

Robust Constant-Time Cryptography

Matthew Kolosick^{*} Basavesh Ammanaghata Shivakumar[†] Sunjay Cauligi[‡]
Marco Patrignani[‡] Marco Vassena[§] Ranjit Jhala^{*} Deian Stefan^{*}
^{*}UC San Diego [†]MPI-SP [‡]University of Trento [§]Utrecht University

ABSTRACT

Cryptographic *library* developers take care to ensure their library does not leak secrets even when there are (inevitably) exploitable vulnerabilities in the *applications* the library is linked against. To do so, they choose some class of application vulnerabilities to defend against and hardcode protections against those vulnerabilities in the library code. A single set of choices is a poor fit for all contexts: a chosen protection could impose unnecessary overheads in contexts where those attacks are impossible, and an ignored protection could render the library insecure in contexts where the attack is feasible.

We introduce **ROBOCOP**, a new methodology and toolchain for building secure *and* efficient applications from cryptographic libraries, via four contributions. First, we present an operational semantics that describes the behavior of a (cryptographic) library executing in the context of a potentially vulnerable application so that we can precisely specify what different attackers can observe. Second, we use our semantics to define a novel security property, *Robust Constant Time* (RCT), that defines when a cryptographic library is secure *in the context of* a vulnerable application. Crucially, our definition is parameterized by an attacker model, allowing us to factor out the classes of attackers that a library may wish to secure against. This refactoring yields our third contribution: a compiler that can synthesize bespoke cryptographic libraries with security tailored to the specific application context against which the library will be linked, guaranteeing that the library is RCT in that context. Finally, we present an empirical evaluation that shows the **ROBOCOP** compiler can automatically generate code to efficiently protect a wide range (over 540) of cryptographic library primitives against three classes of attacks: read gadgets (due to application memory safety vulnerabilities), speculative read gadgets (due to application speculative execution vulnerabilities), and concurrent observations (due to application threads), with performance overhead generally under 2%, thus freeing library developers from making one-size-fits-all choices between security and performance.

1 INTRODUCTION

“Don’t roll your own crypto” is a well known adage directed at application developers when they are considering using cryptography in their code. Instead developers are exhorted to use cryptographic libraries written by experts who have hopefully learned the hard-won lessons of decades of cryptographic and software security research and practice. For such cryptographic library developers, as well as algorithm designers, one of those hard-won lessons is that cryptographic code must be *constant time* [8]. More recently (due to microarchitectural attacks like Spectre [32]) this lesson has been expanded to the requirement that, under certain circumstances, it is also crucial for cryptographic code to be *speculatively constant time* [13]. Together these ensure that the code in the library does not leak

```
1 static int stream_ref(u8 *c, u64 clen, u8 *n, u8 *k) {
2     ... u8 kcopy[32]; ...
3     for (i = 0; i < 32; i++) { kcopy[i] = k[i]; }
4     ...
5     while (clen >= 64) {
6         crypto_core_salsa20(c, in, kcopy, NULL); ...
7     }
8     ...
9     sodium_memzero(kcopy, sizeof kcopy);
10    return 0;
11 }
```

Figure 1: An excerpt from the reference implementation of Salsa20 in LibSodium.

secrets (such as cryptographic keys) through either timing channels or speculative execution, respectively, and significant work has gone into developing theory, tools, and rules of thumb to ensure that cryptographic code is (speculatively) constant time [5, 14].

Attacks via application vulnerabilities. Sadly, this is not enough. While a constant time cryptographic library will not *itself* leak secrets, it is but one component executing within the context of a larger *application*. Security vulnerabilities in application code could themselves lead to inadvertent disclosure of secrets, no matter how careful library authors were to avoid vulnerabilities.

Library authors are keenly aware of this problem and routinely add protections to harden their code against application vulnerabilities. For example, consider the excerpt of the implementation of the Salsa20 stream cipher [9] from LibSodium [17] shown in Figure 1. We might hope that the fact that the (elided) body is constant time, suffices to ensure that the secrets in `kcopy` are not leaked. However the *classic* constant time guarantee only ensures that `stream_ref` does not leak the secrets: it makes no guarantees about what happens if there is a memory safety vulnerability in the application linked against the library. Such a vulnerability can lead to a *read-gadget* that may be used to exfiltrate the secrets in `kcopy`! To defend against such gadgets, LibSodium’s developers take care to zero the intermediate memory used in `stream_ref` (Line 9) to ensure those secrets cannot be leaked through exploitable memory vulnerabilities that reside in the application. Unfortunately, compiler optimizations like dead-store elimination can remove secret scrubbing code and nullify the efforts of LibSodium’s developers [62].

Security vs. performance. Read gadgets are but one of several possible classes of attacks that library developers must defend against. For each class of attacks, the library developer must either manually add the relevant defences to their code, or make an explicit choice *not* to do so (typically due to prohibitive overheads). Consequently, library code “bakes in” some subset of possible protections against application vulnerabilities: which may be either unnecessary or insufficient depending upon the application context. For example,

LibSodium’s zeroization protection against read gadgets is unnecessary when linked against a memory safe Rust application. On the other hand, LibSodium’s code is insufficient against Spectre attacks [55]. The library’s authors chose to elide an appropriate fence to protect against speculative read gadgets exfiltrating the contents of `kcopy` as *all* the clients of the library would have to suffer the corresponding performance degradation, not just the ones where Spectre was a legitimate concern.

In a nutshell, cryptographic library developers are currently in a difficult position: they must make one-size-fits-all security-performance trade offs by manually inserting fragile protections orthogonal to the cryptographic algorithms and protocols they are implementing. Luckily, secure compilers offer a promising solution to escape this dilemma [43]. Library users are in a better position to make security-performance trade offs and can provide a compiler with security policies to be automatically enforced through inserted protections. To realize this vision, we introduce `ROBOCOP`, a new methodology and toolchain for building secure *and* efficient applications from cryptographic libraries. We develop `ROBOCOP` via four concrete contributions:

1. Abstraction: libraries and attackers (§3, §4.1). Our first contribution is a formal operational semantics describing the behavior of a library executing within a potentially vulnerable application. We further define a semantics capturing a high-level, abstract model of speculative execution, based on the notion of a speculation oracle that “guesses” the result of evaluating an expression and then later rolling back or committing if the guesses were correct. These semantics provide a unified setting where we can precisely state different attackers’ observations, guiding the design of our protections.

2. Specification: robust constant time (§4.2). Our second contribution is a novel security property, *robust constant time* (RCT), using our model to precisely define security for a cryptographic library running *in the context of* a potentially vulnerable application. Crucially, our definition is parameterized by an attacker model, capturing the set of attackers that a library is supposed to be secure against. We further define a speculative version, *robust speculative constant time*, capturing constant time in the presence of Spectre.

3. Implementation: The `ROBOCOP` compiler (§5). By factoring out the security assumptions about the context, our notion of RCT enables our third contribution: a compiler that takes a cryptographic library and synthesizes a bespoke binary tailored to the application context against which the library will be linked. To do so, we show how to map each kind of attacker to a concrete code transform that provably, with respect to our definition of RCT, protects against that attacker. Thus, our `ROBOCOP` compiler lets library developers focus on implementing constant-time cryptographic algorithms, without having to worry about baking in a fixed set of potentially inefficient or insecure protections against application vulnerabilities. Instead, protections can be automatically inserted depending on the application context, thereby ensuring the same library code can be securely *and* efficiently reused in all contexts.

4. Evaluation: `SUPERCOP` (§6). Finally, our fourth contribution is an empirical evaluation that shows that our `ROBOCOP` compiler can automatically generate protections for a wide range of cryptographic library code defending against a variety of attacks. Here we modify the `SUPERCOP` [57] cryptographic benchmarking suite.

We instrument `SUPERCOP` to generate and measure the overhead of protecting against three classes of attacks: read gadgets (due to memory safety vulnerabilities in the application), speculative read gadgets (due to speculative execution in the application), and concurrent observations (due to threads in the application). In our test suite of 542 different implementations of cryptographic operations, we show that our `ROBOCOP` compiler can automatically generate code that is secure against application vulnerabilities with the majority of overheads under 2%, thereby demonstrating that RCT reconciles the tension between security and efficiency when reusing cryptographic libraries in different application contexts.

2 OVERVIEW

Every application that works with sensitive or personal user data uses cryptography to ensure the confidentiality or integrity of the data. These cryptographic operations typically rely upon sophisticated mathematics and are notoriously difficult to get right: bugs or security vulnerabilities within cryptographic code risk leaking critical secret keys or data, which could compromise the security of the whole application. Thus, cryptographic operations are typically implemented and encapsulated within *libraries* that are carefully authored and audited by cryptographic experts. These libraries provide trusted implementations of cryptographic operations (e.g. LibSodium [17]) or protocols (e.g. OpenSSL [4]), that can then be widely reused by developers—without requiring cryptographic expertise—to build secure applications.

In addition to correctly implementing cryptographic algorithms, the developers of cryptographic libraries must carefully ensure their code meets certain *generic* security requirements. For example, they must ensure that their libraries are *constant time*, meaning that the execution time of the library must be independent of the values of the secret data that the library operates over. Otherwise, an attacker can measure the timing variations to learn whether a secret conditioned branch or operation went one way or the other, and from that, eventually recover the data itself. There has been significant work [13, 15] on characterizing when a cryptographic library is constant-time and designing recipes to write constant time code, and this work guides both the design of cryptographic algorithms and their implementation in cryptographic libraries.

2.1 Application (attacker) assumptions

Sadly, timing leaks are not the only attacker capability cryptographic library developers need worry about. Recall that, ultimately, libraries are not executed in isolation: they are linked against *applications* written by non-expert developers, leading to another source of vulnerabilities through which secrets can be leaked.

Defending against application vulnerabilities. Consider an application written in C that uses the LibSodium library. If this application has a *buffer overflow* leading to a memory read then an attacker targeting the whole program now has the ability to read any cryptographic secrets that are left accessible in memory, completely bypassing the need for timing channel-based attacks.

In fact, the developers of LibSodium are keenly aware of the need to defend against these potential vulnerabilities in applications where the library may be used. Figure 1 shows an excerpt of LibSodium’s reference implementation of the Salsa20 stream

cipher. `stream_ref` takes as input a key `k` and a nonce `n` and outputs a length `clen` stream of pseudo random bytes in the buffer `c`. The developers make a copy, `kcop`, of the key buffer, `k`. This copied buffer contains secret data that must not be left on the stack, as otherwise a buffer overflow in the application would let an attacker read the secret key off the stack. Thus, in anticipation of LibSodium being used in the context of a vulnerable application, on [Line 9](#) the LibSodium developers invoke `sodium_memzero` to zero out the contents of `kcop`, thereby ensuring that its value is inaccessible regardless of memory vulnerabilities from the application.

Attacker (application) assumptions. In general, cryptographic libraries must defend against vulnerabilities in their host applications. To capture these threats we classify the application vulnerabilities (or attacker capabilities) as *assumptions* about application behavior.

1. Attacks via memory unsafety. The first class of assumptions, illustrated by the code in [Figure 1](#), is that owing to memory unsafety, there exist *memory read gadgets* within the application. Zeroing buffers (like `kcop`) is one key defense against such gadgets, but does nothing to prevent the attacker from reading the original version of the key `k` which remains unzeroed. To protect `k` from read gadgets, libraries like Libsodium offer memory protection APIs which must be manually inserted and toggled on and off by application developers, and hence, are prone to incorrect usage.

2. Attacks via speculation. If the application is written in a memory safe language like Rust, then the library developer need not fret about read gadgets, and can avoid the overhead of zeroing out secrets. However, even in this setting, the library developer must contend with the spectre of hardware speculation and the host of attendant vulnerabilities [24, 32, 33, 37, 52, 61]. There has been work on extending constant time protections to speculation based vulnerabilities, but this work focuses solely on protecting cryptographic code *itself* from leaking secrets due to speculation. Owing to the various overheads imposed by such protection, it is currently unreasonable to apply the same protections to the entirety of application code. As such, library developers may have to contend with attacks based on Spectre vulnerabilities *in the application* [38].

Indeed, in the case of `stream_ref` in [Figure 1](#), Spectre leads to a potential security issue with the clearing of `kcop`: The LibSodium developers forgo an appropriate memory fence in `sodium_memzero`, leading to the possibility of the zeroed memory being read by speculatively executing application code *before* it is zeroed [55].¹ The fence was eschewed due to its performance overhead: *all* clients of Libsodium would suffer the performance degradation for protecting against Spectre. By manually adding or changing protections, Libsodium’s developers are implicitly restricting their defenses to certain classes of attackers: i.e. they are assuming only that the vulnerable application may contain *non-speculative* read gadgets.

3. Attacks via concurrency. Finally, the last category of application assumptions that we consider is whether the application is *concurrent*. In a concurrent context, work like Spectre-Declassified [53] has shown that an attacker can recover secret information if they can observe intermediate results, thus enhancing the reach of (speculative) read gadgets. In fact, the possibility of such observations

¹There are also well-documented issues with implementing zeroing functions in high-level code that mean that even without Spectre the zeroing is only best effort [46].

```

1  static int stream_ref(u8 *c, u64 clen, u8 *n, u8 *k) {
2      mpk_allow_access();
3      switch_to_protected_stack();
4      u8 *c_internal = mpk_malloc(clen);
5      int result = stream_ref_cloned(c_internal, clen, n, k);
6      memcpy(c_internal, c, clen);
7      switch_to_unprotected_stack();
8      clear_scratch_registers();
9      mpk_disable_access();
10     return result;
11 }

```

Figure 2: ROBOCOP protections applied to LibSodium’s Salsa20. Colors correspond to [Figure 11](#).

are cited by the LibSodium developers as a reason to forgo the performance penalty of a memory fence in `sodium_memzero` [55].

2.2 Robust constant time

Unlike with the property of constant time, there does not exist a security property capturing when a cryptographic library is secure when running *in the context of* a potentially vulnerable application. To address this and to capture the different application assumptions that a cryptographic library developer needs to consider we introduce the notion of *robust (speculative) constant time* (RCT). Like other robustness properties [1, 19, 20, 44, 45, 50, 54], a cryptographic library being robustly constant time captures that the library does not leak secrets *when linked against a context (application)*.

RCT serves as a formal security model for how cryptographic code is actually developed and used: protections are applied to the libraries with the goal that their use within an as yet unknown application will remain secure. Attacker capabilities can then be directly expressed as assumptions about the contexts in which the library code will run. For instance, LibSodium’s implicit protections can be explicitly specified as providing RCT assuming only the presence of single-threaded attacks based on memory unsafety (i.e. the presence of gadgets that can perform out-of-bounds reads).

2.3 A robust constant time compiler

By factoring out security assumptions about the context, our notion of RCT lets us design and develop a compiler that takes as input: (1) an implementation of a cryptographic library² and (2) an explicit set of *assumptions* about application/attacker capabilities and then synthesizes a protected library that is guaranteed to be robustly constant time with respect to the given attacker. Library developers can then focus on implementing cryptographic algorithms and their library can be used securely and efficiently in a variety of contexts.

Bespoke protection. Cryptographic libraries like LibSodium are designed to be used in a broad range of applications. As we saw with `libsodium_memzero`, the current state of manually baking protections into library code means that a single decision is made trading off between performance and which contexts the library

²Our compiler assumes that the library is already (speculatively) constant-time. In §4.2 we will discuss how our definition of RCT guarantees that our robust protections are orthogonal to existing (speculative) constant time protections, thus allowing library developers to use any automated tool or manual technique to meet this assumption.

is robust against. Instead, by factoring the protections out of the library and placing them in the compiler, **ROBOCOP**, developers can tailor protections to exactly the level required by the particular application context. For example, when used in a memory safe Rust application we can omit the unnecessary zeroing from `LibSodium`.

