

A Formal Foundation for Secure Remote Execution of Enclaves

Pramod Subramanyan
University of California, Berkeley
pramod@berkeley.edu

Rohit Sinha
University of California, Berkeley
rsinha@eecs.berkeley.edu

Ilia Lebedev
Massachusetts Institute of Technology
ilebedev@mit.edu

Srinivas Devadas
Massachusetts Institute of Technology
devadas@mit.edu

Sanjit A. Seshia
University of California, Berkeley
sseshia@eecs.berkeley.edu

ABSTRACT

Recent proposals for trusted hardware platforms, such as Intel SGX and the MIT Sanctum processor, offer compelling security features but lack formal guarantees. We introduce a verification methodology based on a trusted abstract platform (TAP), a formalization of idealized enclave platforms along with a parameterized adversary. We also formalize the notion of secure remote execution and present machine-checked proofs showing that the TAP satisfies the three key security properties that entail secure remote execution: integrity, confidentiality and secure measurement. We then present machine-checked proofs showing that SGX and Sanctum are refinements of the TAP under certain parameterizations of the adversary, demonstrating that these systems implement secure enclaves for the stated adversary models.

CCS CONCEPTS

• **Security and privacy** → **Formal methods and theory of security**; **Security in hardware**; **Trusted computing**; **Information flow control**;

KEYWORDS

Enclave Programs; Secure Computation; Formal Verification; Confidentiality; Integrity; Remote Attestation

1 INTRODUCTION

An application executing on a typical computing platform contains large privileged software layers in its trusted computing base (TCB), which includes the operating system (OS), hypervisor and firmware. Vulnerabilities in these privileged software layers have been exploited to enable the execution of privileged malware with disastrous consequences [1, 26, 35, 52, 64]. To address this problem, processor designers and vendors are now developing CPUs with hardware primitives, such as Intel SGX and MIT Sanctum, which isolate sensitive code and data within protected memory regions, called *enclaves*, that are inaccessible to other software running on the machine. With these primitives, enclaves become the only

trusted components of an application, and a carefully programmed enclave can execute safely even when the privileged software layers are compromised.

Despite growing interest, there has only been informal security analysis of these new platforms, which we refer to as *enclave platforms*. This lack of formalization has a number of consequences. Developers of enclave programs cannot formally reason about security of their programs; incorrect use of hardware primitives can render enclave programs vulnerable to security breaches. Hardware designers cannot formally state security properties of their architectures and are unable to reason about potential vulnerabilities of enclaves running on their hardware.

Going beyond the difficulty of reasoning about specific programs and platforms, this lack of formalization also makes it difficult to compare and contrast potential improvements to these platforms. For example, a number of proposals have developed software defenses to strengthen the security guarantees of Intel SGX enclaves [55, 68, 69]. However, comparing the security guarantees and adversary models (of say T-SGX [68] and Shinde et al. [69]) is difficult without a unified framework for such reasoning. Furthermore, as we move towards a world in which enclave platforms are widespread, it is conceivable that future data centers will support a number of different enclave platforms. Developers will likely tackle this diversity by relying on tools and libraries that provide a common application programming interface (API) for enclave programs while supporting different target platforms. Reasoning about the security guarantees of such toolchains is also challenging without a unified framework. This paper bridges each of the above gaps by presenting a unified formal framework to specify and verify the security properties of enclave platforms.

We address the formal modeling and verification of enclave platforms in three parts. First, we define the properties required for *secure remote execution* of enclaves. Next we introduce the *trusted abstract platform (TAP)*, an idealization of enclave platforms along with a parameterized adversary model. We present machine-checked proofs showing that the TAP provides secure remote execution against these adversaries. Finally, we present machine-checked proofs demonstrating that formal models of *Intel SGX* and *MIT Sanctum* are refinements [13] of the TAP for different parameterizations of the adversary, and thus also provide secure remote execution.

Secure Remote Execution of Enclaves: This paper first formalizes a model of computation of an enclave program, and the attacker's operations and observations — we assume a privileged software adversary that has compromised the host OS, hypervisor, network, and persistent storage. The execution model allows the developer (or user) to precisely define the expected runtime

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CCS '17, October 30–November 3, 2017, Dallas, TX, USA

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-4946-8/17/10...\$15.00

<https://doi.org/10.1145/3133956.3134098>

behavior of an enclave in the presence of a privileged software adversary. When the user outsources an enclave to a remote platform, she seeks a guarantee that the enclave be executed according to the expected behavior, i.e., the platform must respect the enclave's semantics. We term this property *secure remote execution (SRE)*: any execution of that enclave on the trusted enclave platform must be one of enclave's expected executions (formalized in § 3). Any enclave platform, such as SGX or Sanctum, must guarantee SRE to the user, and this paper describes a formal framework and methodology to prove secure remote execution. SRE is decomposed into lower-level properties — specifically, integrity, confidentiality, and secure measurement — for which we develop machine-checked proofs on models of SGX and Sanctum.

Trusted Abstract Platform: We develop an idealized abstraction of the aforementioned enclave platforms, named Trusted Abstract Platform (TAP), consisting of a small set of formally-specified primitives sufficient to implement enclave execution. As a precursor to proving that TAP satisfies SRE, we define a parameterized model of a privileged software attacker with varying capabilities. We then present machine-checked proofs showing that TAP satisfies the integrity, confidentiality, and secure measurement properties for these attackers. The TAP is a general framework for reasoning about and comparing different enclave platforms, adversary models and security guarantees. For enclave platform implementers, the TAP is a golden model or specification of platform behavior. From the perspective of enclave program developers, the TAP provides a means of reasoning about program security without being bogged down by implementation details of individual enclave platforms.

Refinement: Next, we use the TAP to reason about the security of Intel SGX and MIT Sanctum. We develop formal models of SGX and Sanctum and present machine-checked proofs showing that SGX and Sanctum are *refinements* of our idealized TAP: every operation on SGX and Sanctum can be mapped to a corresponding TAP operation. Since all executions of an enclave on SGX and Sanctum can be simulated by a TAP enclave, and because TAP guarantees SRE, it follows that the SGX and Sanctum models also guarantee SRE. There is a caveat that SGX only refines a version of TAP which leaks some side channel observations to the attacker (see § 5.3), therefore providing a weaker confidentiality guarantee. This form of parameterization demonstrates that the TAP allows us to develop a taxonomy of enclave platforms, each of which provides varying guarantees against different threat models.

1.1 Contributions

This paper makes the following key contributions:

- (1) A formalization of enclave execution in the presence of a privileged software adversary, and the notion of secure remote execution (SRE) of an enclave program.
- (2) The Trusted Abstract Platform, which formally specifies the semantics of a small set of trusted primitives for safely creating and destroying enclaves, entering and exiting from enclaves, and generating attested statements from within enclaves. We also define a parameterized attacker model, for which the TAP gives varying degrees of confidentiality.
- (3) Decomposition of the SRE property into a triad of integrity, confidentiality, and secure measurement properties, which are more easily formally verified using a theorem prover.
- (4) A refinement-based methodology for proving the SRE guarantee of enclave platforms and machine-checked proofs of refinement for models of SGX and Sanctum.

All our models and proof scripts are being made open-source [75]. These models are designed to be modular and amenable to extension by the community.

