrow{l'_1} (I', s' \cup \{t_1 \mapsto v\}, C', F') = \sigma_2$.

(Case CMT)(a, v) In this case $o = st\ \mathcal{M}\ t_a\ t_v$. The hypothesis of the rule ensure that $s_0(t_1) = v$, $t_1 \notin C_0$, $bn(str\text{-}may(\sigma_0, t_1)) \subseteq C_0$, and $s_0(t_a) = a$. The conclusion of the rule ensures that $\hat{s}_1 = \emptyset$, $\hat{C}_1 = \{t_1\}$, $\hat{F}_1 = \emptyset$, $I_{t_1} = \emptyset$, and $l_1 = ds\ a$.

Property (2) and $t_2 < t_1$ ensure that $s'(t_1) = v$ and $s'(t_a) = a$. To show that $bn(str\text{-}may(\sigma', t_1)) \subseteq C'$ we use Lemma A.2. Finally, $t_1 \notin C'$, since $t_1 > t_2$. To prove that $\alpha_2 \neq \text{CMT}(a, v')$ we proceed by contradiction. If $\alpha_2 = \text{CMT}(a, v')$ then $t_2 \notin C_0$. However, $t_2 \in bn(str\text{-}may(\sigma_0, t_1))$, because they write the same address a and $t_2 < t_1$. This contradict that $bn(str\text{-}may(\sigma_0, t_1)) \subseteq C_0$.

Therefore we can apply rule CMT to show that $\sigma' \xrightarrow{l_1} (I', s', C' \cup \{t_1\}, F') = \sigma_2$.

(Case Frc) In this case $o = st\ \mathcal{P}C\ t_v$. The hypothesis of the rule ensure that $s_0(t_1) = v$, $t_1 \notin F_0$, and $bn(str\text{-}may(\sigma_0, t_1)) \subseteq F_0$. The conclusion of the rule ensures that $\hat{s}_1 = \emptyset$, $\hat{C}_1 = \emptyset$, $\hat{F}_1 = \{t_1\}$, $I_{t_1} = \text{translate}(v, \max(I_1))$, and $l_1 = il\ a$.

Property (2) and $t_2 < t_1$ ensure that $s'(t_1) = v$. To show that $bn(str\text{-}may(\sigma', t_1)) \subseteq F'$ we use Lemma A.2.

Finally, $t_1 \notin F'$, since $t_1 > t_2$.

To complete the proof we must show that $\hat{I}_2 = \emptyset$. We proceed by contradiction: if $\hat{I}_2 \neq \emptyset$ then $\alpha_2 = \text{Frc}$, hence this transition fetched t_2 and $t_2 \notin F_0$. However, $t_2 < t_1$ and both update the program counter, therefore $t_2 \in bn(str\text{-}may(\sigma_0, t_1))$. This contradicts the hypothesis that $bn(str\text{-}may(\sigma_0, t_1)) \subseteq F_0$.

Finally, we can apply rule Frc to show that $\sigma' \xrightarrow{l_1} (I' \cup \text{translate}(v, \max(I'), s', C', F' \cup \{t_1\})) = (I' \cup \text{translate}(v, \max(I_0), s', C', F' \cup \{t_1\})) = \sigma'$. \square

A.3 Memory Consistency of Speculative Semantics: Theorem 6.4

We reduce memory consistency for the speculation model to the OoO case using Theorem 6.3. Since the OoO semantics already takes care of reordering, to prove Theorem 6.4 a bisimulation argument suffices. Intuitively, referring to Figure 2, the states “decoded”, “predicted”, “speculated” and “speculatively fetched” in the speculative semantics all correspond in some sense to the state “decoded” in the OoO semantics, in that any progress can still be undone to return to the “decoded” state. In a similar vein, the state “retired” corresponds to “executed” in the OoO semantics, “fetched” to “fetched” and “committed” to “committed”. The only exception is states that are speculatively fetched. In this case there is an option to directly retire the fetched state, without passing through “retired” first. The proof reflects this intuition.