Efficient protection via wrapping and MPK. Our notion of RCT lets us use modern hardware memory protections, in particular Intel™ Memory Protection Keys (MPK), to protect library code incredibly efficiently. Our key insight in **ROBOCOP** is to allocate secret keys and carry out secret computation within a protected memory region. To do so, **ROBOCOP** wraps the library’s external API functions, as illustrated by the protected version of `stream_ref` in [Figure 2](#). We first enable access to the protected memory region and switch to a stack within protected memory. This protected stack ensures all intermediate computation remains protected. [Line 4](#) shows the parameterized nature of **ROBOCOP** where we allocate a copy of the output buffer within protected memory. This allocation is *only* needed when (1) the underlying implementation uses the buffer for intermediate computation, and (2) the application context is *concurrent*. If these conditions are not met, **ROBOCOP** does not generate the extra allocation as the intermediate results *cannot* be observed by an attacker, and hence, do not need to be protected. The wrapper function then calls the original library implementation, `stream_ref_cloned`. After encryption, we copy the internal output buffer back to the publicly visible `c` buffer (again, only if the attacker assumptions require doing so.) The wrapper then switches back to the unprotected stack, clears scratch registers (if in Spectre protection mode where these registers may contain secret values that could be speculatively read), and finally disables access to the protected memory region before returning to the application.

3 SECURITY SEMANTICS

To formally ground the varying notions of robust constant time and the attacker models we develop a high-level, stateful calculus, $\lambda_{\bar{\Delta}}$, whose syntax is shown in [Figure 3](#). Syntactically, $\lambda_{\bar{\Delta}}$ is relatively standard, following λ_{rust} , `CompCert`, and others [30, 36] in employing a block-based memory model, i.e. memory is structured as “disjoint”, fixed width blocks addressed via a block label and offset ($z_b[z_o]$). New blocks are allocated with `newp e` with e the size of the block and the protection label p determining whether the block is allocated in a protected or unprotected memory “page”. The size, set of values, and memory page are tracked in the second component of the non-speculative states (S). Correspondingly there is a protection operation, `protectp`, which models hardware memory protection. `protectp` sets the memory access policy in the first component of the state: `public` only allows access to the public memory page whereas `protected` allows access to all memory. Pointers can be decomposed using `get-block e` and `get-offset e`. Dereferences are written `!e` and assignments are written `eptr := eval`.

We equip $\lambda_{\bar{\Delta}}$ with these semantics: a non-speculative semantics capturing a trace of events that we use to define our attacker models (§3.1) a novel, high-level speculative semantics (§3.2), and (speculative) concurrent semantics capturing a passive observer.

labels	ℓ	::=	app lib
protection	p	::=	public protected
values	v	::=	$z \mid z[z] \mid \lambda_{\ell} \bar{x}.e$
expressions	e	::=	$v \mid x \mid x\{v\} \mid \text{op}(\bar{e}) \mid e[e] \mid !e \mid \text{new}_p e \mid e := e \mid \text{get-block } e \mid \text{get-offset } e \mid e(\bar{e}) \mid \text{protect}_p \mid \text{if } e \text{ then } e \text{ else } e \mid \text{fence } e \mid e; e$
non-speculative states	S	:	$p \times (\mathbb{Z} \rightarrow \{\text{size} : \mathbb{Z}, p : p, v : [\text{size}] \rightarrow v\})$
events	ϵ	::=	$\delta^{\ell} \mid \tau^{\ell \rightarrow \ell}$
domain events	δ	::=	$\mu \mid \text{call } f(\bar{v}) \mid \text{branch } v \mid \text{fence} \mid 0$
transition events	τ	::=	$\text{call } f(\bar{v}) \mid \text{ret } v \mid \text{begin} \mid \text{end } v$
memory events	μ	::=	$\text{new}_p z@z \mid \text{read } v \leftarrow z[z] \mid \text{write } v \mapsto z[z] \mid \text{protect}_p$
constant time events	c	::=	$0 \mid \text{branch } v \mid \text{read } \leftarrow z[z] \mid \text{write } \mapsto z[z] \mid \text{call } f \mid \text{end } v$
speculation directives	d	::=	nonspec spec v fence
speculation frame	Ξ	:	$(\bar{S}, e) \mid (S, e, \bar{\delta}, \bar{\mu})$
speculation oracles	A	:	$\{A : \text{Type}, a : A, \text{spec} : A \times S \times e \rightarrow A \times d\}$
speculative states	Φ	:	$\{S : S, A : A, \Xi : \Xi\}$

Figure 3: $\lambda_{\bar{\Delta}}$ syntax

3.1 Nonspeculative trace semantics

Before describing the nonspeculative semantics we go over the structure of the traces that the semantics is designed to capture.

Traces. We are interested in capturing three aspects of the execution: (1) *who* (which party, app or lib) is executing code at a given time as well as the transfer of control between the parties, (2) *what memory* each party accesses, and (3) the internal branching which will be used to define the various constant time properties. To this end each reduction is labeled with an event, ϵ , whose syntax is shown in [Figure 3](#). Labels are also attached to every function ($\lambda_{\ell} \bar{x}.e$) to distinguish application and library code. An event is either a labeled *domain event*, δ^{ℓ} , which captures an event executed by the party ℓ or a *transition event*, $\tau^{\ell \rightarrow \ell}$, which captures the transfer of control between the two parties.

The main transition events are call $f(\bar{v})$ and `ret v` which represent a call to the function f with arguments \bar{v} and returning from a function with return value v . Beyond the call and return events, there are `begin` and `end v` events that capture the (implicit) beginning and end of a trace. Domain events are either a *memory event* (μ), a call $f(\bar{v})$ event (capturing an function call that stays within a single domain), a `branch v` event (capturing branching on the value v), or the empty event 0 . Memory events are one of an allocation, a protection operation, a read, or a write and track the data associated with each operation (e.g. the value read/written and the location it was read/written from/to).

Transition operational semantics. Our nonspeculative semantics are split between two labeled reduction judgments: top-level transition reductions ([Figure 4](#)) and domain reductions (an excerpt is shown in [Figure 5](#)). Transition reductions are of the form

$$\boxed{\langle S \mid \overline{K}^\ell :: e^\ell \rangle \xRightarrow{\epsilon} \langle S \mid \overline{K}^\ell :: e^\ell \rangle}$$

RED-CALL

$$\epsilon = \begin{cases} \text{call } (\lambda \bar{x}. e)(\bar{v})^\ell & \text{when } \ell_f = \ell \\ (\text{call } (\lambda \bar{x}. e)(\bar{v}))_{\ell \rightarrow \ell_f} & \text{otherwise} \end{cases}$$

$$\frac{}{\langle S \mid \overline{K}^{\ell'} :: K[(\lambda \bar{x}. e)(\bar{v}))^\ell] \rangle \xRightarrow{\epsilon} \langle S \mid \overline{K}^{\ell'} :: K^\ell :: e[\bar{v}/\bar{x}]^{\ell'} \rangle}$$

RED-RET

$$\epsilon = \begin{cases} 0^\ell & \text{when } \ell_K = \ell \\ (\text{ret } v)_{\ell \rightarrow \ell_K} & \text{otherwise} \end{cases}$$

$$\frac{}{\langle S \mid \overline{K}^{\ell'} :: K^{\ell_K} :: v^\ell \rangle \xRightarrow{\epsilon} \langle S \mid \overline{K}^{\ell'} :: K[v]^{\ell_K} \rangle}$$

RED- β

$$\frac{\langle S \mid e \rangle \xrightarrow{\delta} \langle S' \mid e' \rangle}{\langle S \mid \overline{K}^{\ell'} :: K[e]^\ell \rangle \xRightarrow{\delta^\ell} \langle S' \mid \overline{K}^{\ell'} :: K[e']^\ell \rangle}$$

Figure 4: Non-speculative trace semantics transition reductions

$\langle S \mid \overline{K}^\ell :: e^\ell \rangle \xRightarrow{\epsilon} \langle S \mid \overline{K}^\ell :: e^\ell \rangle$. The state S tracks the memory and the current memory access level. To explain the control stack $\overline{K}^\ell :: e^\ell$ we examine the rule **RED-CALL**. The top of our stack (e^ℓ) is the (potentially mid-reduction) body of the currently executing function (with label ℓ). For **RED-CALL** we are reducing a function call in the evaluation context K . We push this continuation (i.e. where the called function will return to) onto the stack of labeled continuations ($\overline{K}^{\ell'}$) and then use the substituted body as the new execution frame on the stack. If the current label, ℓ , and the label of the function we are calling, ℓ_f , differ then we emit a transition event with label $\ell \rightarrow \ell_f$ otherwise we omit a domain event for the call.

Dually, rule **RED-RET** handles returning from a function. This is a transfer of control flow from ℓ , the label of the currently executing function, back to ℓ_K , the function caller. The return value is plugged into the top continuation on the stack and a corresponding return event is generated (we ignore same domain returns). The last “transition” reduction rule, **RED- β** , dispatches to the domain reduction relation, and labels the domain event δ with the current label.

Domain operational semantics. The majority of the domain rules are standard and produce an empty trace event. Conditionals are also standard, but produce a branch event based on the condition. The rule **β -NEW** takes a protection domain p in which to allocate a new block of size z , checking that our current access level allows writing to the domain p using the “can-access” judgment $S.p \sqsubseteq p$. The rule **β -PROTECT** handles setting this access level.

The most notable of the reduction rules are those related to dereferencing and writing through pointers. The rules mirror each other so we will focus on dereferencing as it is simpler. In the rule **β -DEREF** we are dereferencing the pointer $z_b[z_o]$ with block label z_b and offset into that block z_o . To actually obtain the value at that location the following must hold: (1) the block must be accessible ($\text{accessible}(S, z_b)$), (2) the offset must be within the allocated size

$$\boxed{\langle S \mid e \rangle \xrightarrow{\delta} \langle S \mid e \rangle}$$

β -DEREF

$$\frac{\text{accessible}(S, z_b) \quad z_o \in [S(z_b).size] \quad v = S(z_b).v(z_o)}{\langle S \mid !(z_b[z_o]) \rangle \xrightarrow{\text{read } v \leftarrow z_b[z_o]} \langle S \mid v \rangle}$$

β -DEREF-OOB

$$\frac{z_b \notin \text{dom}(S) \vee z_o \notin [S(z_b).size] \quad z'_b \in \text{dom}(S) \quad z'_o \in [S(z'_b).size] \quad v = S(z'_b).v(z'_o) \quad \text{accessible}(S, z'_b)}{\langle S \mid !(z_b[z_o]) \rangle \xrightarrow{\text{read } v \leftarrow z'_b[z'_o]} \langle S \mid v \rangle}$$

β -SUBST

$$\frac{}{\langle S \mid x\{v\} \rangle \xrightarrow{0} \langle S \mid v \rangle}$$

β -PROTECT

$$\frac{}{\langle S \mid \text{protect}_p \rangle \xrightarrow{\text{protect}_p} \langle S[p := p] \mid 0 \rangle}$$

β -NEW

$$\frac{z > 0 \quad z_b = \text{fresh}(S) \quad S.p \sqsubseteq p \quad S' = S[z_b := \{size = z, v = \perp, p = p\}]}{\langle S \mid \text{new}_p z \rangle \xrightarrow{\text{new}_p z @ z_b} \langle S' \mid z_b[0] \rangle}$$

Figure 5: Non-speculative trace semantics domain reduction excerpts

of the block ($z_o \in [S(z_b).size]$), and (3) a value must have been written to the block at the offset z_o ($v = S(z_b).v(z_o)$). When these conditions are met the value is read and a corresponding read trace event is generated. On the other hand, if either of the latter two conditions are not met the dereference is considered out-of-bounds and the rule **β -DEREF-OOB** applies instead. Here, we model the out-of-bounds read as a nondeterministic read from an arbitrary, valid, and accessible location $z'_o[z'_b]$.

3.2 Speculative semantics

To model Spectre and speculative execution broadly we define a second operational semantics for λ_{CT} . Instead of modeling a specific version of speculative execution we seek to capture a high-level essence of speculation, namely that speculation is the combination of guessing how an expression might evaluate and then either rolling back or committing if the guess was correct. That is, different types of speculation can be captured as the following sequence: First, instead of evaluating an expression, make up the value you think it would evaluate to and continue running using that value instead. While running “under speculation” check if any operation invalidates the speculative guess. Then, once you hit a “fence”, rollback to the speculation point if the guess was invalid, or commit the effects of the skipped expression and continue evaluating.

We capture (excerpts of) these notions in Figures 6 and 7. Figure 6 defines the top-level reduction relation $\langle \Phi \mid e \rangle \xrightarrow{\delta} \langle \Phi \mid e \rangle$. Speculative states, Φ , extend the nonspeculative state with a speculation oracle (A) and a stack of speculation frames (Ξ). Speculation frames

$$\begin{array}{c}
\boxed{\langle \Phi \mid e \rangle \xrightarrow{\bar{\delta}} \langle \Phi \mid e \rangle} \\
\text{SPEC-NONSPEC} \\
(a', \text{nonspec}) = \Phi.A.\text{spec}(\Phi.A.a, \Phi.S, e) \\
\frac{\langle \Phi \mid K[e] \rangle \xrightarrow{\bar{\delta}} \langle \Phi' \mid K[e'] \rangle}{\langle \Phi \mid K[e] \rangle \xrightarrow{\bar{\delta}} \langle \Phi' [a := a'] \mid K[e'] \rangle} \\
\text{SPEC-SPEC} \\
(a', \text{spec } v) = \Phi.A.\text{spec}(\Phi.A.a, \Phi.S, e) \quad \text{nonfinal}(\langle \Phi \mid K[e] \rangle) \\
\frac{\langle \Phi \mid e \rangle \xrightarrow{\bar{\delta}^*} \langle \Phi' \mid v' \rangle \quad \bar{\Xi}' = \text{makeFrame}_{v=v'}(\Phi.S, K[e], \bar{\delta}) :: \Phi.\Xi}{\langle \Phi \mid K[e] \rangle \xrightarrow{0} \langle \Phi[\Xi := \bar{\Xi}', a := a'] \mid K[v] \rangle} \\
\text{SPEC-FENCE} \\
(a', \text{fence}) = \Phi.A.\text{spec}(\Phi.A.a, \Phi.S, e) \\
\frac{\text{fence } \langle \Phi \mid K[e] \rangle \text{ to } \langle \Phi' \mid e' \rangle_{\bar{\delta}}}{\langle \Phi \mid K[e] \rangle \xrightarrow{\bar{\delta}} \langle \Phi' [a := a'] \mid e' \rangle} \\
\text{SPEC-}\beta \\
\frac{\langle \Phi.S \mid \bullet :: K[e]^\ell \rangle \xrightarrow{\epsilon} \langle S' \mid \bar{K}' :: e'^{\ell'} \rangle \quad \neg\text{stalled}(\Phi.\Xi, \delta(\epsilon)) \quad \bar{\Xi} = \text{addEvent}(\Phi.\Xi, \delta(\epsilon))}{\langle \Phi \mid K[e] \rangle \xrightarrow{\delta(\epsilon)} \langle \Phi[S := S', \Xi := \bar{\Xi}] \mid \bar{K}'[e'] \rangle}
\end{array}$$

Figure 6: Small step speculative semantics

come in two forms, a “mispeculation” frame, (\widehat{S}, e) , capturing that the current speculation is invalid and will be rolled back to the state S and expression e and “in progress” frames, $(S, e, \bar{\delta}, \bar{\mu})$, capturing that the current speculation is potentially valid. The $\bar{\delta}$ and $\bar{\mu}$ components capture the trace of events of the skipped expression and the subsequent memory events, respectively. The skipped events are used if we commit a speculative frame, “replaying” the events that were speculatively passed over, and the memory events are used to determine whether speculation is invalid (discussed below).

The speculation oracle parameter of the semantics captures the attacker control over speculation and has three components: the type (A) of the “microarchitectural state”, the current microarchitectural state (a), and a microarchitectural state “speculation” function ($\text{spec} : A \times S \times e \rightarrow A \times d$) capturing attacker decisions about when to speculate. The function takes the current microarchitectural state, the current nonspeculative state, and the current expression, updates the microarchitectural state, and returns a speculation directive, d . This directive says whether we will be speculating with the guessed value v ($\text{spec } v$), not speculating (nonspec), or performing a fence operation (fence).