2 FORMAL MODEL OF ENCLAVE EXECUTION

An *enclave platform* implements primitives to create protected memory regions, called *enclaves*, that contain both code and data and are isolated from all other software in the system. The processor monitors all accesses to the enclave: only code running in the enclave can access the enclave's memory. As an example, Intel's SGX instructions enable the creation of user-mode enclaves in the hosting application's address space. Since the privileged software layers (OS/Hypervisor) cannot be trusted to modify enclave memory (which is why system calls are disabled in enclave mode), the enclave platform allows the enclave to access the entire address space of the hosting application. This enables efficient I/O interaction with the external world. The external world can only transfer control to the enclave at statically-defined locations called *entrypoints*. In addition to enabling isolated execution, the enclave platform implements primitives for generating attested statements: code inside an enclave can get messages signed using a per-processor private key along with a hash-based measurement of the enclave. This allows other trusted entities to verify that messages originated from the desired enclave running on a genuine platform. Finally, we assume the enclave platform implements a cryptographically secure random number generator which the enclave can use for cryptographic operations such as key generation.




Figure 1: Execution of Enclave Program

To outsource the enclave's execution, the user sends the enclave program to a remote machine over an untrusted channel (Figure 1). The untrusted OS invokes the enclave platform's primitives to launch an enclave containing the program. While running, an enclave may invoke the enclave platform's primitives to get attested statements (with hash-based measurement $\mu(e)$) and random bits.

The enclave may also send outputs to the user by proxying them via the host application's unprotected memory. We present a sample application in § 3.3 to discuss details of secure remote execution of an enclave program.

2.1 Formal Model of Enclave Programs

An enclave is launched with a set of code pages, containing user-mode instructions, and a set of data pages, which can be used for private heap and stack space. This forms the initial state of the enclave — the enclave platform includes a hash-based measurement of this initial state in all attested statements, which allows the user to verify that the enclave was launched in an expected initial state. The enclave's configuration also includes (1) the entrypoint, (2) the virtual address range *evrange* which maps the enclave's protected memory (i.e., each address in *evrange* either maps to a physical address that is owned by the enclave or is inaccessible due to page permissions) and (3) permissions for each address within *evrange*.

Enclave Programs: The user ships an enclave $e = (init_e, config_e)$. $config_e$ defines the enclave's entrypoint: $config_e.entrypoint$, its virtual address range: $config_e.evrange$, and its access permissions: $config_e.acl$. $init_e$ specifies the enclave's initial state at launch, including values of code and data pages within $config_e.evrange$.

We assume that the enclave's virtual address range and corresponding access permissions ($config_e.evrange$, $config_e.acl$) are the same for local and remote executions. The assumption is required in order to ensure that enclave *measurement* (see § 3.1.1) returns a predictable value.¹

Enclave State: At any point of time, the machine is in some state σ , which we leave abstract for now and defer its definition to § 4.1. The enclave's state $E_e(\sigma)$ is a projection of the machine state, and specifies a valuation of the following state variables: (1) memory $vmem : VA \rightarrow W$ which is a partial map from virtual addresses (within $config_e.evrange$) to machine words, (2) general-purpose registers $regs : \mathbb{N} \rightarrow W$ which are indexed by a natural number, (3) program counter $pc : VA$, and (4) configuration $config_e$, which is copied verbatim from e and remains constant throughout the enclave's execution. $init_e$ specifies the state of enclave's memory ($vmem$) at the time of launch. We abuse notation to also write $init_e(E_e(\sigma))$ to mean that $E_e(\sigma)$ is in its launch-time state (prior to first entry into the enclave).

Enclave Inputs and Outputs: In addition to the private state $E_e(\sigma)$, an enclave accesses non-enclave memory for reading inputs and writing outputs, both of which are under adversary control — since addresses outside *evrange* can be read and modified by both the adversary and the enclave, we assume reads from these addresses return unconstrained values. The enclave may also invoke the enclave platform's primitive to get random numbers (which we also treat as inputs). Therefore, we define enclave's input $I_e(\sigma) \doteq \langle I_e^R(\sigma), I_e^U(\sigma) \rangle$, where $I_e^R(\sigma)$ is any random number that is provided by the platform in that step (ϵ if randomness was not requested), and $I_e^U(\sigma)$ is the projection of the machine state that e may read and the attacker may write; specifically, $I_e^U(\sigma)$ specifies an evaluation of non-enclave memory: a partial map from virtual

addresses (outside $config_e.evrange$) to machine words. Similarly, we define enclave's output $O_e(\sigma)$ to be a projection of the machine state that e can write and the attacker may read; specifically, $O_e(\sigma)$ specifies an evaluation of non-enclave memory: a partial map from virtual addresses (outside $config_e.evrange$) to words. Furthermore, an attacker runs concurrently with e and maintains state that cannot be accessed by the enclave (e.g., hypervisor private memory, other malicious enclaves), which we denote as $A_e(\sigma)$.

Enclave Execution: Our semantics of enclave execution assumes that an enclave program is deterministic modulo the input $I_e(\sigma)$, i.e., the next state of the enclave is a function of the current state $E_e(\sigma)$ and input $I_e(\sigma)$ (which includes $I_e^U(\sigma)$ and $I_e^R(\sigma)$). This is not a restriction in practice as both Sanctum and SGX enclaves interact with the external world via memory-based I/O ($I_e^U(\sigma)$), and besides the random bits ($I_e^R(\sigma)$) from an entropy source, there are no other sources of non-determinism — we only assume single threaded enclaves in this work. Since we make this determinism assumption while defining an enclave's semantics, we must prove that the platform does not allow non-determinism.

The enclave computes by performing steps, where in each step the enclave platform first identifies the instruction to execute (based on the current state of *vmem* and *pc*), and then transitions to the next state based on the operational semantics of the platform's instructions. The instructions include bitvector operations on registers, memory accesses, and the enclave primitives for generating attested statements, randomness, and exiting enclave mode — the platform also includes privileged instructions (e.g., MSR instructions in x86) which can cause arbitrary updates to $A_e(\sigma)$. This model of execution lets us define the platform's transition relation \leadsto , where $(\sigma_i, \sigma_j) \in \leadsto$ indicates that the platform can transition from σ_i to σ_j ; from hereon, we write this in infix form as $\sigma_i \leadsto \sigma_j$. We treat \leadsto as a relation as opposed to a function to allow non-determinism for the platform. The non-determinism includes randomness from entropy sources, direct memory accesses from I/O peripherals, etc. Let us refer to these bits of non-determinism in a particular state as $I^P(\sigma)$, which is only available to the privileged software — the enclave's source of non-determinism is captured in $I_e(\sigma)$. We require that the platform be deterministic relative to $I^P(\sigma)$. A secure platform must also ensure an enclave program e is deterministic relative to its input $I_e(\sigma)$. We state two properties in § 3 that imply these determinism guarantees.

2.2 Formal Model of the Adversary

An enclave executes in the presence of a privileged adversary that has compromised all software layers including the OS, except for the enclave platform. In this section, we abstractly define the effect of the adversary's operations and observations. § 4.2 precisely defines the set of attacker operations and observations for the TAP.