The main challenges in defining the bisimulation are i) to pin down the non-speculated instructions in the speculative semantics and relate them correctly to instructions in the OoO semantics, and ii) account for speculatively fetched instructions. The latter issue arises when retiring an instruction in the speculative semantics that has earlier been speculatively fetched. In that case, the corresponding decoded microinstructions are already in flight, although the bisimilar OoO state will have no trace of this. Then the OoO microinstruction will have to be first executed and then fetched. The following definitions make this intuition precise.

First say that a name t' is *produced* by the PC store microinstruction $t \leftarrow c?st\ \mathcal{P}C\ t_v \in I$, if $t' < t$, i.e. $t \in \text{dom}(\delta(t'))$. We would like to conclude that t is uniquely determined, as we need this to properly relate the speculative and OoO states. However, this does not hold in general. For a counterexample consider the PC store microinstruction $t \leftarrow c?st\ \mathcal{P}C\ t_v$. Suppose that the fetch from t causes a new instruction I' to be allocated with another PC store instruction $t' \leftarrow c'?st\ \mathcal{P}C\ t'_v$ followed by a PC load $t'' \leftarrow c''?ld\ \mathcal{P}C$, i.e. such that $t' < t''$. At this point, $\delta(t') = [t \mapsto v]$ and $\delta(t'') = [t \mapsto v]$. After executing the fetched PC store t' and then the PC load, $s(t') = v'$ and $s(t'') = v'$. At this point, $\delta(t'')$ will map t to v and t' to v' . But then t'' is produced by both t and t' . The same property holds if t' is used as an argument to operations other than a PC load. This causes us to impose the following wellformedness condition on instruction translations:

Definition A.5 (Wellformed instruction translation). The translation function $translate$ is wellformed if $translate(v, t) = I$ implies:

- (1) $t \leftarrow c?st \mathcal{PC} \ t_v, t' \leftarrow c'?ld \mathcal{PC} \in I$ implies $t' < t$.
- (2) $t \leftarrow c?st \mathcal{PC} \ t_v, t \in I$ implies $t \notin fn(i)$.
- (3) For all s there is a unique $t \leftarrow c?st \mathcal{PC} \ t_v$ such that $[c]s$.

Condition A.5.1 and 2 can be imposed without loss of generality since any occurrence of t bound to the microinstruction $t \leftarrow c?st \mathcal{PC} \ v$ can be replaced by v itself, and condition A.5.3 is natural to ensure that any linear control flow gives rise to a correspondingly linear flow of instructions. We obtain:

PROPOSITION A.6. *If t' is produced by t_1 and t' is produced by t_2 then $t_1 = t_2$.* \square

Consider now microinstructions (images of *translate*) I_1 and I_2 such that $bn(I_1) \cap bn(I_2) = \emptyset$. Say that I_1 produces $I_2, I_1 < I_2$, if there is $t \leftarrow c?st \mathcal{PC} \ t_v \in I_1$ such that for each $t' \in I_2$, t' is produced by t . Clearly, if I' is added to the set of microinstructions due to a fetch from I then $I < I'$. Say then that I (and by extension states containing I) is *wellformed* by the partitioning $I_{retired}, I_1, \dots, I_n$, if $I = \bigcup \{I_{retired}, I_1, \dots, I_n\}$, $I_{retired}$ is retired, and for each $I_i, 1 \leq i \leq n$ there is $I \in \{I_{retired}, I_1, \dots, I_n\}$ such that $I < I_i$. Moreover we require that $<^*$ on the partitions $\{I_{retired}, I_1, \dots, I_n\}$ is well-founded and that the partitions are maximal. Note that if I is wellformed by $I_{retired}, I_1, \dots, I_n$ then the partitioning is unique. We note also that all reachable states in the speculative semantics are wellformed and each partition corresponds to the translation of one single ISA instruction. We say that an ISA instruction I_i in the partitioning I_1, \dots, I_n is *unconditionally fetched*, if $I_{retired} < I_i$ and let I_{uf} be the union of $I_{retired}$ and the instructions that have been unconditionally fetched.