To see how speculation plays out we will go over the three corresponding rules: **SPEC-NONSPEC**, **SPEC-SPEC**, and **SPEC-FENCE** and how they capture speculation in the classic Spectre exploit of bypassing a bounds check when evaluating `if $i\{100\} < 10$ then $!b[i\{100\}]$ else 0` (where the size of the block b is 10). The term $i\{100\}$ captures a delayed substitution: earlier in the execution the value 100 was

$$\begin{array}{c}
\boxed{\text{stalled}(\bar{\Xi}, \delta)} \\
\text{STALL-FENCE} \\
\frac{\Phi.\Xi \neq \bullet}{\text{stalled}(\Phi, \text{fence})} \\
\text{STALL-READ} \\
\frac{\Phi.\Xi = \bar{\Xi} :: \bar{\Xi} \quad (\text{protect}_p \in \bar{\Xi}.\bar{\delta} \wedge \Phi.S(z_b).p = \text{protected}) \quad \vee (\text{stalled}(\Phi[\Xi := \bar{\Xi}], \text{read } v \leftarrow z_b[z_o]))}{\text{stalled}(\Phi, \text{read } v \leftarrow z_b[z_o])} \\
\boxed{\bar{\Xi} = \text{addEvent}(\bar{\Xi}, \delta)} \\
\text{ADD-BAD-READ} \\
\frac{v_r \in \text{writeLocs}(\bar{\delta})}{(\widehat{S}, e) = \text{addEvent}((S, e, \bar{\delta}, \bar{\mu}) :: \bar{\Xi}, \text{read } v \leftarrow v_r)} \\
\text{ADD-WRITE} \\
\frac{\bar{\mu}' = \bar{\mu} \diamond \text{write } v \mapsto v_w}{(S, e, \bar{\delta}, \bar{\mu}') = \text{addEvent}((S, e, \bar{\delta}, \bar{\mu}) :: \bar{\Xi}, \text{write } v \mapsto v_w)} \\
\boxed{\text{fence } \langle \Phi \mid e \rangle \text{ to } \langle \Phi \mid e \rangle_{\bar{\delta}}} \\
\text{FENCE-ROLLBACK} \\
\frac{\Phi.\Xi = (\widehat{S}, e') :: \bar{\Xi}}{\text{fence } \langle \Phi \mid e \rangle \text{ to } \langle \Phi[S := S', \Xi := \bar{\Xi}] \mid e' \rangle_{\bullet}} \\
\text{FENCE-COMMIT} \\
\frac{\Phi.\Xi = (S, e', \bar{\delta}, \bar{\mu}) :: \bar{\Xi} \quad S' = \text{commit}(\Phi.S, \bar{\delta} \diamond \bar{\mu}) \quad \bar{\Xi}' = \text{addEvents}(\bar{\Xi}, \bar{\mu})}{\text{fence } \langle \Phi \mid e \rangle \text{ to } \langle \Phi[S := S', \Xi := \bar{\Xi}'] \mid e \rangle_{\bar{\delta}}}
\end{array}$$

Figure 7: Speculative semantics auxiliary definitions

substituted for i .³ We first evaluate the term $i\{100\}$ and apply the rule **SPEC-NONSPEC**: the speculator returns that it does not want to speculate on this expression and evaluation proceeds normally (via the **SPEC- β** rule) so our branch condition is now `100 < 10`. Here the speculator consults its microarchitectural state and sees that every other time we have done this check the result has been true. The speculator thus returns `spec 1` (and a new microarchitectural state) and the rule **SPEC-SPEC** applies. **SPEC-SPEC** runs the skipped expression capturing any memory events (writes, reads, etc.) but does not commit them, instead saving them in a new speculation frame via `makeFrame`. `makeFrame` checks if the speculated value matches the real value: in this case it does not and therefore our new frame is a mispeculation frame saving the current state and continuation on our stack $\bar{\Xi}$. From here evaluation continues with the speculated value 1 and another two **SPEC-NONSPEC** steps evaluate `!b[100]`.

In these **SPEC-NONSPEC** steps the rule **SPEC- β** handles two additional aspects beyond the evaluation. Firstly it checks that the instruction is not stalled ($\neg\text{stalled}(\Phi.\Xi, \delta(\epsilon))$).⁴ This captures that fence or protection instructions will not execute speculatively (**STALL-FENCE**

³We use delayed substitutions (**β -SUBST**) to capture the fact that, when executing on hardware, argument substitution will be compiled to a register access or memory lookup and as such should not be treated as an immediate value.

⁴The $\bar{\delta}$ function removes return events and is defined in [Appendix A](#).

$$\begin{array}{c}
\boxed{\langle \Phi \mid e \rangle \xrightarrow{c} \langle \Phi \mid e \rangle} \\
\langle \Phi \mid e \rangle \xrightarrow{\bar{c}} \langle \Phi' \mid e' \rangle \\
\frac{\bar{\mu} = [\mu \mid \langle \Phi[S.p := \text{public}] \mid !z_b[z_o] \rangle] \xrightarrow{\mu} \langle \Phi'' \mid v \rangle}{\langle \Phi \mid e \rangle \xrightarrow{\bar{\mu} \circ \bar{c}} \langle \Phi' \mid e' \rangle}
\end{array}$$

Figure 8: Speculative semantics with concurrent observer

shows the rule for fences) and that, with MPK, writes and reads to protected memory will not execute speculatively (`STALL-READ` shows the rule for reads). `SPEC-β` also adds the new event to the speculative stack ($\bar{\Xi} = \text{addEvent}(\Phi, \Xi, \delta(\epsilon))$), updating whether the current speculation is invalid or not. Excerpts of `addEvent` are shown in Figure 7. `ADD-BAD-READ` checks if a read is to the location of a speculatively skipped write: if so the current speculation is invalid and will be rolled back (but continues executing until a fence). `ADD-WRITE` shows how other events are added to an in progress frame to be used to check the validity of previous speculation.

Back in our example, were we to continue evaluating we would reach the classic Spectre out of bounds read, however we will instead assume that the speculator decides to stop speculating and returns the fence directive. The rule `SPEC-FENCE` thus applies and we turn to the judgment fence $\langle \Phi \mid !b[100] \rangle$ to $\langle \Phi_1 \mid e_1 \rangle_{\bar{c}}$. This judgment, defined in Figure 7, returns the state and expression with which we will continue evaluation as well as the trace of events from evaluating the speculatively skipped expression. If the speculation was invalidated (`FENCE-ROLLBACK`) then we will return to the continuation where speculation began (this applies to our example and we return to the saved continuation if $100 < 10$ then $!b[i\{100\}]$ else 0 and the respective state at the time of speculation). If the speculation had been valid then `FENCE-COMMIT` would apply. This commits any memory events that were speculatively skipped and checks whether the new, now committed events invalidate previous speculation (we allow nested speculation thus the speculative “stack”).

3.3 Concurrent observer semantics

To capture concurrent observer capabilities we layer another semantics on top of both the speculative and non-speculative semantics. We show the new judgment for the concurrent speculative semantics in Figure 8. It consists of a singular rule that, before any step, adds a read event for every memory location visible to a concurrent thread. As MPK guarantees thread local protection this consists of every location that is not in the protected memory region, even if our “main” thread currently has access to protected memory. The non-speculative version is defined similarly.

4 ROBUST CONSTANT TIME

Operational semantics in hand we may now turn to our objects of study: cryptographic libraries and the applications that use them. Figure 9 defines the syntax, starting with *API contexts*, Γ . These map function names to the number of arguments that function takes and define the external API for a library. Given an API context Γ a *library*, L , is a set of `lib` labeled functions for each of the external names in Γ . We capture this with the well-formedness judgment

API contexts	$\Gamma ::= \bullet \mid f \mapsto z :: \Gamma$
libraries	$L ::= \bullet \mid (f, \lambda_{\text{lib}} \bar{x}. e) :: L$
secret contexts	$\Delta ::= \bullet \mid x \mapsto (z_b, z) :: \Delta$
sets of exposed blocks	$\sigma : 2^{\mathbb{Z}}$
application traces	$A ::= \tau^{\text{lib} \rightarrow \text{app}} \diamond \overline{\delta^{\text{app}}} \diamond \tau^{\text{app} \rightarrow \text{lib}}$
library traces	$L ::= \tau^{\text{app} \rightarrow \text{lib}} \diamond \overline{\delta^{\text{lib}}} \diamond \tau^{\text{lib} \rightarrow \text{app}}$
program traces	$T ::= \tau^{\text{lib} \rightarrow \text{app}} \diamond \overline{\delta^{\text{app}}} \diamond (\text{end } v)^{\text{app} \rightarrow \text{lib}} \mid A \circ L \circ T$

Figure 9: Syntax of programs and traces

$\Gamma \vDash L$, which additionally allows L to contain internal functions (defined in Figure 20 in Appendix B).

To define applications we first define *secret contexts*, Δ . We assume that secrets are stored in memory, so secret contexts capture the locations and lengths of the secret blocks as well as a variable to use to refer to that block in application code. Given an API context Γ and a secret context Δ , an *application* is then an expression with free variables from Γ and Δ and no `lib` labeled subterms, written $\Gamma, \Delta \vdash e$ and defined in Figure 20 in Appendix B.

To define programs we require another judgment, $\Delta \vDash S$, capturing that an initial program state contains all secret blocks defined by Δ . To do so we define an equality up to secrets judgment (akin to the low-equivalence definition used in work on noninterference [26]), $\Delta \vDash S = S'$ (defined in Figure 20 in Appendix B). $\Delta \vDash S = S'$ requires that both S and S' contain a block z_b corresponding to every $x \mapsto (z_b, z)$ in Δ , that the block lengths are z , and that S and S' are exactly equal everywhere outside the domain of Δ . For an API context Γ and secret context Δ we can then build a whole program from a library $\Gamma \vDash L$, an application $\Gamma, \Delta \vdash e$, and an initial state $\Delta \vDash S$ by substituting the block numbers in Δ for their names in e and the functions in L for their names in Γ , written $e[\Delta][L]$.

Our operational semantics captures a sequence of events, but for programs we factor this sequence into a *program trace* (T) with additional structure (shown in Figure 9). This trace captures the decomposition of program execution into alternating sequences of application and library domain events, with transition events the boundaries between them. Application traces A are thus a transition event from library to application, followed by a sequence of application domain events, and then a transition back to the library (we write concatenation as \diamond). Library traces are defined similarly and the trace of an entire program is then alternating sequences of these application and library traces. For program traces we define a gluing concatenation operator $Rx \circ xR'$ which matches the sequence RxR' , capturing that the transition event ending an application trace and starting the subsequent library trace are in fact the same transition event. We then define trace and speculative trace metafunctions capturing the set of all program traces for a given program (the corresponding concurrent versions are as expected):

$$\begin{array}{l}
\text{traces}(\langle S \mid e \rangle) \triangleq \{ T \mid \langle S \mid \bullet :: e \rangle \xrightarrow{T}^* \langle S' \mid \bullet :: v \rangle \} \\
\text{specTraces}(\langle \Phi \mid e \rangle) \triangleq \{ \bar{\delta} \mid \langle \Phi \mid e \rangle \xrightarrow{\bar{\delta}}^* \langle \Phi' \mid v \rangle \wedge \Phi'. \Xi = \bullet \}
\end{array}$$

4.1 Attackers

Traces capture the back and forth of actions of the application and library and the hand-offs between them. Using this we model attackers as assumptions about the behavior of an application. We instantiate this idea to five concrete attacker models: read-only attackers (corresponding to the attacker model libraries like LibSodium assume), memory-safe attackers (corresponding to applications written in memory-safe languages like Rust), speculative attackers, and speculative and non-speculative concurrent observers.

We define read-only and memory-safe attackers by restricting the set of unsafe application behavior. Read-only attackers can read memory they were not given access to but cannot write to it (and thus cannot carry out *active* attacks). In contrast, memory-safe attackers may neither read nor write memory they were not given access to. We capture these notions as trace properties, but an application is only a partial program and cannot be run. As such we must first link with a library before we can assess the application’s (mis)behavior. But we cannot simply take an arbitrary library: an ill-formed library can break application invariants. To untangle this knot we simultaneously define the relevant *restrictions* on the application we are classifying with the *assumptions* that it may make about the library that it is running against. Read-only attackers are thus defined as follows:

DEFINITION 1 (READ-ONLY ATTACKERS). *We say an application $\Gamma, \Delta \vdash e$ is a read-only attacker if, for all libraries $\Gamma \vDash L$, initial states $\Delta \vDash S$, and $T \in \text{traces}(\langle S \mid e[\Delta][L] \rangle)$, $(\text{dom}(\Delta), \emptyset) \vdash \text{read-only } T$.*

Figure 10 shows the judgment $(\sigma_A, \sigma_L) \vdash \text{read-only } T$ which handles the restrictions on the application and assumptions on the library components of the trace T . σ_A (σ_L) is the set of *exposed memory locations* for the application (library). In our model, any memory location that is passed as an argument or return value is considered as “exposed” to the party it is passed to and may be safely read and written to. The rule **READ-ONLY-REC** captures the back and forth of assumptions and restrictions: it decomposes the next application and library trace sequences, imposes the read-only restrictions on the application events ($\text{wf-read-only } \overline{\delta_A}$), and then, under the assumption that the library events do not write out of bounds, inductively requires that the rest of the trace is read-only. The definition of wf-read-only can be found in Figure 22 in Appendix B: it captures that write events are only allowed when within the set of exposed blocks.

Memory-safe attackers are defined similarly, with the predicate adjusted to also preclude reads from unexposed blocks. To capture speculative attackers we restrict to applications that are non-speculatively memory-safe, but consider speculative traces when defining our security properties below. Because our speculative semantics are parameterized by a speculation oracle, we can also consider different *classes* of speculative attacks (such as Spectre v1 vs. Spectre v2) by adding restrictions to the speculation oracle. To capture non-speculative concurrent attackers we restrict to non-speculative, read-only applications but consider the non-speculative concurrent traces.

4.2 Robust constant time

We are at last prepared to define our core security property: robust constant time. Much like classic constant time properties it comes in two flavors: speculative and non-speculative (we also separate the associated concurrent versions). We parameterize our non-speculative robust constant time property by a predicate that captures the application assumptions: in this work we will instantiate this predicate with the read-only and memory-safe properties from above. Robust constant time can thus be defined as follows:

DEFINITION 2 (ROBUST CONSTANT TIME). *We say a library $\Gamma \vDash L$ is robustly constant time for an attacker class pred if, for all secret contexts Δ , applications $\Gamma, \Delta \vDash e$ such that $\text{pred}(\Gamma, \Delta \vdash e)$, and initial states $\Delta \vDash S = S'$ we have that $\text{ct}(\text{traces}(\langle S \mid e[\Delta][L] \rangle)) = \text{ct}(\text{traces}(\langle S' \mid e[\Delta][L] \rangle))$.*

Much like classic constant time the requirement is a form of noninterference: the trace of leakage events (c in Figure 3, ct removes the extra events) cannot vary when varying the secret values (varying secret values are captured by the judgment $\Delta \vDash S = S'$). Unlike classic constant time this invariance must hold when the library is linked with *any* application satisfying our attacker model, thus making this a *robust* security property. Robust speculative constant time is defined similarly, however we consider the speculative traces and consider a speculation model (oracle) A :

DEFINITION 3 (ROBUST SPECULATIVE CONSTANT TIME). *We say a library $\Gamma \vDash L$ is robustly speculative constant time with respect to a speculation model A , if, for all secret contexts Δ , memory-safe applications $\Gamma, \Delta \vDash e$, initial states $\Delta \vDash S = S'$, $\Phi = \{S = S, A = A, \Xi = \bullet\}$, and $\Phi = \{S = S', A = A, \Xi = \bullet\}$, we have that $\text{ct}(\text{specTraces}(\langle \Phi \mid e[\Delta][L] \rangle)) = \text{ct}(\text{specTraces}(\langle \Phi' \mid e[\Delta][L] \rangle))$.*

The concurrent versions simply use the set of concurrent traces. Proof sketches are provided in Appendix C.

A secure compiler property. On its own (concurrent) robust (speculative) constant time serves as a useful security property for understanding when library protections are secure against different attackers, however we are interested in developing a compiler that automatically provides robust protections. To do so we need to understand our “source language”: what assumptions we may make about the source of our compilation. For our compiler we assume that the source is classically (speculative) constant time, thus showing that robustness can be made orthogonal to guaranteeing (speculative) constant time. This gives library developers flexibility: they can safely use any tool or handwritten technique to guarantee that their library implementation is constant time in conjunction with **ROB-COP** providing *robust* (speculative) constant time.