2.2.1 Adversary Tampering. The privileged adversary may pause the enclave at any time, and execute arbitrary instructions that modify the attacker's state $A_e(\sigma)$, enclave's input $I_e^U(\sigma)$ for any enclave e , and launch or destroy any number of enclaves. We model the adversary's effect through the *tamper* relation over pairs of states: $(\sigma_1, \sigma_2) \in \text{tamper}$ if the attacker can change the machine's state from σ_1 to σ_2 , with the constraint that $E_e(\sigma_1) = E_e(\sigma_2)$. That is, the attacker may modify $A_e(\sigma_1)$, $I_e(\sigma_1)$, and $O_e(\sigma_1)$, but not the

¹This is not a restrictive assumption because techniques such as Address Space Layout Randomization (ASLR) of the enclave's virtual memory layout can be implemented after the enclave is initialized à la SGX Shield [67].

enclave's state $E_e(\sigma_1)$. The constraint that *tamper* does not affect an enclave's state is assumed while defining the enclave's semantics. Our integrity property (§ 3) for enclave platforms must prove that this constraint is forced upon the attacker's operations.

$\textit{tamper} \subset \rightsquigarrow$ because a software attacker uses the platform's instructions to update the platform's state. Furthermore, *tamper* is reflexive because the adversary can always leave state unmodified: $\forall \sigma. (\sigma, \sigma) \in \textit{tamper}$. In addition to running concurrently with an enclave e , the adversary may *tamper* the machine's state prior to launch and modify e 's launch state \textit{init}_e and configuration \textit{config}_e .

2.2.2 Adversary Observations. Untrusted software may also *observe* an enclave's execution, depending on the confidentiality guarantees provided by the enclave platform. At the very least, the adversary observes any output, but it may also observe certain side channels such as memory access patterns. Observations are performed by executing arbitrary instructions (e.g., any x86 instruction) and invoking the platform's primitives (e.g., launching other enclaves), and then *observing* the results of these operations. Let $\textit{obs}_e(\sigma)$ denote the result of an observation for the machine state σ . For instance, an attacker that only observes outputs enjoys the observation function $\textit{obs}_e(\sigma) \doteq (O_e(\sigma))$. We specify the observation functions in detail in § 4.2.

2.3 Enclave Execution with an Attacker

An execution trace of the platform is an unbounded-length sequence of states denoted $\pi = \langle \sigma_0, \sigma_1, \dots, \sigma_n \rangle$, such that $\forall i. \sigma_i \rightsquigarrow \sigma_{i+1}$; $\pi[i]$ refers to the i th element of the trace. Since the attacker may pause and resume e at any time, we define e 's execution to be the subsequence of states from π where e is executing. To that end, let the function $\textit{curr}(\sigma)$ denote the current mode of the platform, where $\textit{curr}(\sigma) = e$ iff the platform executes enclave e in state σ . Using this function, we can filter out the steps in π where e is not executing. We write the resulting sequence as $\langle \sigma'_0, \sigma'_1, \dots, \sigma'_m \rangle^2$ where $\textit{init}_e(E_e(\sigma'_0)) \wedge \forall i. \textit{curr}(\sigma'_i) = e$. This subsequence is the enclave's *execution trace*: $\langle (I_e(\sigma'_0), E_e(\sigma'_0), O_e(\sigma'_0)), \dots, (I_e(\sigma'_m), E_e(\sigma'_m), O_e(\sigma'_m)) \rangle$. Since an execution trace of e only includes the steps where e invokes an instruction, the attacker may perform *tamper* between any two consecutive steps of e 's execution trace. Therefore, we also have the property that $\forall i. (\sigma'_i, \sigma'_{i+1}) \in \textit{tamper}$. This has the effect of havocing $A_e(\sigma)$ and $I_e^U(\sigma)$ in all these steps, thus supplying the enclave with fresh inputs at each step.

The semantics of an enclave e , denoted $\llbracket e \rrbracket$, is the set of finite or infinite execution traces, containing an execution trace for each input sequence, i.e., for each value of non-enclave memory and randomness at each step of execution.

$$\llbracket e \rrbracket = \{ \langle (I_e(\sigma'_0), E_e(\sigma'_0), O_e(\sigma'_0)), \dots \rangle \mid \textit{init}_e(E_e(\sigma'_0)) \} \quad (1)$$

We must account for all potential input sequences in $\llbracket e \rrbracket$ because e may receive any value of input at any step. We note that $\llbracket e \rrbracket$ may contain traces of any length, and also contain prefixes of any other trace in $\llbracket e \rrbracket$, i.e., it is prefix-closed. We adopt this definition of $\llbracket e \rrbracket$ because the attacker can pause and destroy the enclave at any time; denial of service is not in scope. Due to the determinism property of enclave programs, a specific sequence of inputs $\langle I_e(\sigma'_0), I_e(\sigma'_1),$

$\dots, I_e(\sigma'_m) \rangle$ uniquely identifies a trace from $\llbracket e \rrbracket$ and determines the expected execution trace of e under that sequence of inputs.

3 SECURE REMOTE EXECUTION OF ENCLAVES

Imagine a user who wishes to outsource the execution of an enclave program e onto a remote platform. The user desires that the platform respect the semantics $\llbracket e \rrbracket$ by executing trace(s) from $\llbracket e \rrbracket$. However, the privileged software layers on the platform are untrusted, therefore the user's trust is based on guarantees provided by the enclave platform. We propose the following notion of secure remote execution (SRE) of enclaves:

Definition 3.1. Secure Remote Execution of Enclaves. A remote platform performs secure execution of an enclave program e if any execution trace of e on the platform is contained within $\llbracket e \rrbracket$. Furthermore, the platform must guarantee that a privileged software attacker only observes a projection of the execution trace, as defined by the observation function \textit{obs} .

It is important to note that SRE does not force the platform to execute e — the attacker may deny service, and this is easily detectable by the user because the attacker cannot forge attested statements as if they originated from the user's enclave. Nor are we forcing the platform to execute e a fixed number of times. The attacker has the capability to execute e as many times as it wishes, and a user can easily defend against these attacks by refusing to provision secrets to other copies of the enclave. With that said, SRE requires the platform to execute traces from $\llbracket e \rrbracket$, and recall that $\llbracket e \rrbracket$ only contains enclave executions that start in the initial state of the enclave (see Equation 1). Furthermore, this definition of $\llbracket e \rrbracket$ assumes secure execution of e in that the attacker only affects e 's execution by affecting the inputs, which are assumed to be untrusted anyway — we later state an integrity property that validates this assumption of the enclave platform.

3.1 Proof Decomposition of SRE

A rational user will outsource the enclave only to a platform that provides a formal guarantee of SRE. To that end, we describe a method for formally verifying that an enclave platform provides SRE to any enclave program. We provide machine-checked proofs of SRE for the TAP in § 4, and show how this applies to models of Intel SGX and MIT Sanctum in § 5. The key idea is to decompose the SRE property into the following set of properties.

- **Secure Measurement:** The platform must *measure* the enclave program to allow the user to detect any changes to the program prior to execution, i.e., the user must be able to verify that the platform is running an unmodified e .
- **Integrity:** The enclave program's execution cannot be affected by a privileged software attacker beyond providing inputs, i.e., the sequence of inputs uniquely determines the enclave's execution trace, and that trace must be allowed by the enclave's semantics $\llbracket e \rrbracket$.
- **Confidentiality:** A privileged software attacker cannot distinguish between the executions of two enclaves, besides what is already revealed by \textit{obs} .