We can now proceed to define the bisimulation R . We restrict attention to reachable states in both the OoO and speculative semantics in order to keep the definition of R manageable and be able to implicitly make use of simple invariant properties such as $\text{dom}(\delta) \cap \text{dom}(\delta(t)) \Rightarrow t \notin C$ (no instruction with a speculated dependency is committed). Let $(I_1, s_1, C_1, F_1) R (I_2, s_2, C_2, F_2, \delta_2, P_2)$ if

- (1) I_2 is wellformed by the partitioning $I_{2,retired}, I_{2,1}, \dots, I_{2,n}$.
- (2) There is a bijection $\tilde{\cdot}$ from $I_{2,uf}$ to I_1 .
- (3) $C_2 = \tilde{C}_1$,
- (4) $s_2 \setminus \text{dom}(\delta_2) = \tilde{s}_1$,
- (5) $F_2 \setminus \text{dom}(\delta_2) = \tilde{F}_1$,

In 3.-5. the bijection $\tilde{\cdot}$ is pointwise extended to sets and expressions.

Note that, from 2. and 4. we get that a microinstruction \tilde{t} in I_1 has been executed iff $t \notin \text{dom}(\delta_2)$.

We prove that R is a weak bisimulation in two steps. We first show that all speculative transitions up until retire or non-speculative fetch are reversible. To prove this it is sufficient to show that each of the conditions 1.-4. is invariant under PRD , EXE , PEXE , RBK , and FTC , the latter under the condition that the fetched instruction is in δ_2 . These transitions are identified by $T1$ in the following picture:



LEMMA A.7. *If $\sigma_1 R (\sigma_2, \delta_2, P_2)$ and $(\sigma_2, \delta_2, P_2) \twoheadrightarrow (\sigma'_2, \delta'_2, P'_2)$ is an instance of PRD , EXE , PEXE , RBK , or speculative FTC then $\sigma_1 R (\sigma'_2, \delta'_2, P'_2)$.*

PROOF. Let $\sigma_2 = (I_2, s_2, C_2, F_2)$ and $(\sigma'_2, \delta'_2, P'_2) = (I'_2, s'_2, C'_2, F'_2, \delta'_2, P'_2)$. **(Case PRD)** We get $\sigma'_2 = (I_2, s_2[t \mapsto v], C_2, F_2, \delta_2, P_2 \cup \{t\})$ and note that conditions 1.-5. are trivially satisfied by the assumptions. **(Case EXE)** We get $\sigma_2 \twoheadrightarrow \sigma'_2$ with $\text{step-param}(\sigma_2, \sigma'_2) = (\text{EXE}, t)$, $t \notin P_2$, $\delta'_2 = \delta_2 \cup \{t \mapsto s_2|_{\text{deps}(t, \sigma_2)}\}$ and $P'_2 = P_2$. Cond. 1 and 2 are straightforward since $I_{2,uf}$ and I_1 are not affected by the transition. For cond. 3 and 5 we get $C'_2 = C_2$ and

$$F'_2 \cap \overline{\text{dom}(\delta'_2)} = F_2 \cap \overline{\text{dom}(\delta_2)},$$

since $\text{dom}(\delta'_2) \setminus \text{dom}(\delta_2) = \{t\}$ and $t \notin \text{dom}(F_2)$. For cond. 4,

$$\begin{aligned}
 s'_2 \setminus \text{dom}(\delta'_2) &= s_2[t \mapsto v] \setminus \text{dom}(\delta_2 \cup \{t \mapsto s_2|_{\text{deps}(t, \sigma_2)}\}) \\
 &= s_2[t \mapsto v] \setminus \text{dom}(\delta_2) \cup \{t\} \\
 &= s_2 \setminus \text{dom}(\delta_2).
 \end{aligned}$$

(Case PEXE) In this case $\sigma_2|_{\{t\}} \twoheadrightarrow \sigma'_2$ with $\text{step-param}(\sigma_2, \sigma'_2) = (\text{EXE}, t)$, $t \in P_2$, $\delta'_2 = \delta_2 \cup \{t \mapsto s_2|_{\text{deps}(t, \sigma_2)}\}$ and $P'_2 = P_2 \setminus \{t\}$. Cond. 1 and 2 again are immediate. For cond. 3, $C'_2 = C_2$ and for cond. 5,

$$F'_2 \cap \overline{\text{dom}(\delta'_2)} = F_2 \cap \overline{\text{dom}(\delta_2)}$$

since, again, $\text{dom}(\delta'_2) \setminus \text{dom}(\delta_2) = \{t\}$ and $t \notin \text{dom}(F_2)$. Finally for cond. 2,

$$\begin{aligned}
 s'_2 \setminus \text{dom}(\delta'_2) &= s_2[t \mapsto v] \setminus \text{dom}(\delta_2) \cup \{t\} \\
 &= s_2[t \mapsto v] \setminus \text{dom}(\delta_2) \cup \{t\} \\
 &= s_2 \setminus \text{dom}(\delta_2)
 \end{aligned}$$