In defining a “source language” that captures classic constant time, note that the classic notion of constant time is that executing a library function produces invariant traces. As such classic constant time is in fact a restricted form of robust constant time, where, for a given API context Γ , the applications are defined by the grammar $e_\Gamma ::= f(\bar{v})$ for $f \in \Gamma$. That is, “source” “applications” are simply individual function calls into the library. For classical speculative constant time a slight complication arises: we must restrict the speculation models to those that do not directly record secrets from execution (captured by the extra condition $\Phi_2.A.a = \Phi'_2.A.a$):

$$(\sigma_A, \sigma_L) \vdash \text{read-only } T$$

READ-ONLY-REC

$$\frac{\sigma_A \cup \text{exposed}(\tau_1) \vdash \text{wf-read-only } \overline{\delta}_A + \sigma'_A \quad \sigma_L \cup \text{exposed}(\tau_2) \vdash \text{wf-read-only } \overline{\delta}_L + \sigma'_L \implies (\sigma'_A, \sigma'_L) \vdash \text{read-only } T}{(\sigma_A, \sigma_L) \vdash \text{read-only } (\tau_1^{\text{lib} \rightarrow \text{app}} \diamond \overline{\delta}_A^{\text{app}} \diamond \tau_2^{\text{app} \rightarrow \text{lib}}) \circ (\tau_2^{\text{app} \rightarrow \text{lib}} \diamond \overline{\delta}_L^{\text{lib}} \diamond \tau_3^{\text{lib} \rightarrow \text{app}}) \circ T}$$

Figure 10: Excerpt of read-only attacker model definition

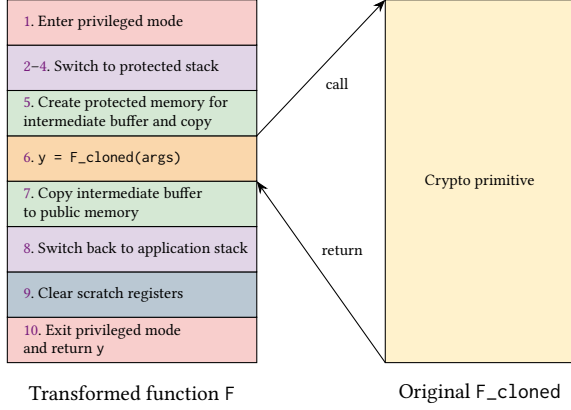


Figure 11: RoboCop wrapping of cryptographic function

DEFINITION 4 (CLASSICAL SPECULATIVE CONSTANT TIME). *We say a library $\Gamma \models L$ is classically speculative constant time with respect to a speculation oracle spec with microarchitectural state type A if, for all secret contexts Δ , classical “applications” $\Gamma, \Delta \models e_\Gamma$, initial states $\Delta \models S = S',$ microarchitectural state $a : A, \Phi_1 = \{S = S, A = \{A = A, a = a, \text{spec} = \text{spec}\}, \Xi = \bullet\},$ and $\Phi'_1 = \{S = S', A = \{A = A, a = a, \text{spec} = \text{spec}\}, \Xi = \bullet\},$ we have that for all traces $\langle \Phi_1 \mid e_\Gamma[\Delta][L] \rangle \xrightarrow{\overline{\delta}}^* \langle \Phi_2 \mid v \rangle$ there exists a trace $\langle \Phi'_1 \mid e_\Gamma[\Delta][L] \rangle \xrightarrow{\overline{\delta'}}^* \langle \Phi'_2 \mid v \rangle$ such that $\text{ct}(\overline{\delta}) = \text{ct}(\overline{\delta}'), \Phi_2.\Xi = \Phi'_2.\Xi = \bullet,$ and $\Phi_2.A.a = \Phi'_2.A.a.$*

In protecting cryptographic libraries our compiler correctness property is then secure preservation of robust constant time from the restricted, classical constant time “source” “applications” to the full set of applications for a particular attacker model. We instantiate this in Section 5.2.

5 A ROBUST COMPILER

We develop RoboCop, a compiler providing robust (speculative) constant time protections for cryptographic libraries against varying attackers. Built on top of the LLVM framework [35], RoboCop uses Intel™ Memory Protection Keys (MPK) to guarantee that secret data (cryptographic secrets and the data derived from them) is only accessible while executing trusted cryptographic library code. We discuss the design of these robust protections in Section 5.1.

While the library protections guarantee that cryptographic code executes within the protected domain, cryptographic secrets often originate and are managed by application code. As such it is necessary to allocate these secrets within protected memory. To do so we provide manual MPK allocation APIs and also adapt techniques

from CryptoMPK [29] to provide an alternative, automatic transformation that securely allocates secrets. Further, for efficiency, we reuse a stack allocated in protected memory and develop a simple LLVM pass that allocates this stack on program entry.

5.1 Making libraries robust

Cryptographic code operates directly on secret data and, to prevent timing-based leaks, is required to be constant time. With this baseline security requirement we operate under the assumption that cryptographic code is *trusted*. The task of RoboCop then is to ensure that the secret data remains inaccessible even if there are vulnerabilities within the client application code. These protections are provided in three steps: (1) Cryptographic developers label the external API functions. (2) RoboCop wraps these API functions to handle the memory isolation. (3) RoboCop replaces all dynamic memory functions (`malloc` and similar) with custom MPK compatible versions, ensuring that all memory allocated within the cryptographic library is kept within the protected domain.

Figure 11 shows our wrapping of cryptographic API functions. For every exported function F in the library, a clone, `F_cloned`, is generated containing the original implementation of F . Internal calls to F are replaced with calls to `F_cloned`: F becomes the external API wrapper for use by the client. F is responsible for switching into and out of the protected memory region.

The new F takes the following domain switching steps: (1) We enable access to the protected memory region with a `wrpkru` instruction. The specification for MPK [27] ensures this is speculatively secure: `wrpkru` will not execute speculatively and protected memory cannot be accessed until `wrpkru` is committed. (2) We get the address of the protected stack and save the current stack pointer. (3) We copy any stack arguments from the unprotected frame to the new protected stack. (4) We switch the stack pointer to the protected stack frame. (5) If concurrent protections are enabled, we allocate an internal copy buffer for the external buffers (discussed below). (6) We call `F_cloned`. (7) After the cryptographic function returns, we copy any internal buffers back to the original, public buffers. (8) We restore the previous stack pointer. (9) We clear all scratch registers which may potentially contain transient secret computation. (10) We disable access to protected memory and return to the application. Together these ensure that all data produced and used by the cryptographic code is within the protected memory region and the region is inaccessible to application code.

Concurrent protections. Rather than allocate extra memory, cryptographic algorithms sometimes carry out intermediate computations within output (often ciphertext) buffers. In a single-threaded context, this is safe as there is no way for client code to access these

buffers before they contain their final, declassified (cryptographically secure) value. In a concurrent context, work like Spectre-Declassified [53] has shown that an attacker can recover secret information if they can observe intermediate results. To defend against these attacks we add a concurrent protection option to RoboCop. Here library developers additionally annotate API arguments that are used for intermediate computation, and RoboCop allocates memory for the arguments within protected memory, performs the intermediate work within the protected domain, and then copies the declassified result back out to the unprotected memory.

5.2 Proving RoboCop secure

As discussed in Section 4.2, the relevant compiler security property is secure preservation of robust constant time between a source language of classical constant time “applications” to the target language of our attacker model. Formally, we represent RoboCop as a parameterized compiler $C_{\text{attacker}} : (\Gamma \vDash L) \rightarrow L$. In our formal setting we have that C_{spec} and $C_{\text{read-only}}$ are the same: both take a library, transform all internal uses of $\text{new}_p e$ into $\text{new}_{\text{protected}} e$, and wrap each external API function with $\text{protect}_{\text{protected}}$ and $\text{protect}_{\text{public}}$. To capture that the application manages the secret buffers and must allocate them in protected memory we slightly modify the initial state well-formedness judgment to $\Delta \vDash_{\text{protected}} S = S'$ to capture that each secret block in Δ is allocated in the protected memory region. We additionally assume that applications and libraries do not contain any protect_p expressions prior to being run through our compiler. We model $C_{\text{concurrent}}$ as $C_{\text{read-only}}$ plus a protected allocation and copy for all API arguments that are used internally. We prove that each compiler is secure and guarantees its corresponding robust constant time property. The theorem statement for $C_{\text{read-only}}$ is shown below (the remaining theorems are in Appendix C):

THEOREM 1 ($C_{\text{read-only}}$ GUARANTEES READ-ONLY ROBUST CONSTANT TIME). *If $\Gamma \vDash L$ is classically constant time and does not contain any protect_p subterms, then $C_{\text{read-only}}(\Gamma \vDash L)$ is robustly constant time for read-only attackers (that do not contain protect_p) as defined in Definition 1.*

Notably, the corresponding speculative theorem (Theorem 3 in Appendix C) ensures that C_{spec} guarantees robust speculative constant time against any class of speculation for which the library is classically speculatively constant time.

6 EVALUATION

To evaluate the cost of guaranteeing robust constant time we ask the following questions:

- Q1:** What is the overhead of *robust* constant time against read-only and speculative attackers? (§6.1)
- Q2:** What is the overhead of *robust* constant time against concurrent observers? (§6.2)

Benchmarks. To study the performance of RoboCop on a wide range of cryptographic code we modify the SUPERCOP [57] cryptographic benchmarking tool. SUPERCOP’s benchmarking suite is broken down into operations: we focus on its collections of implementations of authenticated encryption (**aead**), Diffie-Hellman (**dh**), public key encryption (**encrypt**), key encapsulation mechanism

(**kem**), public key signatures (**sign**), and stream cipher (**stream**) algorithms. Within each of these operations SUPERCOP collects multiple implementations of each algorithm: e.g. the **stream** data set contains several implementations of both Salsa20 and ChaCha20.

The design intent of SUPERCOP is to find the fastest each cryptographic algorithm can run on a given machine. To this end its benchmarking has two phases: a *try* phase and a *measure* phase. In the try phase every implementation of an algorithm is compiled with every compiler (in our case Clang with the -O3, -O2, -Os, and -O flags). These implementations are then benchmarked a single time on a small test set, with the fastest combination of implementation and compiler being selected for the measure phase. In the measure phase, these implementations are benchmarked multiple times and, if applicable, across a range of input data sizes. Due to the nature of our implementation of RoboCop we restrict one of these axes by removing all non-C/C++ implementations.

With its broad suite of algorithms and its find the fastest methodology, we believe SUPERCOP is a more *robust* means of benchmarking cryptographic software security techniques and encourage future designs to use it for benchmarking. We found its selecting from multiple compilation levels particularly beneficial as, due to the cascading effects of optimizations, comparing at the same optimization level is not truly a head-to-head comparison. Indeed we found that sometimes the fastest optimization level would differ between RoboCop and the baseline, with several instances of -O3 producing significantly slower code in combination with the protections than -O.

Machine and software setup. We run all benchmarks on a 13th Gen Intel® Core™ i9-13900KS, with 125GB RAM, and running Linux kernel version 6.3.0. We run SUPERCOP configured to collect data only on cores with the same frequency and our data is collected from CPUs running at 5.6 GHz. RoboCop adds new passes to LLVM 16.0.2 and is split into two passes: the library pass which adds the bulk of the protections and the application pass which allocates a stack in protected memory on program entry. SUPERCOP defines API functions for each operation and these are annotated as the external API for RoboCop. We manually label secret key buffers and insert protected allocations for them in the protected versions. For the concurrent protection benchmarks we label the ciphertext arguments on the **dh** and **stream** APIs as being used for internal computation. We use a modified version of jemalloc 5.2.0 [29] patched to provide MPK allocation functions. Our baseline replaces the libc malloc implementation with an unpatched version of jemalloc.

Summary of results. We find that robust constant time protections can be generally be guaranteed with minimal overhead (less than 5% in almost all cases), though there do exist a small number of outliers with up to 20% overhead for key encapsulation mechanisms and up to 40% overhead for AEAD and signature algorithms. At small data sizes highly optimized stream ciphers also carry a large overhead, however these workloads take on the order of a few hundred cycles. We also find that there is a higher, but still minimal overhead for protecting against concurrent attackers.

Benchmark	N	Data size (bytes)	Q_1	Median overhead	Q_3
aead	385	2048	-0.17%	0.17%	0.73%
dh	9	fixed	-0.19%	0.23%	0.51%
encrypt	15	4237	-1.89%	-0.02%	0.80%
kem	84	fixed	-0.62%	0.05%	1.99%
sign	38	4237	-0.27%	0.02%	0.92%
stream	11	4096	0.86%	0.96%	1.32%

Table 1: Overheads for read-only/speculative protections vs. unprotected baseline. N is the number of algorithms in the dataset, Data size is the size of the operation’s input, Q_1 and Q_3 are the first and third quartile overheads, and the median overhead is of the overheads of the mean runtimes.




Figure 12: Box plot of overheads for read-only/speculative protections compared to unprotected baseline. Cutoff at 10% overhead, outliers beyond this point are discussed directly.⁵

6.1 Read-only and speculative attackers

We measure the cost of ensuring robust constant time against read-only and speculative attackers across six data sets (shown in Table 1 and Figure 12). Assuming the library is classically speculatively constant time, the only additional protections needed to move from robust constant time against read-only attackers to speculative robust constant time are the clearing of scratch registers when exiting cryptographic code. As such we combine the measurements for both protections into one and treat the extra clearing of registers as defense in depth in the case of read-only attackers.

For the largest data size for each benchmark we find that the median overhead across all benchmarked algorithms is below 1% and that 75% of all implementations for each primitive have overheads under 2%. For algorithms with varying input lengths, SUPERCOP measures performance across a wide range of lengths. We show the varying overheads across these sizes for stream ciphers in Table 2. The overhead increases as data sizes get smaller, with a median overhead of 34% for encrypting a single byte. Fortunately, at this data size encryption only takes a few hundred cycles so this high relative overhead remains a minimal raw cost.

Outliers. While for 75% of implementations in our data set, robust protections have overheads below 2%, there are some notable outliers above 10%.⁶ Firstly, two **kem** implementations have reported overheads around 400%. This is due to the two-stage nature of SUPERCOP: the initial measurement in the try-phase incorrectly identifies the fastest optimization level when using RoboCOP, choosing -0 as the optimization level for RoboCOP and -03 for the baseline. When we run the full measurements for RoboCOP with -03 on these two **kem** implementations, the overhead drops to sub 2%.

⁵For clarity: 30/385 **aead** and 11/84 **kem** implementations lie above the respective whiskers and below the cutoff.

⁶Figure 12 shows all outliers below 10%.

Data size (bytes)	Q_1	Median overhead	Q_3	Baseline cycles
1	18.46%	34.08%	52.29%	4.45e+02
128	11.37%	15.68%	23.23%	8.10e+02
256	8.49%	9.23%	16.25%	1.41e+03
512	4.78%	6.00%	8.45%	2.12e+03
1024	3.16%	3.82%	5.64%	4.08e+03
2048	1.63%	2.46%	4.03%	7.72e+03
4096	0.86%	0.96%	1.32%	1.54e+04

Table 2: Read-only/speculative protections overheads for stream ciphers with varying plaintext sizes. Baseline cycles is the mean number of cycles for the unprotected baseline.

There are also three **aead** algorithms whose measured overheads exceed 20% and whose selected optimization levels differ, however unlike the **kem** implementations here SUPERCOP’s selection of the fastest implementation parameters is correct: the overhead more than doubles when -03 is used with RoboCOP. This illustrates a significant benefit of using SUPERCOP compared to benchmarking with the same optimizations applied in all cases.

After reanalysis there is one **kem** implementation with an overhead of 17%, two **sign** implementations with overheads of 21% and 36%, and four **aead** implementations with overheads of 28%, 36%, 36%, and 37%. We do not have explanations for these higher overheads, however we observe that they call OpenSSL’s cryptographic implementations and use internal randomness and dynamic allocation, though they are not the only algorithms that do so.

6.2 Concurrent attackers

To protect against concurrent attackers it is necessary for all buffers containing intermediate values derived from secrets to remain within the MPK protected memory. In its read-only attacker protections mode RoboCOP ensures all memory originating from cryptographic code meets this requirement, however some cryptographic implementations use external, public buffers as internal, intermediate (private) buffers. To measure the cost of protecting these intermediate buffers we benchmark the **dh** and **stream** data sets with RoboCOP’s concurrent protections mode. Here we annotate the top-level SUPERCOP API functions `crypto_dh` and `crypto_stream` to mark the ciphertext argument as being used for intermediate computation. In the case of **stream** the size of this buffer is dynamically determined so we mark the plaintext length argument as the size for allocating a secure intermediate buffer. Table 3 shows the median overheads for the read-only protections compared to the median overheads for the concurrent protections. For our **dh** data set protecting these intermediate buffers increases the median overhead from 0.23% to 0.46%. For our stream cipher data set at the largest data size the increase is greater, with the median increasing from 0.96% to 1.74%. We hypothesize that the higher overhead for stream ciphers is due to the dynamic allocation whereas the statically sized buffer of the Diffie-Hellman API allows optimizations in both allocation and copying back to the external buffer.