² $\langle \sigma'_0, \sigma'_1, \dots, \sigma'_m \rangle = \text{filter}(\lambda \sigma. \textit{curr}(\sigma) = e, \pi)$.

3.1.1 Secure Measurement. During launch, the platform computes a hash of the enclave's initial contents (*init*) along with relevant configuration bits (*config*). The hash-based measurement acts as a unique identity for the enclave, which follows directly from the collision resistance assumption of the cryptographic hash function, and therefore finds use in authenticating the enclave. Any deviation from the desired enclave program will be detected when the enclave sends an attested statement to the user — we assume that attested statements are produced using a quoting scheme that is unforgeable under chosen message attacks (UF-CMA); we do not model the cryptography of this scheme, and refer the reader to [58] for a formal treatment of this subject. The secure measurement property states that any two enclaves with the same measurement must also have the same semantics: they must produce equivalent execution traces for equivalent input sequences.

Let $\mu(e)$ be the measurement of enclave e , computed when launching the enclave. The operation must be such that two enclaves with the same measurement have identical initial states.

$$\begin{aligned} \forall \sigma_1, \sigma_2. \text{init}_{e_1}(E_{e_1}(\sigma_1)) \wedge \text{init}_{e_2}(E_{e_2}(\sigma_2)) \Rightarrow \\ \mu(e_1) = \mu(e_2) \iff E_{e_1}(\sigma_1) = E_{e_2}(\sigma_2) \end{aligned} \quad (2)$$

Next we need to ensure that if two enclaves e_1 and e_2 have the same state, then they produce equivalent execution traces for equivalent input sequences. This is the determinism property we assumed (while defining $\llbracket e \rrbracket$ in § 2.1) of the enclave platform, so we must prove it here.

$$\begin{aligned} \forall \pi_1, \pi_2. \quad (3) \\ (E_{e_1}(\pi_1[0]) = E_{e_2}(\pi_2[0])) \quad \wedge \\ \forall i. (\text{curr}(\pi_1[i]) = e_1) \iff (\text{curr}(\pi_2[i]) = e_2) \quad \wedge \\ \forall i. (\text{curr}(\pi_1[i]) = e_1) \implies I_{e_1}(\pi_1[i]) = I_{e_2}(\pi_2[i]) \quad \implies \\ (\forall i. E_{e_1}(\pi_1[i]) = E_{e_2}(\pi_2[i]) \wedge O_{e_1}(\pi_1[i]) = O_{e_2}(\pi_2[i])) \end{aligned}$$

Equation 3 states that if: (i) the two traces π_1 and π_2 start with the same initial state for enclaves e_1 and e_2 , (ii) if π_1 and π_2 enter and exit the enclaves in lockstep (i.e., the two enclaves execute for the same number of steps), (iii) if the input sequence to the two enclaves is the same, then the two enclaves execute identically in both traces: they have the same sequence of state and output values.

3.1.2 Integrity. The integrity guarantee ensures that the execution of the enclave in the presence of attacker operations is identical to the execution of the program without the attacker's operations. In other words, the attacker only impacts an enclave's execution by controlling the sequence of inputs — all other operations, such as controlling I/O peripherals and executing supervisor-mode instructions, has no effect on the enclave's execution. Any two traces (of the same enclave program) that start with equivalent enclave states and have the same input sequence will produce the same sequence of enclave states and outputs, even though the attacker's operations may differ in the two traces.

$$\begin{aligned} \forall \pi_1, \pi_2. \quad (4) \\ (E_e(\pi_1[0]) = E_e(\pi_2[0])) \quad \wedge \\ \forall i. (\text{curr}(\pi_1[i]) = e) \iff (\text{curr}(\pi_2[i]) = e) \quad \wedge \\ \forall i. (\text{curr}(\pi_1[i]) = e) \implies I_e(\pi_1[i]) = I_e(\pi_2[i]) \quad \implies \\ (\forall i. E_e(\pi_1[i]) = E_e(\pi_2[i]) \wedge O_e(\pi_1[i]) = O_e(\pi_2[i])) \end{aligned}$$




Figure 2: Integrity property

Figure 2 shows the two traces from the integrity property. The adversary's steps are labelled A_1 and A_2 while the enclave's steps are labelled e . Assumptions are annotated in blue, and proof obligations are shown in red. The enclave's inputs are assumed to be the same in both traces; this is shown by the \approx_I symbol. The initial state of the two enclaves is assumed to be the same and this is shown by the \approx_E symbol. The adversary's actions are defined by the *tamper* function (see § 2.2.1), and these actions may differ between the two traces. The integrity proof must show that the enclave's state and outputs do not differ despite this: $\forall i. E_e(\pi_1[i]) = E_e(\pi_2[i]) \wedge O_e(\pi_1[i]) = O_e(\pi_2[i])$. These proof obligations are denoted by the red \approx_E and \approx_O symbols. We assume that the adversary executes for the same number of steps in both traces. This does not restrict the adversary's power as any attack that requires the adversary to execute for a different number of steps in the two traces can be simulated in our model by padding the adversary's shorter trace with the appropriate number of "no-ops."

The theorem states that, given the above assumptions, enclave state and outputs are identical in the two traces at every step: $\forall i. E_e(\pi_1[i]) = E_e(\pi_2[i]) \wedge O_e(\pi_1[i]) = O_e(\pi_2[i])$.

3.1.3 Confidentiality. The enclave platform must ensure that the attacker does not observe the enclave's execution beyond what is allowed by the observation function *obs*, which must include the *init* and *config* components of the enclave's description, outputs to non-enclave memory, exit events from enclave mode to untrusted code, and other side channels leakage as permitted by the observation function *obs*. The privileged software attacker may use any combination of machine instructions to perform an attack, and the attacker should observe the same results from executing these instructions for any pair of enclaves that produce the same observation via the function *obs*. In other words, all enclave traces with the same observations, but possibly different enclave states, must be indistinguishable to the attacker. If a platform allows the attacker to observe different values in two such traces, we say that such a platform does not provide confidentiality.

Formally, to prove that an attacker learns no other information beyond *obs*, we must prove that for any two traces that have equivalent attacker operations and equivalent observations of the enclave

execution, but possibly different enclave private states and executions, the attacker's execution (its sequence of states) is identical.

$$\begin{aligned}
 & \forall \pi_1, \pi_2. & (5) \\
 & (A_{e_1}(\pi_1[0]) = A_{e_2}(\pi_2[0])) & \wedge \\
 & \forall i. \text{curr}(\pi_1[i]) = \text{curr}(\pi_2[i]) \wedge I^P(\pi_1[i]) = I^P(\pi_2[i]) & \wedge \\
 & \forall i. \text{curr}(\pi_1[i]) = e \implies \text{obs}_{e_1}(\pi_1[i+1]) = \text{obs}_{e_2}(\pi_2[i+1]) \implies & \\
 & (\forall i. A_{e_1}(\pi_1[i]) = A_{e_2}(\pi_2[i])) &
 \end{aligned}$$




Figure 3: Confidentiality property

Figure 3 depicts the confidentiality property. As in Figure 2, the attacker's steps are labelled A while the enclave's steps are labelled e_1 and e_2 . The two traces start off in equivalent states (shown by the blue \approx_A) but diverge (at state σ_i) because the two enclaves may perform different computation. The enclave's adversary-visible observations are assumed to be the same in both traces when the enclave is executing (shown by the blue obs_{e_1}). Adversary non-determinism, $I^P(\sigma)$, is assumed to be the same in both traces and this ensures that adversary actions (defined by the *tamper* function from § 2.2.1) are the same in both traces. The theorem states that adversary state is identical at every step: $\forall i. A_{e_1}(\pi_1[i]) = A_{e_2}(\pi_2[i])$ and is illustrated by the red \approx_A . The theorem implies that adversary state is a deterministic function of only adversary actions and initial state. In particular, adversary state does not depend on enclave state. This shows that the attacker has not learned any additional information beyond the observation function.