(Case RBK) We get that $(I_2, s_2, C_2, F_2) \not\sim_t (I_2, \delta_2(t), C_2, F_2)$ and $t \notin P_2$. We get $I'_2 = I_2 \setminus \Delta^+$, $s'_2 = s_2|_{\overline{\Delta^+}}$, $C'_2 = C_2$, $F'_2 = F_2 \setminus \Delta^*$, $\delta'_2 = \delta_2 \setminus \Delta^*$, and $P'_2 = P_2 \setminus \Delta^*$. For cond. 1 and 2 first note that $t \in \text{dom}(\delta_2)$. If t is not a PC store the effect of RBK is to remove t from s_2, F_2, δ_2, P_2 . This does not affect the bijection $\tilde{\cdot}$, so 1 and 2 remain valid also for $(\sigma'_2, \delta'_2, P'_2)$. If t is a PC store then we need to observe the following: Since t is speculated, t is a member of some "macro"-instruction (= partition) $I_{2,i}$. The set Δ^+ contains all instructions/partitions $I_{2,j}$ such that $I_{2,i} <^+ I_{2,j}$. In particular, no such $I_{2,j}$ is in $I_{2,uf}$, since otherwise $I_{2,j}$ would have been added by a retired PC store microinstruction. It follows that the bijection $\tilde{\cdot}$ is not affected by the removal of Δ^+ , and 1 and 2 are reestablished for the new speculative state.

For cond. 3, $C'_2 = C_2$. For cond. 4, we calculate:

$$\begin{aligned}
 s'_2|_{\overline{\text{dom}(\delta'_2)}} &= (s_2|_{\overline{\Delta^+}})|_{\overline{\text{dom}(\delta'_2)}} \\
 &= s_2|_{\overline{\text{dom}(\delta'_2) \cup P'_2 \cup \Delta^*}} \\
 &= s_2|_{\overline{\text{dom}(\delta_2 \setminus \Delta^*) \cup \Delta^*}} \\
 &= s_2|_{\overline{\text{dom}(\delta_2) \cup \Delta^*}} \\
 &= s_2|_{\overline{\text{dom}(\delta_2)}}.
 \end{aligned}$$

Note that the final step uses prop. A.8.

PROPOSITION A.8. *The relation \prec^* is a partial order.*

PROOF. By induction in the length of derivation. \square

Finally for cond. 5:

$$\begin{aligned} F'_2 \setminus \text{dom}(\delta'_2) &= (F_2 \setminus \Delta^*) \setminus (\text{dom}(\delta_2) \setminus \Delta^*) \\ &= F_2 \setminus (\text{dom}(\delta_2)). \end{aligned}$$

(Case speculative Frc) We get that $\sigma_2 \twoheadrightarrow \sigma'_2$, $\text{step-param}(\sigma_2, \sigma'_2) = (\text{Frc}(I'_2), t)$, $P'_2 = P_2$, $\delta'_2 = \delta_2 \cup \{t' \mapsto s|_{\{t\}} \mid t' \in I'_2\}$, and, since t is speculated, $t \in \text{dom}(\delta_2)$. Also, we find $t \leftarrow c?st \mathcal{PC} t_v \in I_2$, $s_2(t) = v$, $t \notin F_2$, $\text{bn}(\text{str-may}(\sigma_2, t)) \subseteq F_2$, $I'_2 = I_2 \cup I''_2$, $s'_2 = s$, $C'_2 = C_2$, and $F'_2 = F \cup \{t\}$. For cond. 1 and 2 observe that no instruction I' added by the fetch can belong to $I'_{2, \text{fu}}$, since all such instructions are produced by a retired PC store instruction.