Our benchmarking treats every algorithm within each data set as if they use public buffers for intermediate computations. In practice RoboCOP lets library designers label specifically which/if buffers are used for intermediate computations. This ensures that code that does not use the external buffers never has to pay the price for the extra allocation and copying and that the same library code base

NON-CONCURRENT			CONCURRENT		
Q_1	median	Q_3	Q_1	median	Q_3
-0.19%	0.23%	0.51%	0.33%	0.46%	0.95%

(a) *dh*

Data size (bytes)	NON-CONCURRENT			CONCURRENT		
	Q_1	median	Q_3	Q_1	median	Q_3
1	18.46%	34.08%	52.29%	24.99%	42.05%	72.92%
128	11.37%	15.68%	23.23%	18.0%	20.02%	34.45%
256	8.49%	9.23%	16.25%	10.97%	13.12%	22.70%
512	4.78%	6.00%	8.45%	6.68%	9.11%	11.84%
1024	3.16%	3.82%	5.64%	4.41%	5.63%	8.04%
2048	1.63%	2.46%	4.03%	2.61%	3.22%	4.15%
4096	0.86%	0.96%	1.32%	1.37%	1.74%	2.55%

(b) *stream***Table 3: Overhead of read-only/speculative protections vs. overhead with concurrent protections.**

can be used in all contexts: in single-threaded mode ROBOCOP can omit the allocation and copy but in concurrent mode it handles the creation of a protected intermediate buffer.

7 RELATED WORK

Memory isolation. MPK has been used to provide in-process memory isolation [25, 56], including between safe and unsafe Rust code [23, 31, 48]. MPK protections can be modified by unprivileged instructions, as such ROBOCOP assumes that applications do not have access to these instructions and including through control-flow hijacks. Countermeasures like binary rewriting [56], system call filtering [51, 59], and CFI schemes [10, 11] could be used to lift these assumptions.

Similar to ROBOCOP, CryptoMPK [29] leverages MPK to protect the confidentiality of secret cryptographic data from memory disclosure vulnerabilities. We believe that it can guarantee robust constant time against read-only attackers, though their work is not framed in this manner. CryptoMPK differs from ROBOCOP in several fundamental ways: First, it is instead a whole-program analysis and transformation, identifying “crypto buffers” throughout the program and toggling MPK protection around use sites. As such it inserts significantly more context switches than ROBOCOP, and dynamically allocates and frees secret stack buffers. ROBOCOP’s trusted library model avoids these costs, allowing a single context switch on library entry and exit, completely avoiding the need to dynamically allocate stack buffers and leading to significant performance gains. Second, by avoiding a whole program analysis ROBOCOP’s model allows a much simpler implementation. This allows ROBOCOP to be applied to more complicated code and even hard to analyze assembly code (though we have not implemented this). Lastly, CryptoMPK does not handle robust speculative protections nor robust protections against concurrent attackers. As an optimization, CryptoMPK chooses not to securely allocate small secret stack buffers and instead inserts zeroing code. This zeroing code has the same trade offs between protecting against Spectre and performance as in LibSodium, and the buffers are also visible to concurrent attackers. CryptoMPK provides a `mxor` annotation to ensure that eventually declassified buffers are not marked as tainted when they are mixed with secret key data. This has the result of exposing these buffers to concurrent attackers.

Secure zeroization is often deployed as a manual protection in cryptographic libraries. Unfortunately, implementing secure memory zeroing in a high-level language is essentially impossible [46, 49, 62]. Recent work [42] shows how to develop a compiler pass to implement secure zeroization. ROBOCOP avoids these issues by restricting all secret data to a protected memory region, thus avoiding the need for zeroization (apart from register zeroing).

(Speculative) constant time. Many domain-specific languages and compilers have been developed to produce high-assurance cryptographic code [3, 6, 15, 47, 60]. Spectre attacks [32] significantly reduced the guarantees of these tools and prompted defenses against speculative leaks [14]. Jasmin implements Selective Speculative Load Hardening to protect against Spectre-PHT [53], performing stack zeroization and register clearing [42]. Blade inserts a minimum number of protections to prevent leaks via Spectre-PHT gadgets [58]. Swivel hardens WebAssembly sandboxes against speculative sandbox breakout and sandbox poisoning attacks [41, 63]. Serberus mitigates all currently known Spectre variants in code that follows the *static* constant-time discipline, which additionally prohibits secret function arguments and return values [40]. These tools serve as complements to ROBOCOP: in combination they can be used to guarantee end-to-end protections against speculative attackers as discussed in §4.2. Notably, there is an exception to this statement in the case of Serberus: In handling Spectre-RSB [33], Serberus makes an implicit assumption that the return stack buffer is empty when entering cryptographic code. Much like the issues with prior work on constant time protections, Serberus is assuming that the cryptographic code represents the entire program. This can be remedied by RSB filling on entry to cryptographic code.

Foundations for cryptographic software security. Researchers have developed trace-based leakage models to reason about timing leaks in cryptographic code [7, 39]. These models have then been extended with prediction oracles [22], speculative semantics and directives [13, 16, 21] to capture leaks via (combinations of) different Spectre gadgets [18]. Our formal approach builds on these well-established practices. Our notion of RCT is inspired by previous work on secure compilers [2], which are formally defined as compilers that preserve classes of (hyper)-properties in adversarial contexts. Patrignani and Guarnieri [45] develop secure robust compilation criteria to formally examine the security guarantees of protections inserted by major compilers against Spectre-PHT, our approach to a secure compiler property follows this line of work.

8 LIMITATIONS

Our implementation of ROBOCOP has a few limitations: (1) While ROBOCOP handles attackers in concurrent threads it does not itself handle running concurrent *cryptographic library* code. To handle this we could allocate a single protected stack per thread. It would be safe to reuse a single MPK protection key across all threads as concurrent threads would only have access while running cryptographic library code. (2) ROBOCOP assumes that no other code is using MPK. (3) ROBOCOP does not currently handle ensuring that protected memory does not get dumped on crashes or written to disk nor does it ensure that it is cleared at the end of program execution. These could be handled in one place by operating on the pages assigned to the MPK protected domain. (4) ROBOCOP assumes

the speculative behavior of MPK follows Intel’s specification and `wrpkru` will never execute speculatively and protected memory will never be accessed speculatively until protections have been fully committed [27]. Hardware bugs such as `melt-down-pk` [12] violate these assumptions, and we rely on hardware fixes for such bugs (for instance official patches have already been provided for the machine used for our evaluation). (5) `ROBOCOP` does not currently toggle Intel™ `DOITM` [28]. Currently this mode has no effect, but it will become necessary to ensure certain instructions run in constant time. As it is not intended for always on use [34], `ROBOCOP` is ideal for managing its use for cryptographic library code.

REFERENCES

- [1] Martin Abadi. 1999. Secrecy by Typing in Security Protocols. *J. ACM* 46, 5 (Sept. 1999), 749–786. <https://doi.org/10.1145/324133.324266>
- [2] Carmine Abate, Roberto Blanco, Deepak Garg, Catalin Hritcu, Marco Patrignani, and Jeremy Thibault. 2019. Journey Beyond Full Abstraction: Exploring Robust Property Preservation for Secure Compilation. In *2019 IEEE 32nd Computer Security Foundations Symposium (CSF)* (Hoboken, NJ, USA, 2019-06). IEEE, 256–25615. <https://doi.org/10.1109/CSF.2019.00025>
- [3] José Bacelar Almeida, Manuel Barbosa, Gilles Barthe, Arthur Blot, Benjamin Grégoire, Vincent Laporte, Tiago Oliveira, Hugo Pacheco, Benedikt Schmidt, and Pierre-Yves Strub. 2017. Jasmin: High-Assurance and High-Speed Cryptography. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security* (Dallas, Texas, USA) (CCS ’17). Association for Computing Machinery, New York, NY, USA, 1807–1823. <https://doi.org/10.1145/3133956.3134078>
- [4] OpenSSL Project Authors. [n. d.]. *OpenSSL*. OpenSSL Project. <https://www.openssl.org/>
- [5] Manuel Barbosa, Gilles Barthe, Karthik Bhargavan, Bruno Blanchet, Cas Cremers, Kevin Liao, and Bryan Parno. 2021. SoK: Computer-Aided Cryptography. In *2021 IEEE Symposium on Security and Privacy (SP)* (2021-05). IEEE, 777–795. <https://doi.org/10.1109/SP40001.2021.00008> ISSN: 2375-1207.
- [6] Gilles Barthe, Sandrine Blazy, Benjamin Grégoire, Rémi Hutin, Vincent Laporte, David Pichardie, and Alix Trieu. 2019. Formal Verification of a Constant-Time Preserving C Compiler. *Proc. ACM Program. Lang.* 4, POPL, Article 7 (dec 2019), 30 pages. <https://doi.org/10.1145/3371075>
- [7] Gilles Barthe, Benjamin Grégoire, and Vincent Laporte. 2018. Secure Compilation of Side-Channel Countermeasures: The Case of Cryptographic “Constant-Time”. In *2018 IEEE 31st Computer Security Foundations Symposium (CSF)* (2018-07). IEEE, 328–343. <https://doi.org/10.1109/CSF.2018.00031> ISSN: 2374-8303.
- [8] Daniel J. Bernstein. 2005. *Cache-timing attacks on AES*. Technical Report. The University of Illinois at Chicago. <https://api.semanticscholar.org/CorpusID:2217245>
- [9] Daniel J. Bernstein. 2008. The Salsa20 Family of Stream Ciphers. In *New Stream Cipher Designs: The eSTREAM Finalists* (Berlin, Heidelberg) (*Lecture Notes in Computer Science*), Matthew Robshaw and Olivier Billet (Eds.). Springer, 84–97. https://doi.org/10.1007/978-3-540-68351-3_8
- [10] Nathan Burow, Scott A. Carr, Joseph Nash, Per Larsen, Michael Franz, Stefan Brunthaler, and Mathias Payer. 2017. Control-Flow Integrity: Precision, Security, and Performance. *Comput. Surveys* 50 (April 2017), 16:1–16:33. <https://doi.org/10.1145/3054924>
- [11] Nathan Burow, Xiping Zhang, and Mathias Payer. [n. d.]. SoK: Shining Light on Shadow Stacks. In *2019 IEEE Symposium on Security and Privacy (SP)* (2019-05). IEEE, 985–999. <https://doi.org/10.1109/SP.2019.00076> ISSN: 2375-1207.
- [12] Claudio Canella, Jo Van Bulck, Michael Schwarz, Moritz Lipp, Benjamin von Berg, Philipp Ortner, Frank Piessens, Dmitry Evtvyushkin, and Daniel Gruss. [n. d.]. A Systematic Evaluation of Transient Execution Attacks and Defenses. 249–266. <https://www.usenix.org/conference/usenixsecurity19/presentation/canella>
- [13] Sunjay Cauligi, Craig Disselkoben, Klaus v. Gleissenthall, Dean Tullsen, Deian Stefan, Tamara Rezk, and Gilles Barthe. [n. d.]. Constant-time foundations for the new spectre era. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation* (New York, NY, USA, 2020-06-11) (PLDI 2020). Association for Computing Machinery, 913–926. <https://doi.org/10.1145/3385412.3385970>
- [14] Sunjay Cauligi, Craig Disselkoben, Daniel Moghimi, Gilles Barthe, and Deian Stefan. 2022. SoK: Practical Foundations for Software Spectre Defenses. In *2022 IEEE Symposium on Security and Privacy (SP)* (2022-05). IEEE, 666–680. <https://doi.org/10.1109/SP46214.2022.9833707> ISSN: 2375-1207.
- [15] Sunjay Cauligi, Gary Soeller, Brian Johannesmeyer, Fraser Brown, Riad S. Wahby, John Renner, Benjamin Gregoire, Gilles Barthe, Ranjit Jhala, and Deian Stefan. 2019. FaCT: A DSL for timing-sensitive computation. In *Programming Language Design and Implementation (PLDI)*. ACM SIGPLAN.
- [16] Kevin Cheang, Cameron Rasmussen, Sanjit Seshia, and Pramod Subramanyan. 2019. A Formal Approach to Secure Speculation. In *2019 IEEE 32nd Computer Security Foundations Symposium (CSF)*. IEEE, 288–28815. <https://doi.org/10.1109/CSF.2019.00027>
- [17] Frank Denis. [n. d.]. *libsodium*. libsodium. <https://doc.libsodium.org/>
- [18] Xavier Fabian, Marco Guarnieri, and Marco Patrignani. 2022. Automatic Detection of Speculative Execution Combinations. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security* (Los Angeles, CA, USA) (CCS ’22). Association for Computing Machinery, New York, NY, USA, 965–978. <https://doi.org/10.1145/3548606.3560555>
- [19] Cédric Fournet, Andrew D. Gordon, and Sergio Maffei. 2007. A Type Discipline for Authorization Policies. *ACM Trans. Program. Lang. Syst.* 29, 5, Article 25 (Aug. 2007). <https://doi.org/10.1145/1275497.1275500>
- [20] Andrew D. Gordon and Alan Jeffrey. 2003. Authenticity by Typing for Security Protocols. *J. Comput. Secur.* 11, 4 (July 2003), 451–519. <http://dl.acm.org/citation.cfm?id=959088.959090>
- [21] Marco Guarnieri, Boris Köpf, Jan Reineke, and Pepe Vila. 2021. Hardware-Software Contracts for Secure Speculation. In *42nd IEEE Symposium on Security and Privacy, SP 2021, San Francisco, CA, USA, 24-27 May 2021*. IEEE, 1868–1883. <https://doi.org/10.1109/SP40001.2021.00036>
- [22] Marco Guarnieri, Boris Köpf, José F. Morales, Jan Reineke, and Andrés Sánchez. 2020. Spectector: Principled Detection of Speculative Information Flows. In *2020 IEEE Symposium on Security and Privacy (SP)* (2020-05). IEEE, 1–19. <https://doi.org/10.1109/SP40000.2020.00011> ISSN: 2375-1207.
- [23] Merve Gülmez, Thomas Nyman, Christoph Baumann, and Jan Tobias Mühlberg. 2023. Friend or Foe Inside? Exploring In-Process Isolation to Maintain Memory Safety for Unsafe Rust. *arXiv preprint arXiv:2306.08127* (2023).
- [24] Enes Göktas, Kaveh Razavi, Georgios Portokalidis, Herbert Bos, and Cristiano Giuffrida. 2020. Speculative Probing: Hacking Blind in the Spectre Era. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security* (New York, NY, USA, 2020-11-02) (CCS ’20). Association for Computing Machinery, 1871–1885. <https://doi.org/10.1145/3372297.3417289>
- [25] Mohammad Hedayati, Spyridoula Gravani, Ethan Johnson, John Criswell, Michael L Scott, Kai Shen, and Mike Marty. 2019. Hodor: Intra-process isolation for high-throughput data plane libraries. In *2019 USENIX Annual Technical Conference, USENIX ATC 2019, Renton, WA, USA, July 10-12, 2019*. USENIX Association.
- [26] Daniel Hedin and A. Sabelfeld. [n. d.]. A Perspective on Information-Flow Control.
- [27] Intel 2020. Intel® 64 and IA-32 Architectures Software Developer’s Manual.
- [28] Intel. 2023. *Data Operand Independent Timing ISA Guidance*. Technical Report. <https://www.intel.com/content/www/us/en/developer/articles/technical/software-security-guidance/best-practices/data-operand-independent-timing-isa-guidance.html>
- [29] Xuancheng Jin, Xuangan Xiao, Songlin Jia, Wang Gao, Dawu Gu, Hang Zhang, Siqi Ma, Zhiyun Qian, and Juanru Li. 2022. Annotating, tracking, and protecting cryptographic secrets with CryptoMPK. In *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, IEEE, 650–665.
- [30] Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer. 2017. RustBelt: securing the foundations of the rust programming language. 2 (2017), 1–34. Issue POPL. <https://doi.org/10.1145/3158154>
- [31] Paul Kirth, Mitchel Dickerson, Stephen Crane, Per Larsen, Adrian Dabrowski, David Gens, Yeoul Na, Stijn Volckaert, and Michael Franz. 2022. PKRU-safe: automatically locking down the heap between safe and unsafe languages. In *Proceedings of the Seventeenth European Conference on Computer Systems*. 132–148.
- [32] 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)* (2019-05). IEEE, 1–19. <https://doi.org/10.1109/SP.2019.00002> ISSN: 2375-1207.
- [33] Esmail Mohammadian Koruyeh, Khaled N. Khasawneh, Chengyu Song, and Nael Abu-Ghazaleh. 2018. Spectre Returns! Speculation Attacks using the Return Stack Buffer. <https://www.usenix.org/conference/woot18/presentation/koruyeh>
- [34] Michael Larabel. 2023. Intel’s “DOITM” Security Feature Not Intended For Always-On Use, Linux Patches To Be Revised. <https://www.phoronix.com/news/Intel-DOITM-Not-Always-On>
- [35] Chris Lattner and Vikram Adve. 2004. LLVM: A Compilation Framework for Lifelong Program Analysis and Transformation. In *CGO*. San Jose, CA, USA, 75–88.
- [36] Xavier Leroy. 2009. A Formally Verified Compiler Back-end. 43, 4 (2009), 363–446. <https://doi.org/10.1007/s10817-009-9155-4>
- [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* (New York, NY, USA, 2018-10-15) (CCS ’18). Association for Computing Machinery, 2109–2122. <https://doi.org/10.1145/3243734.3243761>