3.2 Soundness of SRE Decomposition

THEOREM 3.2. *An enclave platform that satisfies secure measurement, integrity, and confidentiality property for any enclave program also satisfies secure remote execution.*

Proof Sketch: Suppose that the user sends an arbitrary enclave e to a remote server for execution, and the platform launches enclave e_r some time later — because e is sent over an untrusted channel, e may or may not equal e_r . If the user finds $\mu(e_r) \neq \mu(e)$, then the platform has no obligations to execute a trace from $\llbracket e \rrbracket$. Otherwise, if $\mu(e_r) = \mu(e)$, we have that $\llbracket e_r \rrbracket = \llbracket e \rrbracket$ thanks to the measurement and integrity properties. So, the two programs have identical runtime behaviors, which is a prerequisite for SRE. Finally, confidentiality implies that the attacker's observation is restricted to *obs*.

3.3 Application of Secure Remote Execution

The measurement and integrity properties guarantee that the remote platform executes a trace from $\llbracket e \rrbracket$, while the confidentiality

property ensures that the attacker does not learn more information than what the enclave wishes to reveal. Together, these three properties imply SRE.

SRE is useful for building practical trusted applications. Consider the problem of executing a batch job securely in the cloud. The user sends an enclave program, which implements a function on sensitive data, to an enclave platform in the cloud. The protocol includes the following steps:

- (1) The user sends an enclave program e to the cloud provider, which launches the program on an enclave platform.
- (2) The user and enclave establish an authenticated TLS channel via an ephemeral Diffie-Hellman (D-H) exchange.
 - (a) User sends her public parameter g^x to the enclave, where x is a randomly generated fresh value.
 - (b) Enclave sends its public parameter $\text{attest}(g^y)$ to the user, in the form of an attested statement, thus guaranteeing that a genuine enclave platform launched the expected enclave e .
 - (c) User and enclave compute a shared master secret g^{xy} , and derive symmetric session keys s_{ue} and s_{eu} , a key for each direction.
- (3) The user now sends encrypted input to enclave using this shared secret: $\{in\}_{s_{ue}}$.
- (4) The enclave decrypts its input, performs the computation and returns the encrypted result to the user: $\{out\}_{s_{eu}}$.

Consider the following security property: the attacker neither learns secret input $\{in\}$ nor the secret output $\{out\}$. To that end, the user 1) develops an enclave program that only accepts encrypted inputs and sends encrypted outputs, and 2) specifies an observation function (*obs*) where a privileged software adversary is only allowed to view the enclave's outputs to non-enclave memory — this is acceptable because e encrypts its outputs.

The measurement guarantees that the user will only establish a channel with the expected enclave on a genuine enclave platform. Integrity ensures that the platform will execute a trace from $\llbracket e \rrbracket$, thus respecting e 's semantics. The platform may choose to not launch e or prematurely terminate e , but such executions will not generate $\{out\}_{s_{eu}}$ and hence can be trivially detected by the user. Integrity also ensures that the platform does not rollback the contents of enclave's memory while it is alive (i.e., not destroyed) as such attacks will cause the enclave's execution to proceed differently from $\llbracket e \rrbracket$, and SRE guarantees $\llbracket e \rrbracket$. SRE does not require the platform to defend against rollback attacks on persistent storage. This protection is not needed for the batch service because the enclave does not update $\{in\}_{s_{ue}}$, and any tampering to $\{in\}_{s_{ue}}$ will fail the cryptographic integrity checks. Finally, confidentiality ensures that the enclave platform only reveals the *obs* function of enclave's execution to the software attacker, which only includes the encrypted outputs. We now have our end-to-end security property.

Should the enclave require state beyond enclave's memory to perform the job, it would require integrity, freshness, and confidentiality for non-enclave state, which is not covered by SRE. The enclave can implement cryptographic protections (e.g., Merkle tree) and techniques for state continuity [57] to address this concern.

State Var.	Type	Description
pc	VA	The program counter.
regs	$\mathbb{N} \rightarrow W$	Architectural registers: map from natural numbers to words.
mem	$PA \rightarrow W$	The memory: a map from physical addresses to words.
addr_map	$VA \rightarrow (ACL \times PA)$	Map from virtual addresses to permissions and physical addresses for current process.
cache	$(Set \times Way) \rightarrow (\mathbb{B} \times Tag)$	Cache: map from a tuple of cache sets and ways to valid bits and cache tags.
current_eid	\mathcal{E}_{id}	Current enclave. $current_eid = OS$ means that no enclave is being executed.
owner	$PA \rightarrow \mathcal{E}_{id}$	Map from physical address to the enclave address is allocated to.
enc_metadata	$\mathcal{E}_{id} \rightarrow \mathcal{E}_M$	Map from enclave ids to metadata record type (\mathcal{E}_M).
os_metadata	\mathcal{E}_M	Record that stores a checkpoint of privileged software state.

Table 1: Description of TAP State Variables

4 THE TRUSTED ABSTRACT PLATFORM

The trusted abstract platform (TAP) consists of a processor with a program counter, general purpose registers, virtual address translation and a set of primitives to support enclave execution. In this section, we first introduce a formal model of the TAP. We present a range of adversary models with varying capabilities. We then present a set of machine-checked proofs showing that the TAP satisfies the properties required for secure remote execution: (i) secure measurement, (ii) integrity and (iii) confidentiality.

4.1 TAP Model

Recall that the TAP is modeled as a finite state transition system: $TAP = (\Sigma, \rightsquigarrow, init)$. Σ is the set of states, \rightsquigarrow is the transition relation and $init \in \Sigma$ is the initial state.

4.1.1 TAP State Variables. The states Σ of the TAP are a defined as a valuation of the state variables *Vars*. These variables are described in Table 1. *pc*, *regs*, *mem* have their usual meanings. *addr_map* maps individual virtual addresses to physical addresses and permissions. This is unlike a typical processor which uses a page table to map virtual *page numbers* to physical *page numbers*. The TAP is an abstraction and must abstract a diverse set of architectures with different page table structures and page sizes. Therefore, it maps each virtual address to a physical address.

The TAP models a cache to show that confidentiality is preserved in the presence of a software adversary attempting cache attacks. The cache model assumes a physically-indexed and physically-tagged cache with symbolic parameters: an unbounded number of sets, unbounded associativity, and an arbitrary deterministic replacement policy. The TAP cache model leaves the mapping of physical addresses to cache sets and mapping of physical addresses to cache tags uninterpreted. In other words, the TAP formalism applies to a broad space of set associative caches, and only requires that the cache set and cache tag for each memory access be deterministic functions of the physical address. The exact functions and parameters are specified by implementations (refinements) of TAP, which are models of SGX and Sanctum in this paper.