For cond. 3, $C'_2 = C_2$ is immediate.

For cond. 4 we calculate:

$$\begin{aligned} s'_2|_{\overline{\text{dom}(\delta'_2)}} &= s_2|_{\overline{\text{dom}(\delta_2) \cup \{t' \mid t' \in I'_2\}}} \\ &= s_2|_{\overline{\text{dom}(\delta_2)}} \end{aligned}$$

Finally for cond. 5:

$$\begin{aligned} F'_2 \setminus \text{dom}(\delta'_2) &= (F_2 \cup \{t\}) \setminus (\text{dom}(\delta_2) \cup \{t' \mid t' \in I'_2\}) \\ &= F_2 \setminus \text{dom}(\delta_2) \end{aligned}$$

\square

The following Lemma handle cases for RET when $t \notin F_2$, CMT, Frc when $t \notin \text{dom}(\delta_2)$ (which are identified by T2 in Figure), and the cases for RET when $t \in F_2$ (which are identified by T3 in Figure).

LEMMA A.9. *Assume that $\sigma_1 R(\sigma_2, \delta_2, P_2)$.*

- (1) *If $(\sigma_2, \delta_2, P_2) \twoheadrightarrow (\sigma'_2, \delta'_2, P'_2)$ is an instance of CMT, Frc, or RET then $\sigma_1 \twoheadrightarrow^* \sigma'_1$ such that $\sigma'_1 R(\sigma'_2, \delta'_2, P'_2)$.*
- (2) *If $\sigma_1 \twoheadrightarrow \sigma'_1$ then $(\sigma_2, \delta_2, P_2) \twoheadrightarrow (\sigma'_2, \delta'_2, P'_2)$ such that $\sigma'_1 R(\sigma'_2, \delta'_2, P'_2)$.*

PROOF. Assume first that $(\sigma_2, \delta_2, P_2) \twoheadrightarrow (\sigma'_2, \delta'_2, P'_2)$. Let $\sigma_i = (I_i, s_i, C_i, F_i)$ and $\sigma'_i = (I'_i, s'_i, C'_i, F'_i)$. We prove that $\sigma'_1 R(\sigma'_2, \delta'_2, P'_2)$ and proceed by cases, first from the speculative to the OoO semantics.

(Case CMT) We get $\sigma_2 \twoheadrightarrow \sigma'_2$, and $t \notin \text{dom}(\delta_2)$, $t \leftarrow c?st M t_a t_v \in I_2$, $s_2(t) = v$, $t \notin C_2$ and $\text{bn}(\text{str-may}(\sigma_2, t)) \subseteq C_2$. Since t is not speculated we get $t \in I_{2, \text{retired}}$. Since $t \notin \text{dom}(\delta_2)$ we get $s_1(\tilde{t}) = \tilde{v}$ by cond. 4, $\tilde{t} \notin \tilde{C}_1$ by cond. 3, $\tilde{t} \in \tilde{I}_2$ by cond. 2, and, since all members of $\text{bn}(\text{str-may}(\sigma_2, t))$ are non-specified, by cond. 2, 3, 4, $\text{bn}(\text{str-may}(\sigma_1, \tilde{t})) \subseteq C_1$. It follows that $\sigma'_1 = (I_1, s_1, C_1 \cup \{\tilde{t}, F_1\})$, and the conditions 1-5. for σ'_1 and σ'_2 are easily verified.