- [38] Andrea Mambretti, Alexandra Sandulescu, Alessandro Sorniotti, William Robertson, Engin Kirda, and Anil Kurmus. 2021. Bypassing memory safety mechanisms through speculative control flow hijacks. In *2021 IEEE European Symposium on Security and Privacy (EuroS&P)* (2021-09). IEEE, 633–649. <https://doi.org/10.1109/EuroSP51992.2021.00048>
- [39] David Molnar, Matt Piotrowski, David Schultz, and David Wagner. 2005. The Program Counter Security Model: Automatic Detection and Removal of Control-Flow Side Channel Attacks. USENIX Association, Baltimore, MD.
- [40] Nicholas Mosier, Hamed Nemati, John C. Mitchell, and Caroline Trippel. 2023. Serberus: Protecting Cryptographic Code from Spectres at Compile-Time. *ArXiv abs/2309.05174* (2023). <https://api.semanticscholar.org/CorpusID:261682113>
- [41] Shravan Narayan, Craig Disselkoe, Daniel Moghimi, Sunjay Cauligi, Evan Johnson, Zhao Gang, Anjo Vahldiek-Oberwagner, Ravi Sahita, Hovav Shacham, Dean Tullsen, and Deian Stefan. 2021. Swivel: Hardening WebAssembly against Spectre. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 1433–1450. <https://www.usenix.org/conference/usenixsecurity21/presentation/narayan>
- [42] Santiago Arranz Olmos, Gilles Barthe, Ruben Gonzalez, Benjamin Grégoire, Vincent Laporte, Jean-Christophe Lechenet, Tiago Oliveira, and Peter Schwabe. 2023. High-assurance zeroization. *Cryptology ePrint Archive, Paper 2023/1713*. <https://eprint.iacr.org/2023/1713> <https://eprint.iacr.org/2023/1713>
- [43] Marco Patrignani, Amal Ahmed, and Dave Clarke. 2019. Formal Approaches to Secure Compilation: A Survey of Fully Abstract Compilation and Related Work. *ACM Comput. Surv.* 51, 6, Article 125 (feb 2019), 36 pages. <https://doi.org/10.1145/3280984>
- [44] Marco Patrignani and Sam Blackshear. 2023. Robust Safety for Move. In *2023 IEEE 36th Computer Security Foundations Symposium (CSF)*. IEEE Computer Society, Los Alamitos, CA, USA, 308–323. <https://doi.org/10.1109/CSF57540.2023.00045>
- [45] Marco Patrignani and Marco Guarnieri. 2021. Exorcising Spectres with Secure Compilers. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer Communications Security (Virtual Event, Republic of Korea) (CCS '21)*. Association for Computing Machinery, New York, NY, USA, 445–461. <https://doi.org/10.1145/3460120.3484534>
- [46] Colin Percival. 2014. *Zeroing buffers is insufficient*. <https://www.daemonology.net/blog/2014-09-06-zeroing-buffers-is-insufficient.html>
- [47] Jonathan Protzenko, Jean-Karim Zinzindohoué, Aseem Rastogi, Tahina Ramananandro, Peng Wang, Santiago Zanella-Béguelin, Antoine Delignat-Lavaud, Cătălin Hrițcu, Karthikeyan Bhargavan, Cédric Fournet, and Nikhil Swamy. 2017. Verified Low-Level Programming Embedded in F*. *Proc. ACM Program. Lang.* 1, ICFP, Article 17 (aug 2017), 29 pages. <https://doi.org/10.1145/3110261>
- [48] Elijah Rivera, Samuel Mergendahl, Howard Shrobe, Hamed Okhravi, and Nathan Burow. 2021. Keeping safe rust safe with galeed. In *Annual Computer Security Applications Conference*. 824–836.
- [49] RustCrypto. [n. d.]. *Zeroize*. RustCrypto. <https://docs.rs/zeroize/1.7.0/zeroize/>
- [50] Michael Sammler, Deepak Garg, Derek Dreyer, and Tadeusz Litak. 2020. The high-level benefits of low-level sandboxing. *PACMPL* 4, POPL (2020), 32:1–32:32. <https://doi.org/10.1145/3371100>
- [51] David Schrammel, Samuel Weiser, Richard Sadek, and Stefan Mangard. 2022. Jenny: Securing Syscalls for {PKU-based} Memory Isolation Systems. In *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, 936–952.
- [52] Basavesh Ammanaghatta Shivakumar, Jack Barnes, Gilles Barthe, Sunjay Cauligi, Chitchanok Chuengsatiansup, Daniel Genkin, Sioli O’Connell, Peter Schwabe, Rui Qi Sim, and Yuval Yarom. 2023. Spectre Declassified: Reading from the Right Place at the Wrong Time. In *2023 IEEE Symposium on Security and Privacy (SP) (2023-05)*. 1753–1770. <https://doi.org/10.1109/SP46215.2023.10179355> ISSN: 2375-1207.
- [53] Basavesh Ammanaghatta Shivakumar, Jack Barnes, Gilles Barthe, Sunjay Cauligi, Chitchanok Chuengsatiansup, Daniel Genkin, Sioli O’Connell, Peter Schwabe, Rui Qi Sim, and Yuval Yarom. 2023. Spectre declassified: Reading from the right place at the wrong time. In *2023 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1753–1770.
- [54] David Swasey, Deepak Garg, and Derek Dreyer. 2017. Robust and Compositional Verification of Object Capability Patterns. In *Proceedings of the 2017 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2017, October 22 - 27, 2017*.
- [55] Reini Urban. [n. d.]. *libsodium_memzero with memory barrier*. Issue #802 · *jedisct1/libsodium*. <https://github.com/jedisct1/libsodium/issues/802>
- [56] Anjo Vahldiek-Oberwagner, Eslam Elnikety, Nuno O Duarte, Michael Sammler, Peter Druschel, and Deepak Garg. 2019. ERIM: Secure, Efficient In-process Isolation with Protection Keys (MPK). In *28th USENIX Security Symposium (USENIX Security 19)*. USENIX Association, 1221–1238.
- [57] VAMPIRE. [n. d.]. *SUPERCOP: System for Unified Performance Evaluation Related to Cryptographic Operations and Primitives*. <https://bench.cr.yt.to/supercop.html>
- [58] Marco Vassena, Craig Disselkoe, Klaus von Gleissenthall, Sunjay Cauligi, Rami Gökhan Kıcı, Ranjit Jhala, Dean Tullsen, and Deian Stefan. 2021. Automatically eliminating speculative leaks from cryptographic code with blade. 5 (2021), 49:1–49:30. Issue POPL. <https://doi.org/10.1145/3434330>
- [59] Alexios Voulimeas, Jonas Vinck, Ruben Mechelinc, and Stijn Volckaert. 2022. You shall not (by) pass! practical, secure, and fast PKU-based sandboxing. In *Proceedings of the Seventeenth European Conference on Computer Systems*. 266–282.
- [60] Conrad Watt, John Renner, Natalie Popescu, Sunjay Cauligi, and Deian Stefan. 2019. CT-Wasm: Type-Driven Secure Cryptography for the Web Ecosystem. *Proc. ACM Program. Lang.* 3, POPL, Article 77 (jan 2019), 29 pages. <https://doi.org/10.1145/3290390>
- [61] Johannes Wikner and Kaveh Razavi. 2022. RETBLEED: Arbitrary Speculative Code Execution with Return Instructions. 3825–3842. <https://www.usenix.org/conference/usenixsecurity22/presentation/wikner>
- [62] Zhaomo Yang, Brian Johannesmeyer, Anders Trier Olesen, Sorin Lerner, and Kirill Levchenko. 2017. Dead Store Elimination (Still) Considered Harmful. 1025–1040. <https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/yang>
- [63] Hosein Yavarzadeh, Mohammadkazem Taram, Shravan Narayan, Deian Stefan, and Dean Tullsen. 2023. Half&Half: Demystifying Intel’s Directional Branch Predictors for Fast, Secure Partitioned Execution. In *IEEE Symposium on Security and Privacy (S&P)*. IEEE.

A LANGUAGE

$$\langle S \mid \overline{K}^\ell :: e^\ell \rangle \xRightarrow{\epsilon} \langle S \mid \overline{K}^\ell :: e^\ell \rangle$$

$$\text{RED-CALL} \quad \epsilon = \begin{cases} \text{call } (\lambda \bar{x}. e)(\bar{v})^\ell & \text{when } \ell_f = \ell \\ (\text{call } (\lambda \bar{x}. e)(\bar{v}))_{\ell \rightarrow \ell_f} & \text{otherwise} \end{cases}$$

$$\frac{}{\langle S \mid \overline{K}^{\ell'} :: K[(\lambda \bar{x}. e)(\bar{v}))^\ell] \rangle \xRightarrow{\epsilon} \langle S \mid \overline{K}^{\ell'} :: K^\ell :: e[v/x]^\ell \rangle}$$

$$\text{RED-RET} \quad \epsilon = \begin{cases} 0^\ell & \text{when } \ell_K = \ell \\ (\text{ret } v)_{\ell \rightarrow \ell_K} & \text{otherwise} \end{cases}$$

$$\frac{}{\langle S \mid \overline{K}^{\ell'} :: K^{\ell_K} :: v^\ell \rangle \xRightarrow{\epsilon} \langle S \mid \overline{K}^{\ell'} :: K[v]^\ell \rangle}$$

$$\text{RED-}\beta$$

$$\frac{\langle S \mid e \rangle \xrightarrow{\delta} \langle S' \mid e' \rangle}{\langle S \mid \overline{K}^{\ell'} :: K[e]^\ell \rangle \xRightarrow{\delta^\ell} \langle S' \mid \overline{K}^{\ell'} :: K[e']^\ell \rangle}$$

$$\langle S \mid e \rangle \xrightarrow{\delta} \langle S \mid e \rangle$$

$$\beta\text{-SUBST} \quad \frac{}{\langle S \mid x\{v\} \rangle \xrightarrow{0} \langle S \mid v \rangle}$$

$$\beta\text{-OP} \quad \frac{v' = \delta(\text{op})(\bar{v})}{\langle S \mid \text{op}(\bar{v}) \rangle \xrightarrow{0} \langle S \mid v' \rangle}$$

$$\beta\text{-DEREF} \quad \frac{\text{accessible}(S, z_b) \quad z_o \in [S(z_b).size] \quad v = S(z_b).v(z_o)}{\langle S \mid !(z_b[z_o]) \rangle \xrightarrow{\text{read } v \leftarrow z_b[z_o]} \langle S \mid v \rangle}$$

$$\beta\text{-DEREF-OOB} \quad \frac{z_b \notin \text{dom}(S) \vee z_o \notin [S(z_b).size] \quad z'_b \in \text{dom}(S) \quad z'_o \in [S(z'_b).size] \quad v = S(z'_b).v(z'_o) \quad \text{accessible}(S, z'_b)}{\langle S \mid !(z_b[z_o]) \rangle \xrightarrow{\text{read } v \leftarrow z'_b[z'_o]} \langle S \mid v \rangle}$$

$$\beta\text{-WRITE} \quad \frac{\text{accessible}(S, z_b) \quad z_o \in [S(z_b).size] \quad S' = S(z_b).v[z_o := v]}{\langle S \mid z_b[z_o] := v \rangle \xrightarrow{\text{write } v \rightarrow z_b[z_o]} \langle S' \mid 0 \rangle}$$

$$\beta\text{-WRITE-OOB} \quad \frac{z_b \notin \text{dom}(S) \vee z_o \notin [S(z_b).size] \quad z'_b \in \text{dom}(S) \quad z'_o \in [S(z'_b).size] \quad S' = S(z'_b).v[z'_o := v] \quad \text{accessible}(S, z'_b)}{\langle S \mid z_b[z_o] := v \rangle \xrightarrow{\text{write } v \rightarrow z'_b[z'_o]} \langle S' \mid 0 \rangle}$$

$$\beta\text{-NEW} \quad \frac{z > 0 \quad z_b = \text{fresh}(S) \quad S.p \sqsubseteq p \quad S' = S[z_b := \{size = z, v = \perp, p = p\}]}{\langle S \mid \text{new}_p z \rangle \xrightarrow{\text{new}_p z @ z_b} \langle S' \mid z_b[0] \rangle}$$

$$\beta\text{-GET-BLOCK} \quad \frac{}{\langle S \mid \text{get-block } (z_b[z_o]) \rangle \xrightarrow{0} \langle S \mid z_b \rangle}$$

$$\beta\text{-GET-OFFSET} \quad \frac{}{\langle S \mid \text{get-offset } (z_b[z_o]) \rangle \xrightarrow{0} \langle S \mid z_o \rangle}$$

$$\beta\text{-IF-FALSE} \quad \frac{}{\langle S \mid \text{if } 0 \text{ then } e \text{ else } e' \rangle \xrightarrow{\text{branch } 0} \langle S \mid e' \rangle}$$

$$\beta\text{-IF-TRUE} \quad \frac{v \neq 0}{\langle S \mid \text{if } v \text{ then } e \text{ else } e' \rangle \xrightarrow{\text{branch } v} \langle S \mid e \rangle}$$

$$\beta\text{-PROTECT} \quad \frac{}{\langle S \mid \text{protect}_p \rangle \xrightarrow{\text{protect}_p} \langle S[p := p] \mid 0 \rangle}$$

$$\beta\text{-SEQ} \quad \frac{}{\langle S \mid v; e \rangle \xrightarrow{0} \langle S \mid e \rangle}$$

$$\beta\text{-FENCE} \quad \frac{}{\langle S \mid \text{fence} \rangle \xrightarrow{\text{fence}} \langle S \mid 0 \rangle}$$

Figure 13: Non-speculative trace semantics

$$\begin{array}{l}
x[v/x] \triangleq x\{v\} \\
y[v/x] \triangleq y \\
(x\{v'\})[v/x] \triangleq x\{v'[v/x]\} \\
(y\{v'\})[v/x] \triangleq y\{v'[v/x]\} \\
\dots
\end{array}$$

$$\begin{array}{ccc}
\frac{}{\text{protected} \sqsubseteq \text{public}} & \frac{}{\text{public} \not\sqsubseteq \text{protected}} & \frac{S.p \sqsubseteq S(z_b).p}{\text{accessible}(S, z_b)}
\end{array}$$

Figure 14: Non-speculative semantics auxiliary definitions

Note: terminal states in the speculative semantics are of the form $\langle \Phi \mid v \rangle$ where $\Phi.\Xi = \bullet$.