The variable *current_eid* tracks the enclave currently being executed. *owner* maps each physical address to the enclave which exclusively “owns” it. If $owner[p] = e$, only enclave *e* can access (fetch/read/write) this word of memory. We abuse notation and use *e* to refer to both the “enclave id,” a unique integer assigned

by the platform to an enclave, as well the enclave itself. Attempts to access physical address *p* by all other enclaves and privileged software are blocked. $owner[p] = OS$ means that address *p* is not allocated to any enclave. *owner* corresponds to the EPCM in SGX and the DRAM bitmap in Sanctum. It is the primary mechanism to enforce isolation of enclave’s private memory. *enc_metadata* stores metadata about each initialized enclave.

State var.	Description
entrypoint	Enclave entrypoint.
addr_map	Virtual to physical mapping/permissions.
excl_vaddr	Set of private virtual addresses.
measurement	Enclave measurement.
pc	Saved PC (in case of interrupt).
regs	Saved registers (in case of interrupt).
paused	Flag set only when enclave is interrupted.

Table 2: Fields of the enc_metadata record

Enclave Metadata: Table 2 lists various fields within the enclave metadata record. It stores the entrypoint to the enclave, its virtual to physical mappings and what set of virtual addresses are private to the enclave. The *pc* and *regs* fields are used to checkpoint enclave state when it is interrupted. The *paused* flag is set to true only when an enclave is interrupted and ensures that enclaves cannot be tricked into resuming execution from invalid state.

Privileged Software Metadata: The *os_metadata* record contains three fields: *pc*, *regs*, and *addr_map*. The *pc* and *regs* fields store a checkpoint of privileged software state. These are initialized when entering enclave state and restored when the enclave exits. The *addr_map* field is the privileged software’s virtual to physical address mapping and associated permissions.

4.1.2 TAP Operations. Table 3 describes the operations supported by the TAP. *fetch*, *load*, *store* work as usual. The platform guarantees that memory owned by enclave *e* is not accessible to other enclaves or privileged software. Each of these operations update cache state, set the access bit in *addr_map*, and return whether the operation was a hit or a miss in the cache.

The virtual to physical mappings of both enclave and privileged software are controlled using *get_addr_map* and *set_addr_map*.

Operation	Description
fetch(v) load(v) store(v)	Fetch/read/write from/to virtual address v . Fail if v is not executable/readable/writable respectively according to the <code>addr_map</code> or if <code>owner[addr_map[v].PA] ≠ current_eid</code> .
get_addr_map(e, v) set_addr_map($e, v, p, perm$)	Get/set virtual to physical mapping and associated permissions for virtual address v .
launch(e, m, x_v, x_p, t) destroy(e) enter(e), resume(e) exit(), pause() attest(d)	Initialize enclave e by allocating <code>enc_metadata[e]</code> . Set <code>mem[p]</code> to 0 for each p such that <code>owner[p] = e</code> . Deallocate enclave <code>enc_metadata[e]</code> . enter enters enclave e at entrypoint, while resume starts execution of e from the last saved checkpoint. Exit enclave. pause also saves a checkpoint of pc and regs and sets <code>enc_metadata[e].paused = true</code> . Return hardware-signed message with operand d and enclave measurement e : $\{d \parallel \mu(e)\}_{SK_p}$.

Table 3: Description of TAP Operations

$(current_eid = OS \wedge e \notin enc_metadata) \wedge$
 $executable(m[t]) \wedge t \in x_v \wedge$
 $\forall p. p \in x_p \implies owner[p] = OS \wedge$
 $\forall v. v \in x_v \implies (valid(m[v]) \implies m[v]_{PA} \in x_p) \wedge$
 $\forall v_1, v_2. (v_1 \in x_v \wedge v_2 \in x_v) \implies (m[v_1]_{PA} \neq m[v_2]_{PA})$
 $\iff (launch_status = success)$

Note: $m[v]_{PA}$ refers to physical address that virtual address v points to under the mapping m .

Figure 4: Conditions for the success of launch

As enclave's memory access pattern can leak via observation of the access/present bits in `addr_map`, `get_addr_map(e, v)` must fail (on a secure TAP) for virtual addresses in the set `enc_metadata[e].evrange` when called from outside the enclave. However, SGX does permit privileged software to access an enclave's private page tables. We introduce a "setting" in the TAP, called `priv_mappings`, and this insecure behavior is allowed when `priv_mappings = false`. **Enclave Creation:** The `launch(e, m, x_v, x_p, t)` operation is used to create an enclave. The enclave's virtual to physical address mapping and associated permissions are specified by m . x_v is the set of enclave-private *virtual* addresses (*evrange*). It corresponds to the base address and size arguments passed to `ECREATE` in SGX and `create_enclave` in Sanctum. x_p is the set of *physical* addresses allocated to the enclave and its entrypoint is the virtual address t .

The launch operation only succeeds if enclave id e does not already exist, if the entrypoint is mapped to an enclave-private executable address, if every virtual address in x_v that is accessible to the enclave points to a physical address in x_p , and if there is no aliasing among the addresses in x_v . A precise statement of the conditions that result in a successful launch shown in Figure 4. These conditions have subtle interactions with the requirements for SRE. For example, if virtual addresses within x_v are allowed to alias, an adversary can construct two enclaves which have the same measurement but different semantics. The potential for such attacks emphasizes the need for formal modeling and verification.

Enclave Destruction: An enclave is deallocated using `destroy`, which zeroes out the enclave's memory so that its private state is not leaked when the privileged software reclaims the memory. This is necessary for confidentiality because untrusted privileged software can destroy an enclave at any time.

Enclave Entry/Exit: The operation `enter(e)` enters enclave e by setting the pc to its entrypoint and `current_eid` to e . `resume` continues execution from a checkpoint saved by `pause`. The enclave may transfer control back to the caller via `exit`. `pause` models forced interruption by privileged software (e.g., device interrupt).

Attestation: The attestation operator provided by the TAP ensures that the user is communicating with a bona fide enclave. The `attest` operation can only be invoked from within the enclave and may be used by the enclave to establish its identity as part of an authentication protocol. `attest` returns a hardware-signed cryptographic digest of data d and a measurement: $d \parallel \mu(e) \parallel \{d \parallel \mu(e)\}_{SK_p}$. The signature uses the processor's secret key SK_p , whose corresponding public key is signed by the trusted platform manufacturer.

4.2 The TAP Adversary

As described in § 2.2, the TAP adversary model is based on a privileged software attacker and consists of two components: (i) a tampering relation that describes how the adversary can change the platform's state, and (ii) an observation function that describes what elements of machine state are visible to the adversary. The adversary model is parameterized and has three instantiations: Adversary *MCP*, adversary *MC* and adversary *M*, which are defined below. *MCP* is the most general adversary and models all capabilities of a privileged software attacker, while *MC* and *M* are restricted.