(Case non-speculative Frc) For a non-speculative fetch we get $t \notin \text{dom}(\delta_2)$. Also, $\sigma_2 \twoheadrightarrow \sigma'_2$, $\text{step-param}(\sigma_2, \sigma'_2) = (\text{Frc}(I'_2), t)$, $t \leftarrow c?st \mathcal{PC} t_v \in I_2$, $s_2(t) = v$, $t \notin F$, $\text{bn}(\text{str-may}(\sigma_2, t)) \subseteq C_2$, $I'_2 = I_2 \cup I''_2$, $s'_2 = s_2$, $C'_2 = C_2$, and $F'_2 = F_2 \cup \{t\}$. Since $t \leftarrow c?st \mathcal{PC} t_v \in I_2$ and $t \notin \text{dom}(\delta_2)$ we obtain that $t \leftarrow c?st \mathcal{PC} t_v \in I_{2, \text{retired}}$ and hence that $\tilde{t} \leftarrow \tilde{c}?st \mathcal{PC} \tilde{t}_v \in I_1$. Also, $s_1(\tilde{t}) = \tilde{v}$ by cond. 4, $t \notin F_1$ by cond. 5, and $\text{bn}(\text{str-may}(\sigma_1, \tilde{t})) \subseteq F_1$. But then $\sigma_1 \twoheadrightarrow \sigma'_1$ where $I'_1 = I_1 \cup \tilde{I}_2$ (by wellformedness), $s'_1 = s_1$, $C'_1 = C_1$, and $F'_1 = F_1 \cup \{\tilde{t}\}$.

(Case RET and $t \notin F_2$) Assume that $\sigma_2 \twoheadrightarrow \sigma'_2$, $t \in \text{dom}(s_2)$, $\text{dom}(\delta_2(t)) \cap \text{dom}(\delta_2) = \emptyset$, $t \notin P_2$, and $(I_2, s_2, C_2, F_2) \sim_t (I_2, \delta_2(t), C_2, F_2)$ such that $I'_2 = I_2$, $s'_2 = s_2$, $C'_2 = C_2$, $F'_2 = F_2$, $\delta'_2 = \delta_2|_{\overline{\{t\}}}$, and $P'_2 = P_2$. By Lemma 5.3 $[l](I_2, s_2, C_2, F_2) = s_2(t)$. By cond. 4, whenever $t' \notin \text{dom}(\delta_2)$, $s_1(\tilde{t}') = s_2(t')$. We know that $s_1(\tilde{t}_i) = s_2(\tilde{t}_i)$ by 4. Let $t \leftarrow c?o \in I_2$ be the microinstruction bound to t in I_2 . Since $\text{dom}(\delta_2) \cap \text{dom}(\delta_2(t)) = \emptyset$ we know that $t \in I_{2, \text{uf}}$ and hence $\tilde{t} \in \text{dom}(I_1)$ by 2. It follows that \tilde{t} can be executed resulting in $I'_1 = I_1$, $s'_1 = s_1[\tilde{t} \mapsto s'_2(t)]$ such that $s'_1(\tilde{t}) = s'_2(\tilde{t})$, $C'_1 = C_1$, and $F'_1 = F_1$. This is sufficient to reestablish R.

(Case RET and $t \in F_2$) In this case we must additionally account for the microinstructions produced by t : i.e. the partition \hat{I}_2 of I_2 such that $\{t \leftarrow c?o\} < \hat{I}_2$. The microinstructions are not in $I_{2, \text{uf}}$, since $t \in \text{dom}(\delta_2)$, hence are not covered by the bijection $\tilde{\cdot}$. However, microinstructions in \hat{I}_2 are in $I'_{2, \text{uf}}$. For this reason, in order to reestablish the bisimulation, the OoO must perform a further step from σ'_1 and apply Frc to t . This allows to extend the bijection $\tilde{\cdot}$ to relate the newly decoded microinstructions to \hat{I}_2 .

(Converse direction) For the converse direction, from the OoO semantics to the speculative semantics the steps for commits and non-speculative retires follow the commit case above closely. The only delicate case is for EXE. So assume $\sigma_1 \rightarrow_{\text{OoO}} \sigma'_1$ such that $\tilde{t} = \tilde{t} \leftarrow \tilde{c}?o \in I_1$, $\tilde{t} \notin \text{dom}(s_1)$, $[\tilde{c}]s_1$ is true, $[l]\sigma_1 = (\tilde{v}, \tilde{I})$, $I'_1 = I_1$, $s'_1 = s_1[\tilde{t} \mapsto \tilde{v}]$, $C'_1 = C_1$, and $F'_1 = F_1$. We get that $t = t \leftarrow c?o \in I_{2, \text{uf}}$. There are two cases. Either $t \notin \text{dom}(s_2)$ (t has not yet been executed), or $t \in \text{dom}(s_2) \cap \text{dom}(\delta_2)$. In the former case, the execution step can be mirrored in the speculative semantics and then retired. In the latter case, the execution step can be retired directly. This completes the proof of lemma A.9. \square