$$\boxed{\langle \Phi \mid e \rangle \xrightarrow{\bar{\delta}} \langle \Phi \mid e \rangle}$$

$$\begin{array}{c}
\text{SPEC-NONSPEC} \\
(a', \text{nonspec}) = \Phi.A.\text{spec}(\Phi.A.a, \Phi.S, e) \\
\frac{\langle \Phi \mid K[e] \rangle \xrightarrow{\bar{\delta}} \langle \Phi' \mid K[e'] \rangle}{\langle \Phi \mid K[e] \rangle \xrightarrow{\bar{\delta}} \langle \Phi'[a := a'] \mid K[e'] \rangle}
\end{array}$$

$$\begin{array}{c}
\text{SPEC-SPEC} \\
(a', \text{spec } v) = \Phi.A.\text{spec}(\Phi.A.a, \Phi.S, e) \quad \text{nonfinal}(\langle \Phi \mid K[e] \rangle) \\
\frac{\langle \Phi \mid e \rangle \xrightarrow{\bar{\delta}^*} \langle \Phi' \mid v' \rangle \quad \bar{\Xi}' = \text{makeFrame}_{v=v'}(\Phi.S, K[e], \bar{\delta}) :: \Phi.\Xi}{\langle \Phi \mid K[e] \rangle \xrightarrow{0} \langle \Phi[\Xi := \bar{\Xi}', a := a'] \mid K[v] \rangle}
\end{array}$$

$$\begin{array}{c}
\text{SPEC-FENCE} \\
(a', \text{fence}) = \Phi.A.\text{spec}(\Phi.A.a, \Phi.S, e) \\
\frac{\text{fence } \langle \Phi \mid K[e] \rangle \text{ to } \langle \Phi' \mid e' \rangle_{\bar{\delta}}}{\langle \Phi \mid K[e] \rangle \xrightarrow{\bar{\delta}} \langle \Phi'[a := a'] \mid e' \rangle}
\end{array}$$

$$\boxed{\langle \Phi \mid e \rangle \xrightarrow{\bar{\delta}} \langle \Phi \mid e \rangle}$$

$$\boxed{\text{stalled}(\bar{\Xi}, \delta)}$$

$$\begin{array}{c}
\text{SPEC-}\beta \\
\langle \Phi.S \mid \bullet :: K[e]^{\ell} \rangle \xrightarrow{\epsilon} \langle S' \mid \bar{K}' :: e'^{\ell'} \rangle \\
\frac{\neg \text{stalled}(\Phi.\Xi, \delta(\epsilon)) \quad \bar{\Xi} = \text{addEvent}(\Phi.\Xi, \delta(\epsilon))}{\langle \Phi \mid K[e] \rangle \xrightarrow{\delta(\epsilon)} \langle \Phi[S := S', \Xi := \bar{\Xi}] \mid \bar{K}'[e'] \rangle}
\end{array}$$

$$\begin{array}{c}
\text{STALL-FENCE} \\
\Phi.\Xi \neq \bullet \\
\text{stalled}(\Phi, \text{fence})
\end{array}$$

$$\begin{array}{c}
\text{STALL-PROTECT} \\
\Phi.\Xi \neq \bullet \\
\text{stalled}(\Phi, \text{protect}_p)
\end{array}$$

$$\begin{array}{c}
\text{STALL-READ} \\
\Phi.\Xi = \Xi :: \bar{\Xi} \\
\frac{(\text{protect}_p \in \Xi.\bar{\delta} \wedge \Phi.S(z_b).p = \text{protected}) \vee (\text{stalled}(\Phi[\Xi := \bar{\Xi}], \text{read } v \leftarrow z_b[z_o]))}{\text{stalled}(\Phi, \text{read } v \leftarrow z_b[z_o])}
\end{array}$$

$$\begin{array}{c}
\text{STALL-WRITE} \\
\Phi.\Xi = \Xi :: \bar{\Xi} \\
\frac{(\text{protect}_p \in \Xi.\bar{\delta} \wedge \Phi.S(z_b).p = \text{protected}) \vee (\text{stalled}(\Phi[\Xi := \bar{\Xi}], \text{write } v \mapsto z_b[z_o]))}{\text{stalled}(\Phi, \text{write } v \mapsto z_b[z_o])}
\end{array}$$

Figure 15: Small step speculative semantics

$$\boxed{\text{fence } \langle \Phi \mid e \rangle \text{ to } \langle \Phi \mid e \rangle_{\bar{\delta}}}$$

$$\overline{\Xi} = \text{addEvent}(\overline{\Xi}, \delta)$$

$\frac{\text{FENCE-NO-SPEC} \quad \Phi, \Xi = \bullet}{\text{fence } \langle \Phi \mid e \rangle \text{ to } \langle \Phi \mid e \rangle \bullet}$	$\frac{\text{FENCE-ROLLBACK} \quad \Phi, \Xi = \overline{(S, e')} :: \overline{\Xi}}{\text{fence } \langle \Phi \mid e \rangle \text{ to } \langle \Phi[S := S, \Xi := \overline{\Xi}] \mid e' \rangle \bullet}$	$\frac{\text{FENCE-COMMIT} \quad \Phi, \Xi = (S, e', \overline{\delta}, \overline{\mu}) :: \overline{\Xi} \quad S' = \text{commit}(\Phi.S, \overline{\delta} \diamond \overline{\mu}) \quad \overline{\Xi}' = \text{addEvents}(\overline{\Xi}, \overline{\mu})}{\text{fence } \langle \Phi \mid e \rangle \text{ to } \langle \Phi[S := S', \Xi := \overline{\Xi}'] \mid e \rangle_{\overline{\delta}}}$
$\frac{\text{ADD-NO-SPEC}}{\bullet = \text{addEvent}(\bullet, \delta)}$	$\frac{\text{ADD-MISPEC}}{\overline{(S, e)} :: \overline{\Xi} = \text{addEvent}(\overline{(S, e)} :: \overline{\Xi}, \delta)}$	$\frac{\text{ADD-NON-MEM} \quad \delta' \neq \mu}{(S, e, \overline{\delta}, \overline{\mu}) :: \overline{\Xi} = \text{addEvent}((S, e, \overline{\delta}, \overline{\mu}) :: \overline{\Xi}, \delta)}$
$\frac{\text{ADD-BAD-READ} \quad v_r \in \text{writeLocs}(\overline{\delta})}{\overline{(S, e)} = \text{addEvent}((S, e, \overline{\delta}, \overline{\mu}) :: \overline{\Xi}, \text{read } v \leftarrow v_r)}$	$\frac{\text{ADD-GOOD-READ} \quad v_r \notin \text{writeLocs}(\overline{\delta})}{(S, e, \overline{\delta}, \overline{\mu}) :: \overline{\Xi} = \text{addEvent}((S, e, \overline{\delta}, \overline{\mu}) :: \overline{\Xi}, \text{read } v \leftarrow v_r)}$	
$\frac{\text{ADD-WRITE}}{(S, e, \overline{\delta}, \overline{\mu}) \diamond \text{write } v \mapsto v_w = \text{addEvent}((S, e, \overline{\delta}, \overline{\mu}) :: \overline{\Xi}, \text{write } v \mapsto v_w)}$		
$\frac{\text{ADD-NEW}}{(S, e, \overline{\delta}, \overline{\mu}) \diamond \text{new}_p z @ z_b = \text{addEvent}((S, e, \overline{\delta}, \overline{\mu}) :: \overline{\Xi}, \text{new}_p z @ z_b)}$	$\frac{\text{ADD-PROTECT}}{(S, e, \overline{\delta}, \overline{\mu}) \diamond \text{protect}_p = \text{addEvent}((S, e, \overline{\delta}, \overline{\mu}) :: \overline{\Xi}, \text{protect}_p)}$	
$\begin{aligned} \delta(\epsilon^\ell) &\triangleq \delta(\epsilon) \\ \delta(\tau^{\ell \rightarrow \ell'}) &\triangleq \delta(\tau) \\ \delta(\text{begin}) &\triangleq 0 \\ \delta(\text{end } v) &\triangleq \text{end } v \\ \delta(\text{call } f(\overline{v})) &\triangleq \text{call } f(\overline{v}) \\ \delta(\text{ret } v) &\triangleq 0 \\ \delta \text{new}_p z @ z_b &\triangleq \text{new}_p z @ z_b \\ \delta(\text{read } v \leftarrow z_b[z_o]) &\triangleq \text{read } v \leftarrow z_b[z_o] \\ \delta(\text{write } v \mapsto z_b[z_o]) &\triangleq \text{write } v \mapsto z_b[z_o] \\ \delta(\text{branch } v) &\triangleq \text{branch } v \\ \delta(\text{fence}) &\triangleq \text{fence} \\ \delta(\text{protect}_p) &\triangleq \text{protect}_p \\ \delta(0) &\triangleq 0 \end{aligned}$		
$\text{writeLocs}(\overline{c}) \triangleq \{v_w \mid \text{write } v \mapsto v_w \in \overline{c}\}$		
$\begin{aligned} \text{commit}(S, \bullet) &\triangleq S \\ \text{commit}(S, \delta \diamond \overline{\delta}') &\triangleq \text{commit}(\text{commit}(S, \delta), \overline{\delta}') \\ \text{commit}(S, \text{write } v \mapsto z_b[z_o]) &\triangleq S[z_b].[z_o := v] \\ \text{commit}(S, \text{protect}_p) &\triangleq S[p := p] \\ \text{commit}(S, \text{new}_p z @ z_b) &\triangleq S[z_b := \{size = z, v = \perp, p = p\}] \\ \text{commit}(S, \delta) &\triangleq S \end{aligned}$		
$\begin{aligned} \text{makeFrame}_\top(S, e, \overline{\delta}) &\triangleq (S, e, \overline{\delta}, \bullet) \\ \text{makeFrame}_\perp(S, e, \overline{\delta}) &\triangleq \overline{(S, e)} \end{aligned}$		

Figure 16: Speculative semantics auxiliary functions

$$\boxed{\langle S \mid \overline{K}^\ell :: e^\ell \rangle \xrightarrow{\bar{\epsilon}}_C \langle S \mid \overline{K}^\ell :: e^\ell \rangle}$$

$$\frac{\langle S \mid \overline{K}^\ell :: e^{\ell'} \rangle \xrightarrow{\bar{\epsilon}} \langle S \mid \overline{K_2^{\ell_2}} :: e_2^{\ell_2'} \rangle}{\overline{\mu} = [\text{read } v \leftarrow z_b[z_o] \mid \text{public} \sqsubseteq S(z_b) \wedge z_o \in [S(z_b).size] \wedge v = S(z_b).v(z_o)]} \langle S \mid \overline{K}^\ell :: e^{\ell'} \rangle \xrightarrow{\overline{\mu}^{\text{app}} \circ \bar{\epsilon}}_C \langle S \mid \overline{K_2^{\ell_2}} :: e_2^{\ell_2'} \rangle}$$

$$\boxed{\langle \Phi \mid e \rangle \xrightarrow{\bar{\delta}}_C \langle \Phi \mid e \rangle}$$

$$\frac{\langle \Phi \mid e \rangle \xrightarrow{\bar{\delta}} \langle \Phi' \mid e' \rangle}{\overline{\mu} = [\mu \mid \langle \Phi[S.p := \text{public}] \mid !z_b[z_o] \rangle \xrightarrow{\mu} \langle \Phi' \mid v \rangle]} \langle \Phi \mid e \rangle \xrightarrow{\overline{\mu} \circ \bar{\delta}}_C \langle \Phi' \mid e' \rangle}$$

Figure 17: Concurrent observer semantics

B ATTACKER MODELS

$$\begin{aligned}
A & ::= \tau^{\text{lib} \rightarrow \text{app}} \diamond \overline{\delta^{\text{app}}} \diamond \tau^{\text{app} \rightarrow \text{lib}} \\
L & ::= \tau^{\text{app} \rightarrow \text{lib}} \diamond \overline{\delta^{\text{lib}}} \diamond \tau^{\text{lib} \rightarrow \text{app}} \\
T & ::= \tau^{\text{lib} \rightarrow \text{app}} \diamond \overline{\delta^{\text{app}}} \diamond (\text{end } v)^{\text{app} \rightarrow \text{lib}} \mid A \circ L \circ T
\end{aligned}$$

Figure 18: Traces

$$\begin{aligned}
\text{API contexts } \Gamma & ::= \bullet \mid f \mapsto z :: \Gamma \\
\text{libraries } L & ::= \bullet \mid (f, \lambda_{\text{lib}} \bar{x}. e) :: L \\
\text{secret contexts } \Delta & ::= \bullet \mid x \mapsto (z_b, z) :: \Delta \\
\text{sets of exposed blocks } \sigma & : \quad 2^{\mathbb{Z}}
\end{aligned}$$

Figure 19: Syntax of programs

$$\begin{array}{c}
\boxed{\Gamma \vDash L} \\
\frac{}{\bullet \vDash \bullet} \qquad \frac{\text{length}(\bar{x}) = z \quad \Gamma \vDash L}{f \mapsto z :: \Gamma \vDash (f, \lambda_{\text{lib}} \bar{x}. e) :: L} \qquad \frac{f \mapsto z :: \Gamma \vDash L}{f \mapsto z :: \Gamma \vDash (g, \lambda_{\text{lib}} \bar{x}. e) :: L} \\
\boxed{\Delta \vDash S = S} \qquad \boxed{\Delta \vDash S} \\
\frac{}{\bullet \vDash S = S} \qquad \frac{S(z_b).size = S'(z_b).size = z \quad \Delta \vDash S - \{z_b\} = S' - \{z_b\}}{x \mapsto (z_b, z) :: \Delta \vDash S = S'} \qquad \frac{\Delta \vDash S = S}{\Delta \vDash S} \\
\boxed{\Gamma, \Delta \vDash e} \\
\frac{\text{fv}(e) \subseteq \text{dom}(\Gamma) \cup \text{dom}(\Delta) \quad \lambda_{\text{lib}} \bar{x}. e \notin e}{\Gamma, \Delta \vDash e}
\end{array}$$

Figure 20: Program well-formedness judgments

$$\begin{aligned}
\text{traces}(\langle S \mid e \rangle) &\triangleq \{T \mid \langle S \mid \bullet :: e \rangle \xrightarrow{T}^* \langle S' \mid \bullet :: v \rangle\} \\
\text{concurrentTraces}(\langle S \mid e \rangle) &\triangleq \{T \mid \langle S \mid \bullet :: e \rangle \xrightarrow{T}^*_C \langle S' \mid \bullet :: v \rangle\} \\
\text{specTraces}(\langle \Phi \mid e \rangle) &\triangleq \{\bar{\delta} \mid \langle \Phi \mid e \rangle \xrightarrow{\bar{\delta}}^* \langle \Phi' \mid v \rangle \wedge \Phi'.\Xi = \bullet\} \\
\text{concurrentSpecTraces}(\langle \Phi \mid e \rangle) &\triangleq \{\bar{\delta} \mid \langle \Phi \mid e \rangle \xrightarrow{\bar{\delta}}^*_C \langle \Phi' \mid v \rangle \wedge \Phi'.\Xi = \bullet\} \\
e[\bullet] &\triangleq e \\
e[x \mapsto (z_b, z) :: \Delta] &\triangleq (e[z_b/x])[\Delta] \\
e[\bullet] &\triangleq e \\
e[(f, \lambda_{1\text{lib}} \bar{x}.e :: L)] &\triangleq (e[\lambda_{1\text{lib}} \bar{x}.e/f])[L]
\end{aligned}$$

Figure 21: Traces

For a given API context Γ , attackers (contexts) are terms $\Gamma \vdash e$. We define four attacker models: unrestricted attackers, read-only attackers, memory-safe attackers, and speculative attackers.

DEFINITION 5 (APPLICATIONS). For a given API context Γ and secret context Δ , an application is an expression e such that $\Gamma, \Delta \vdash e$.

DEFINITION 6 (READ-ONLY ATTACKERS). We say an application $\Gamma, \Delta \vdash e$ is a read-only attacker if, for all libraries $\Gamma \vDash L$, initial states $\Delta \vDash S$, and traces $T \in \text{traces}(\langle S \mid e[\Delta][L] \rangle)$, $(\text{dom}(\Delta), \emptyset) \vdash \text{read-only } T$.

DEFINITION 7 (MEMORY-SAFE ATTACKERS). We say an application $\Gamma, \Delta \vdash e$ is a read-only attacker if, for all libraries $\Gamma \vDash L$, initial states $\Delta \vDash S$, and traces $T \in \text{traces}(\langle S \mid e[\Delta][L] \rangle)$, $(\text{dom}(\Delta), \emptyset) \vdash \text{mem-safe } T$.