4.2.1 Adversary Tampering. The tamper relation for the TAP formalizes all possible actions that may be carried out by a software attacker. It serves two purposes. When reasoning about integrity, the tamper relation defines all operations that an adversary may carry out to interfere with enclave execution. When reasoning about enclave confidentiality, the tamper relation models computation performed by the adversary in order to learn an enclave's private state. The most general tamper relation corresponds to Adversary *MCP*, and is defined as the adversary carrying out an unbounded number of the following actions:

- (1) Unconstrained updates to pc and regs. These are modeled using the havoc statement commonly used in program verification: `havoc pc, regs`.
- (2) Loads and stores to memory with unconstrained address (va) and data ($data$) arguments.
 - $\langle op, hit_f \rangle \leftarrow \text{fetch}(va)$
 - $\langle regs[ri], hit_l \rangle \leftarrow \text{load}(va)$
 - $hit_s \leftarrow \text{store}(va, data)$
- (3) Modification of the adversary's view of memory by calling `set_addr_map` and `get_addr_map` with unconstrained arguments.
 - `set_addr_map($e, v, p, perm$)`
 - `regs[ri] \leftarrow get_addr_map(e, v)`
- (4) The invocation of enclave operations with unconstrained arguments.
 - Launch enclaves: `launch(e, m, x_v, x_p, t)`.
 - Destroy enclaves: `destroy(e)`.
 - Enter and resume enclaves: `enter(e)` and `resume(e)`.
 - Exit (`exit`) from and interrupt (pause) enclaves.

Any adversary program, including malicious operating systems, hypervisors, malicious enclaves and privileged malware, can be modeled using the above actions. Therefore, this adversary model allows us to reason about TAP security guarantees in presence of a general software adversary.

Restricted Adversaries: Adversary MC is restricted to computation based on memory values and cache state; it ignores the value returned by `get_addr_map`. Adversary M only computes using memory values; it ignores hit_f , hit_l and hit_m returned by `fetch`, `load` and `store`, respectively, in addition to the result of `get_addr_map`.

4.2.2 Adversary Observation Model. The observation function captures what state the user expects to be attacker-visible.

Adversary M : The observation function $obs_e^M(\sigma)$ is a partial map from physical addresses to words and allows the adversary to observe the contents of all memory locations not private to enclave e . It is defined as $\sigma(\text{mem}[p])$ when $\sigma(\text{owner}[p]) \neq e$ and \perp otherwise.

$$obs_e^M(\sigma) \doteq O_e(\sigma) \doteq \lambda p. \text{ITE}(\sigma(\text{owner}[p]) \neq e, \sigma(\text{mem}[p]), \perp)$$

Adversary MC : The observation function $obs_e^{MC}(\sigma)$ specifies that besides contents of memory locations that are not private to an enclave, the adversary can also observe whether these locations are cached. It is also a partial map from physical addresses to words and is defined to be the tuple $\langle \sigma(\text{mem}[p]), \text{cached}(\sigma, p) \rangle$ when $\sigma(\text{owner}[p]) \neq e$ and \perp otherwise. $\text{cached}(\sigma, p)$ is true iff physical address p stored in the cache in the machine state σ .

$$obs_e^{MC}(\sigma) \doteq \lambda p. \text{ITE}(\sigma(\text{owner}[p]) \neq e, \langle \sigma(\text{mem}[p]), \text{cached}(\sigma, p) \rangle, \perp)$$

Note that the adversary cannot directly observe whether an enclave's private memory locations are cached. However, unless cache sets are partitioned between the attacker and the enclave, cache attacks [78, 84] allow the adversary to *learn* this information.

Adversary MCP : The MCP adversary's capabilities are extended beyond MC to include observation of the virtual to physical mappings and associated access/permission bits for each virtual address.

$$obs_e^{MCP}(\sigma) \doteq \lambda \sigma. \langle obs_e^{MC}(\sigma), \lambda v. \sigma(\text{get_addr_map}(e, v)) \rangle$$

The notation $\sigma(\text{get_addr_map}(e, v))$ refers to the result of evaluating `get_addr_map(e, v)` in the state σ .

4.2.3 Enclave and Adversary State and Inputs. Recall that the state of an enclave e is $E_e(\sigma)$, which is defined as the tuple $\langle E_{\text{vmem}}(e, \sigma), E_{\text{regs}}(e, \sigma), E_{\text{pc}}(e, \sigma), E_{\text{cfg}}(e, \sigma) \rangle$ if $e \in \text{enc_metadata}$ and \perp otherwise. The components of this tuple are as follows:

- $E_{\text{vmem}}(e, \sigma)$ is a partial map from virtual addresses to words. It is defined to be $\sigma(\text{mem}[\text{enc_metadata}[e].\text{addr_map}[v]_{\text{PA}}])$ if $v \in \sigma(\text{enc_metadata}[e].\text{evrange})$ and \perp otherwise. In other words, E_{vmem} refers to the content of each virtual memory address in the enclave's *evrange*.
- $E_{\text{regs}}(e, \sigma)$ is $\sigma(\text{regs})$ if $\text{curr}(\sigma) = e$ (when the enclave is executing), and $\sigma(\text{enc_metadata}[e].\text{regs})$ otherwise.
- $E_{\text{pc}}(e, \sigma)$ is $\sigma(\text{pc})$ if $\text{curr}(\sigma) = e$ (when the enclave is executing), and $\sigma(\text{enc_metadata}[e].\text{pc})$ otherwise.
- The tuple $E_{\text{cfg}}(e, \sigma)$ consists of the following elements:
 - (i) $\sigma(\text{enc_metadata}[e].\text{addr_map})$
 - (ii) $\sigma(\text{enc_metadata}[e].\text{entrypoint})$
 - (iii) $\sigma(\text{enc_metadata}[e].\text{evrange})$

Recall that the input to enclave e at state σ is $I_e(\sigma)$, which is the tuple $I_e(\sigma) \doteq \langle I_e^R(\sigma), I_e^U(\sigma) \rangle$. $I_e^R(\sigma)$ is the random number provided at the state σ . $I_e^U(\sigma)$ is a partial map from virtual address to words. It is $\sigma(\text{mem}[\text{enc_metadata}[e].\text{addr_map}[v]_{\text{PA}}])$ if each of these conditions hold: (i) enclave e is executing: $\text{curr}(\sigma) = e$, (ii) v is mapped to some physical address: $\sigma(\text{valid}(\text{enc_metadata}.\text{addr_map}[v]))$, and (iii) v is not private: $v \notin \sigma(\text{enc_metadata}[e].\text{evrange})$; it is \perp otherwise. In other words, $I_e^U(\sigma)$ refers to the contents of each virtual address *not* in the enclave's *evrange*.

The output $O_e(\sigma)$ contains memory values outside enclave's *evrange*. $O_e(\sigma)$ is defined identically to $I_e^U(\sigma)$. The adversary's state $A_e(\sigma)$ is modeled as the tuple $\langle S(\sigma), \hat{E}_e(\sigma) \rangle$. $S(\sigma) \doteq \langle S_{\text{vmem}}(\sigma), S_{\text{regs}}(\sigma), S_{\text{cfg}}(\sigma) \rangle$ denotes privileged software state. $S_{\text{vmem}}(\sigma) \doteq \lambda v. \sigma(\text{mem}[\text{os_metadata}.\text{addr_map}[v]_{\text{PA}}])$ is its view of memory. $S_{\text{regs}}(\sigma)$ denotes the privileged software's registers: $\sigma(\text{regs})$ when privileged software is executing, and $\sigma(\text{os_metadata}.\text{regs})$ otherwise. $S_{\text{cfg}}(\sigma)$ is the privileged software's virtual to physical mappings: $\sigma(\text{os_metadata}.\text{addr_map})$. $\hat{E}_e(\sigma)$ is the state of all the other enclaves in the system except for enclave e : $\hat{E}_e(\sigma) \doteq \lambda e'. \text{ITE}(e \neq e', E_{e'}(\sigma), \perp)$, where ITE denotes if-then-else.