We now obtain theorem 6.4 as a corollary of lemma A.7 and A.9. \square

A.4 MIL Constant Time Security: Theorem 7.3

The proof is done by showing that the relation R is a bisimulation for the OoO transition relation, where $\sigma R \sigma'$ iff $\sigma \approx \sigma'$ and there exist $\sigma_0 \sim_L \sigma'_0$ and n such that $\sigma_0 \twoheadrightarrow^n \sigma$ and $\sigma'_0 \twoheadrightarrow^n \sigma'$.

Let $(I, s, C, F) = \sigma$, $(I', s', C', F') = \sigma'$, $\sigma \twoheadrightarrow \sigma_1 = (I \cup I_t, s \cup s_t, C \cup C_t, F \cup F_t)$, and $\text{step-param}(\sigma, \sigma_1) = (\alpha, t)$. The proof is done by case analysis on α .


(Case EXE) The hypothesis of the rule ensure that $t = t \leftarrow c?o \in I$, $s(t) \uparrow$, $[c]s$, and $[l]\sigma = (v, l)$. The conclusion ensures that $\hat{s} = \{t \mapsto v\}$, $\hat{C} = \emptyset$, $\hat{F} = \emptyset$, and $\hat{I} = \emptyset$. Relation \approx ensures that $t \in I'$, $s'(t) \uparrow$, and $[c]s'$.

We must show that exists v' such that $[l]\sigma' = (v', l)$. For $o = e$ and $o = st \tau t_a t_v$ the proof is trivial. In fact, since $[l]\sigma$ then all free names of o are defined in σ and \approx ensures that the same names are defined in σ' . For $o = ld \tau t_a$, $[l]\sigma = (v, l)$ ensures that $\text{bn}(\text{str-act}(\sigma, t)) = \{t_s\}$, $\sigma(t_a) \downarrow$, and $\sigma(t_s) \downarrow$. Relation \approx ensures that $\sigma'(t_a) = \sigma(t_a)$, $\sigma'(t_s) \downarrow$, $I = I'$ (hence there are the same store instructions), and that for every store $t' \leftarrow c?st \tau t'_a t'_v$, $[c']\sigma = [c']\sigma'$, and $\sigma(t'_a) = \sigma'(t'_a)$. Therefore $\text{bn}(\text{str-act}(\sigma', t)) = \text{bn}(\text{str-act}(\sigma, t))$ and $[l]\sigma' = (\sigma'(t_s), l')$. Finally, since relation \approx guarantee that $(t_s \in C) \Leftrightarrow (t_s \in C')$ then $l' = l$.

These properties enable applying rule (EXE) to show that $\sigma' \xrightarrow{l}$

$(I', s' \cup \{t \mapsto s'(t_s)\}, C', F') = \sigma'_1$.

To prove that \approx is preserved we use Theorem 6.3. Let $t' \leftarrow c'?o' \in I$. Let N' be $fn(c')$ if o' is neither a load or a store; $fn(c') \cup \{t'_a\}$ if o' is a memory or register access and t'_a is the corresponding address; $fn(c') \cup \{t'_v\}$ if o' is a program counter update and t'_v is the corresponding value. If $t \notin N'$ then the proof is trivial, since $[c]s \cup \{t \mapsto v\} = [c]s = [c]s' = [c]s' \cup \{t \mapsto s'(t_s)\}$ (the same holds for the address in case of a resource accesses or program counter update).