To capture speculative attackers we consider (non-speculatively) memory safe attackers run in the speculative semantics.

$$\boxed{(\sigma_A, \sigma_L) \vdash \text{read-only } T}$$

$$\frac{A = \tau_1^{\text{lib} \rightarrow \text{app}} \diamond \bar{\delta}_A^{\text{app}} \diamond \tau_2^{\text{app} \rightarrow \text{lib}} \quad L = \tau_2^{\text{app} \rightarrow \text{lib}} \diamond \bar{\delta}_L^{\text{lib}} \diamond \tau_3^{\text{lib} \rightarrow \text{app}} \quad \sigma_A \cup \text{exposed}(\tau_1) \vdash \text{wf-read-only } \bar{\delta}_A + \sigma'_A \quad \sigma_L \cup \text{exposed}(\tau_2) \vdash \text{wf-read-only } \bar{\delta}_L + \sigma'_L \implies (\sigma'_A, \sigma'_L) \vdash \text{read-only } T}{(\sigma_A, \sigma_L) \vdash \text{read-only } A \circ L \circ T}$$

$$\frac{\sigma_A \cup \text{exposed}(\tau) \vdash \text{wf-read-only } \bar{\delta} + \sigma'_A}{(\sigma_A, \sigma_L) \vdash \text{read-only } \tau^{\text{lib} \rightarrow \text{app}} \diamond \bar{\delta}^{\text{app}} \diamond (\text{end } v)^{\text{app} \rightarrow \text{lib}}}$$

$$\boxed{(\sigma_A, \sigma_L) \vdash \text{mem-safe } T}$$

$$\frac{A = \tau_1^{\text{lib} \rightarrow \text{app}} \diamond \bar{\delta}_A^{\text{app}} \diamond \tau_2^{\text{app} \rightarrow \text{lib}} \quad L = \tau_2^{\text{app} \rightarrow \text{lib}} \diamond \bar{\delta}_L^{\text{lib}} \diamond \tau_3^{\text{lib} \rightarrow \text{app}} \quad \sigma_A \cup \text{exposed}(\tau_1) \vdash \text{wf-mem-safe } \bar{\delta}_A + \sigma'_A \quad \sigma_L \cup \text{exposed}(\tau_2) \vdash \text{wf-read-only } \bar{\delta}_L + \sigma'_L \implies (\sigma'_A, \sigma'_L) \vdash \text{read-only } T}{(\sigma_A, \sigma_L) \vdash \text{mem-safe } A \circ L \circ T}$$

$$\frac{\sigma_A \cup \text{exposed}(\tau) \vdash \text{wf-mem-safe } \bar{\delta} + \sigma'_A}{(\sigma_A, \sigma_L) \vdash \text{mem-safe } \tau^{\text{lib} \rightarrow \text{app}} \diamond \bar{\delta}^{\text{app}} \diamond (\text{end } v)^{\text{app} \rightarrow \text{lib}}}$$

$$\boxed{\sigma \vdash \text{wf-read-only } \bar{\delta} + \sigma}$$

$$\frac{}{\sigma \vdash \text{wf-read-only } \text{new}_p \ z @ z_b + \sigma \cup \{z_b\}} \quad \frac{}{\sigma \vdash \text{wf-read-only } \text{read } v \leftarrow z_b[z_0] + \sigma} \quad \frac{z_b \in \sigma}{\sigma \vdash \text{wf-read-only } \text{write } v \mapsto z_b[z_0] + \sigma}$$

$$\frac{}{\sigma \vdash \text{wf-read-only } \text{call } f(\bar{v}) + \sigma} \quad \frac{}{\sigma \vdash \text{wf-read-only } \text{branch } v + \sigma} \quad \frac{}{\sigma \vdash \text{wf-read-only } \text{protect}_p + \sigma} \quad \frac{}{\sigma \vdash \text{wf-read-only } 0 + \sigma}$$

$$\frac{}{\sigma \vdash \text{wf-read-only } \emptyset + \sigma} \quad \frac{\sigma \vdash \text{wf-read-only } \delta + \sigma' \quad \sigma' \vdash \text{wf-read-only } \bar{\delta} + \sigma''}{\sigma \vdash \text{wf-read-only } \delta \diamond \bar{\delta} + \sigma''}$$

$$\begin{array}{c}
\boxed{\sigma \vdash \text{wf-mem-safe } \bar{\delta} \dashv \sigma} \\
\hline
\frac{}{\sigma \vdash \text{wf-mem-safe new}_p z@z_b \dashv \sigma \cup \{z_b\}} \quad \frac{z_b \in \sigma}{\sigma \vdash \text{wf-mem-safe read } v \leftarrow z_b[z_o] \dashv \sigma} \quad \frac{z_b \in \sigma}{\sigma \vdash \text{wf-mem-safe write } v \mapsto z_b[z_o] \dashv \sigma} \\
\hline
\frac{}{\sigma \vdash \text{wf-mem-safe call } f(\bar{v}) \dashv \sigma} \quad \frac{}{\sigma \vdash \text{wf-mem-safe branch } v \dashv \sigma} \quad \frac{}{\sigma \vdash \text{wf-mem-safe protect}_p \dashv \sigma} \quad \frac{}{\sigma \vdash \text{wf-mem-safe } 0 \dashv \sigma} \\
\hline
\frac{}{\sigma \vdash \text{wf-mem-safe } \emptyset \dashv \sigma} \quad \frac{\sigma \vdash \text{wf-mem-safe } \delta \dashv \sigma' \quad \sigma' \vdash \text{wf-mem-safe } \bar{\delta} \dashv \sigma''}{\sigma \vdash \text{wf-mem-safe } \delta \diamond \bar{\delta} \dashv \sigma''}
\end{array}$$

Figure 22: Attacker model judgments

B.1 Security definitions

$$\begin{array}{l}
\text{ct}(\epsilon^\ell) \triangleq \text{ct}(\epsilon) \\
\text{ct}(\tau^{\ell \rightarrow \ell'}) \triangleq \text{ct}(\tau) \\
\text{ct}(\text{begin}) \triangleq 0 \\
\text{ct}(\text{end } v) \triangleq \text{end } v \\
\text{ct}(\text{call } f(\bar{v})) \triangleq \text{call } f \\
\text{ct}(\text{ret } v) \triangleq 0 \\
\text{ct}(\text{new}_p z@z_b) \triangleq 0 \\
\text{ct}(\text{read } v \leftarrow z_b[z_o]) \triangleq \text{read } \leftarrow z_b[z_o] \\
\text{ct}(\text{write } v \mapsto z_b[z_o]) \triangleq \text{write } \mapsto z_b[z_o] \\
\text{ct}(\text{branch } v) \triangleq \text{branch } v \\
\text{ct}(\text{fence}) \triangleq 0 \\
\text{ct}(\text{protect}_p) \triangleq 0 \\
\text{ct}(0) \triangleq 0
\end{array}$$

Figure 23: Constant time events

$e_\Gamma ::= f(\bar{v})$ with $f \in \Gamma$.

DEFINITION 8 (CLASSICAL CONSTANT TIME). We say a library $\Gamma \vDash L$ is classically constant time if, for all secret contexts Δ , classical “applications” Γ , $\Delta \vDash e_\Gamma$, and initial states $\Delta \vDash S = S'$ we have that $\text{ct}(\text{traces}(\langle S \mid e_\Gamma[\Delta][L] \rangle)) = \text{ct}(\text{traces}(\langle S' \mid e_\Gamma[\Delta][L] \rangle))$.

DEFINITION 9 (ROBUST CONSTANT TIME). We say a library $\Gamma \vDash L$ is robustly constant time for an attacker class pred if, for all secret contexts Δ , applications Γ , $\Delta \vDash e$ such that $\text{pred}(\Gamma, \Delta \vdash e)$, and initial states $\Delta \vDash S = S'$ we have that $\text{ct}(\text{traces}(\langle S \mid e[\Delta][L] \rangle)) = \text{ct}(\text{traces}(\langle S' \mid e[\Delta][L] \rangle))$.

DEFINITION 10 (ROBUST CONSTANT TIME FOR CONCURRENT OBSERVERS). We say a library $\Gamma \vDash L$ is robustly constant time for concurrent observers if, for all secret contexts Δ , read-only applications Γ , $\Delta \vDash e$, and initial states $\Delta \vDash S = S'$ we have that $\text{ct}(\text{concurrentTraces}(\langle S \mid e[\Delta][L] \rangle)) = \text{ct}(\text{concurrentTraces}(\langle S' \mid e[\Delta][L] \rangle))$.

DEFINITION 11 (CLASSICAL SPECULATIVE CONSTANT TIME). We say a library $\Gamma \vDash L$ is classically speculative constant time with respect to a speculation oracle spec with microarchitectural state type A if, for all secret contexts Δ , classical “applications” Γ , $\Delta \vDash e_\Gamma$, initial states $\Delta \vDash S = S'$, microarchitectural state $a : A$, $\Phi_1 = \{S = S, A = \{A = A, a = a, \text{spec} = \text{spec}\}, \Xi = \bullet\}$, and $\Phi'_1 = \{S = S', A = \{A = A, a = a, \text{spec} = \text{spec}\}, \Xi = \bullet\}$, we have that for all traces $\langle \Phi_1 \mid e_\Gamma[\Delta][L] \rangle \xrightarrow{\bar{\delta}}^* \langle \Phi_2 \mid v \rangle$ there exists a trace $\langle \Phi'_1 \mid e_\Gamma[\Delta][L] \rangle \xrightarrow{\bar{\delta}'}^* \langle \Phi'_2 \mid v \rangle$ such that $\text{ct}(\bar{\delta}) = \text{ct}(\bar{\delta}')$, $\Phi_2.\Xi = \Phi'_2.\Xi = \bullet$, and $\Phi_2.A.a = \Phi'_2.A.a$.

DEFINITION 12 (ROBUST SPECULATIVE CONSTANT TIME). We say a library $\Gamma \vDash L$ is robustly speculative constant time with respect to a speculation model A , if, for all secret contexts Δ , memory-safe applications Γ , $\Delta \vDash e$, initial states $\Delta \vDash S = S'$, $\Phi = \{S = S, A = A, \Xi = \bullet\}$, and $\Phi' = \{S = S', A = A, \Xi = \bullet\}$, we have that $\text{ct}(\text{specTraces}(\langle \Phi \mid e[\Delta][L] \rangle)) = \text{ct}(\text{specTraces}(\langle \Phi' \mid e[\Delta][L] \rangle))$.

DEFINITION 13 (ROBUST SPECULATIVE CONSTANT TIME FOR CONCURRENT OBSERVERS). We say a library $\Gamma \vDash L$ is robustly speculative constant time for concurrent observers with respect to a speculation model A , if, for all secret contexts Δ , memory-safe applications Γ , $\Delta \vDash e$, initial states $\Delta \vDash S = S'$, $\Phi = \{S = S, A = A, \Xi = \bullet\}$, and $\Phi' = \{S = S', A = A, \Xi = \bullet\}$, we have that $\text{ct}(\text{concurrentSpecTraces}(\langle \Phi \mid e[\Delta][L] \rangle)) = \text{ct}(\text{concurrentSpecTraces}(\langle \Phi' \mid e[\Delta][L] \rangle))$.

C COMPILER PROOFS

THEOREM 2 ($\mathbb{C}_{\text{read-only}}$ GUARANTEES READ-ONLY ROBUST CONSTANT TIME). *If $\Gamma \vDash L$ is classically constant time and does not contain any protect_p subterms, then $\mathbb{C}_{\text{read-only}}(\Gamma \vDash L)$ is robustly constant time for read-only attackers (that do not contain protect_p).*

PROOF SKETCH. We proceed by simultaneous induction over the operational semantics and $(\text{dom}(\Delta), \emptyset) \vdash \text{read-only } T$ for each trace T under S and S' . We derive the latter from our assumption that e is read-only instantiated with L and S or S' . We inductively build a partial isomorphism between the sets of traces, mapping each trace prefix under S to an identical trace prefix under S' and additionally carrying a map of those locations that vary between the two traces' states.

We maintain the invariant that the varying components of the states in the partial isomorphism are in the memory page protected and that the protection level during application execution is set to `public`. This invariant allows us to employ angelic non-determinism during application execution to find an exact equivalent action in the other set of traces: the visible components of the state are exactly equal. We further rely on the fact that allocation is a constant-time observable event and therefore, given our assumption that L is classically constant time, there must exist traces with the exact same library allocated blocks. (This is in fact why allocation is considered a timing side-channel.) For library sub-sequences of the trace, we apply our assumption that the library is classically constant time and therefore the sub-sequence produces invariant traces. \square

THEOREM 3 (\mathbb{C}_{spec} GUARANTEES ROBUST SPECULATIVE CONSTANT TIME). *If $\Gamma \vDash L$ is classically speculative constant time for a speculation oracle spec and does not contain any protect_p subterms, then $\mathbb{C}_{\text{spec}}(\Gamma \vDash L)$ is robustly speculatively constant time for any speculation model A such that $A.\text{spec} = \text{spec}$ (for attackers that do not contain protect_p).*

PROOF SKETCH. The proof proceeds similarly to the proof for our read-only compiler. The key distinction is we rely on the fence behavior of protect_p which, along with our invariant, ensures that varying locations are not visible during speculative execution. We additionally rely on the fact that classical speculative constant time requires calls to the library be noninterferent for any microarchitectural state A and that the microarchitectural state is also noninterferent. \square

THEOREM 4 ($\mathbb{C}_{\text{concurrent}}$ GUARANTEES ROBUST CONSTANT TIME FOR CONCURRENT OBSERVERS). *If $\Gamma \vDash L$ is classically constant time and does not contain any protect_p subterms, then $\mathbb{C}_{\text{concurrent}}(\Gamma \vDash L)$ is robustly constant time for concurrent observers (that do not contain protect_p).*

THEOREM 5 ($\mathbb{C}_{\text{concurrent}}$ GUARANTEES ROBUST SPECULATIVE CONSTANT TIME FOR CONCURRENT OBSERVERS). *If $\Gamma \vDash L$ is classically speculative constant time for a speculation oracle spec and does not contain any protect_p subterms, then $\mathbb{C}_{\text{concurrent}}(\Gamma \vDash L)$ is robustly speculatively constant time for concurrent observers for any speculation model A such that $A.\text{spec} = \text{spec}$ (for attackers that do not contain protect_p).*

PROOF SKETCH. The proofs proceed in parallel to their respective non-concurrent versions. When executing within the application, the concurrent observations have no extra visibility so our invariant's guarantee that all unprotected memory is invariant ensure the concurrent read events are also invariant. When executing with the library, and unlike with $\mathbb{C}_{\text{read-only}}$ and \mathbb{C}_{spec} , the compiler guarantees that all memory that varies is in protected memory and thus does not appear in the concurrent trace events. In the speculative case we use the assumption that the library is classically speculatively constant time. \square

D EVALUATION ADDITIONAL DATA

Data size	Q_1	Median overhead	Q_3	Baseline cycles
1	0.27%	3.04%	8.74%	7.14e+03
128	0.12%	1.48%	4.92%	1.6e+04
256	-0.01%	0.74%	2.60%	2.69e+04
512	-0.01%	0.61%	1.58%	5.04e+04
1024	-0.04%	0.33%	1.09%	9.49e+04
2048	-0.17%	0.17%	0.73%	1.85e+05

Table 4: aead overhead by size

Data size	Q_1	Median overhead	Q_3	Baseline cycles
29	-2.09%	-0.4%	0.56%	5.59e+05
59	-2.51%	-0.57%	0.51%	5.64e+05
117	-2.51%	-0.09%	0.64%	5.53e+05
231	-2.13%	-0.2%	0.66%	5.46e+05
453	-1.51%	-0.03%	0.79%	9.24e+05
709	-1.39%	-0.26%	0.73%	1.16e+06
2711	-1.53%	0.12%	0.70%	4.01e+06
4237	-1.89%	-0.02%	0.80%	6.51e+06

Table 5: encrypt overhead by size

Data size	Q_1	Median overhead	Q_3	Baseline cycles
29	-0.3%	0.09%	0.97%	3.22e+08
59	-0.49%	0.02%	0.84%	3.22e+08
117	-0.46%	-0.04%	0.81%	3.22e+08
231	-0.48%	0.01%	0.64%	3.22e+08
453	-0.37%	0.04%	0.95%	3.08e+08
709	-0.38%	0.0%	0.61%	3.08e+08
2711	-0.7%	-0.01%	0.66%	3.08e+08
4237	-0.27%	0.02%	0.92%	2.99e+08

Table 6: sign overhead by size