For $t \in fn(c')$ we reason as follows. States σ_1 and σ'_1 are the $(n+1)$ -th states of two OoO traces $\rho = \sigma_0 \rightarrow^{n+1} \sigma_1$ and $\rho' = \sigma'_0 \rightarrow^{n+1} \sigma'_1$ such that $\sigma_0 \sim_L \sigma'_0$. There is a trace $\rho_1 = \sigma_0 \rightarrow^{n+1} \sigma_1 \rightarrow^m \sigma_s$ that has prefix ρ , such that $C(\sigma_s, t'')$ for every $t'' \leq \max(bn(I'))$. Notice that this state is “sequential”. Since in the OoO semantics the storage is monotonic then $\sigma_1(t) = \sigma_s(t)$. Theorem 6.3 permits to connect this trace to a sequential trace, which enables to use the MIL constant-time hypothesis. In fact, there exists an ordered execution π of ρ_1 that ends in σ_s : $\pi = \sigma_0 \rightarrow^{n+1+m} \sigma_s$. For the same reason, ρ' is a prefix of a trace ρ'_1 that ends in a sequential state σ'_s , $\sigma'_s(t) = \sigma'_1(t)$, and there exists a sequential trace π' of $n+1+m$ steps that ends in σ'_s . Since $C(\sigma_s, t'')$ and $C(\sigma'_s, t'')$ then $[c'](\sigma_s) \downarrow$ and $[c'](\sigma'_s) \downarrow$. Therefore, we can use the assumption on MIL constant-time to show that $\sigma_s \approx \sigma'_s$, hence $[c']\sigma_s = [c']\sigma'_s$. Since $t \in fn(c)$, either $[c'](\sigma_s) \uparrow$ and $[c'](\sigma'_s) \uparrow$ or $[c']\sigma_1 = [c']\sigma'_1$. The same reasoning is used if t' is a resource accesses and $t = t'_a$, or if t' is program counter update and $t = t'_v$.

(Case CMT) The hypothesis ensure that $t \leftarrow c?st \mathcal{M} t_a t_v \in I$, $s(t) = v$, $t \notin C$, $bn(str\text{-}may(\sigma, t)) \subseteq C$, and $s(t_a) = a$. The conclusion ensures that $s_t = \emptyset$, $C_t = \{t\}$, $F_t = \emptyset$, $I_t = \emptyset$, and $l = ds a$. Also, the invariant guarantees that $[c]s$.

The relation \approx ensures that $t \leftarrow c?st \mathcal{M} t_a t_v \in I'$, $\exists v'. s'(t) = v'$, $t \notin C'$, $s'(t_a) = a$, and $[s'] [c]$.

To show that $bn(str\text{-}may(\sigma', t)) = bn(str\text{-}may(\sigma, t))$ we use the same reasoning used to prove that $bn(str\text{-}act(\sigma', t)) = \{t_s\}$ of case EXE when $o = ld t_a$.

Relation \approx ensures that $bn(str\text{-}may(\sigma', t)) \subseteq C'$. Therefore we can apply rule (CMT) to show that $\sigma' \xrightarrow{l} (I', s', C' \cup \{t\}, F') = \sigma'_1$ hence $\sigma_1 \approx \sigma'_1$.

(Case FRC) The hypothesis of the rule ensure that $t \leftarrow c?st \mathcal{P}C t_v \in I$, $s(t) = v$, $t \notin F$, $bn(str\text{-}may(\sigma, t)) \subseteq F$. The conclusion ensures that $s_t = \emptyset$, $C_t = \emptyset$, $F_t = \{t\}$, $I_t = translate(v, \max(I))$, and $l = il v$. Also, the invariant guarantees that $[c]s$.

The relation \approx ensures that $t \leftarrow c?st \mathcal{P}C t_v \in I'$, $\exists v'. s'(t) = v'$, $t \notin F'$, and $[c]s'$.

To show that $bn(str\text{-}may(\sigma', t)) = bn(str\text{-}may(\sigma, t))$ we use the same reasoning used to prove that $bn(str\text{-}act(\sigma', t)) = \{t_s\}$ of case

EXE when $o = ld t_a$.

Relation \approx ensures that $bn(str\text{-}may(\sigma', t)) \subseteq F'$. We can apply rule (FRC) to show that $\sigma' \xrightarrow{il v'} (I' \cup translate(v', \max(I)), s', C', F' \cup \{t\}) = \sigma'_1$.

To show that $v' = v$ and \approx is reestablished for $translate(v, \max(I))$ we use a similar reasoning to case EXE. We find two sequential traces that end with the fetch of t and use MIL constant time to show that the value used for the PC update must be the same and that the parameter and conditions of the newly decoded microinstructions are equivalent in the two states.