4.3 Proof of Secure Remote Execution for TAP

We proved three machine-checked theorems that correspond to the requirements for secure remote execution as described in § 3.

TAP Integrity: We proved that the integrity result (Equation 4) holds for the TAP for all three adversaries: M , MC and MCP . This shows that these adversaries have no effect on enclave execution beyond providing inputs via non-enclave memory.

TAP Measurement: We showed that Equation 2 and Equation 3 are satisfied by the TAP. The proof for Equation 3 need not include adversarial operations because integrity ensures that an adversary cannot affect enclave's execution beyond providing inputs.

TAP Confidentiality: We showed three confidentiality results, each corresponding to the three TAP adversaries: M , MC , and MCP .

Confidentiality holds unconditionally for adversary M .

For adversary MC , let $pa2set : \text{PA} \rightarrow \text{Set}$ be the function that maps physical addresses to cache sets. This function is uninterpreted (abstract) in the TAP and will be defined by implementation.

We showed that confidentiality holds for adversary *MC* if Equation 6 is satisfied: a physical address belonging to an enclave never shares a cache set with a physical address outside the enclave.

$$\forall p_1, p_2, \sigma, e. \quad (6) \\ \sigma(\text{owner}[p_1] = e \wedge \text{owner}[p_2] \neq e) \implies (pa2set(p_1) \neq pa2set(p_2))$$

Finally, we showed that confidentiality holds for adversary *MCP* if Equation 6 is satisfied by the TAP implementation and the TAP configuration Boolean `priv_mappings` is true.

5 REFINEMENTS OF THE TAP

We prove that models of MIT Sanctum and Intel SGX are *refinements* of the TAP under certain adversarial parameters. Refinement shows that each operation, including all adversarial operations, on Sanctum and SGX processors can be mapped to an “equivalent” TAP action. The refinement proof implies SRE, which was proven on the TAP, also holds for our models of SGX and Sanctum.

5.1 Refinement Methodology

Let $Impl = \langle \Sigma_L, \sim_L, init_L \rangle$ be a transition system with states Σ_L , transition relation \sim_L and initial state $init_L$. We say that *Impl* refines *TAP*, or equivalently *TAP* simulates *Impl*, if there exists a *simulation relation* $R \subseteq (\Sigma_L \times \Sigma)$ with the following property:

$$\begin{aligned} & \left(\forall s_j \in \Sigma_L, s_k \in \Sigma, \sigma_j \in \Sigma. \right. \\ & (s_j, \sigma_j) \in R \wedge s_j \sim_L s_k \implies \\ & \left. ((s_k, \sigma_k) \in R \vee (\exists \sigma_k \in \Sigma. \sigma_j \sim \sigma_k \wedge (s_k, \sigma_k) \in R)) \right) \wedge \\ & (init, init_L) \in R \end{aligned} \quad (7)$$




Figure 5: Illustration of Stuttering Simulation

The condition states that for every pair of states (s_j, σ_j) (belonging to *Impl* and *TAP* respectively) that are related by *R*, if *Impl* steps from s_j to s_k , then either (i) the *TAP* takes no steps, and s_k is related to σ_j , or (ii) there exists a state σ_k of *TAP* such that σ_j steps to σ_k and (s_k, σ_k) are related according to *R*. In addition, the initial states of *Impl* and *TAP* must be related by *R*. This is illustrated in Figure 5. This corresponds to the notion of stuttering simulation [13], and we require stuttering because a single invocation of `launch` corresponds to several API calls in Sanctum and SGX.

Refinement states that every trace of *Impl* can be mapped using relation *R* to a trace of *TAP*; effectively this means that *TAP* has a superset of the behaviors of *Impl*. The security properties of the TAP

are *hyperproperties*, which in general, are not preserved by refinement [19]. However, the properties we consider are *2-safety* properties [77] that are variants of *observational determinism* [48, 62]. These are properties preserved by refinement. Intuitively, refinement asserts that any behavior observable on *Impl* has an equivalent behavior on *TAP* via *R*, so any attacker advantage available on *Impl* but not on *TAP* precludes refinement. The SRE properties proven on the TAP therefore also hold for SGX and Sanctum, as we show that these architectures (conditionally) refine the TAP.

5.2 MIT Sanctum Processor

Sanctum [21] is an open-source enclave platform that provides strong confidentiality guarantees against privileged software attackers. In addition to protecting enclave memory from direct observation and tampering, Sanctum protects against software attackers that seek to observe an enclave’s memory access patterns.

5.2.1 Sanctum Overview. Sanctum implements enclaves via a combination of hardware extensions to RISC-V [5] and trusted software at the highest privilege level, called the *security monitor*. **Sanctum Hardware:** Sanctum minimally extends Rocket Chip [4, 5], an open source reference implementation of the RISC-V ISA [81, 82]. Specifically, Sanctum hardware isolates physical addresses by dividing system memory (DRAM) into *regions*, which use disjoint last level cache sets, and allocating each region to an enclave exclusively. Since enclaves have exclusive control over one or more DRAM regions, there is no leakage of private memory access patterns through the cache. An adversary cannot create a TLB entry that maps its virtual address to an enclave’s private cache set.

Sanctum Monitor: The bulk of Sanctum’s logic is implemented in a trusted *security monitor*. The monitor exclusively operates in RISC-V’s *machine mode*, the highest privilege-level implemented by the processor, and solely able to bypass virtual address translation. Monitor data structures maintain enclave and DRAM region state. The monitor configures the Sanctum hardware to enforce low-level invariants that comprise enclave access control policies. For example, the monitor places an enclave’s page tables within that enclave’s DRAM region, preventing the OS from monitoring an enclave’s page table metadata to infer memory access patterns. The monitor exposes an API for enclave operations, including measurement. A trusted bootloader bootstraps the system, loads the monitor and creates a chain of certificates authenticating the Sanctum chip and the loaded security monitor.

5.2.2 Sanctum Model. Our Sanctum model combines the Sanctum hardware and the reference security monitor, and includes hardware registers, hardware operations, monitor data structures and the monitor API. The hardware registers include the DRAM bitmap which tracks ownership of DRAM regions, page table base pointers, and special regions of memory allocated to the monitor and for direct memory access (DMA). Hardware operations modeled include page tables and address translation, and memory instruction fetch, loads and stores. All monitor APIs are modeled.

5.2.3 Sanctum Model Refines TAP. The Sanctum refinement proof is expressed in three parts.

1. Concrete MMU refines Abstract MMU: We constructed an abstract model of a memory management unit (MMU). The abstract

MMU's address translation is similar to the TAP; it is a single-level map from virtual page numbers to permissions and physical page numbers. In contrast, the concrete Sanctum MMU models a multi-level page table walk in memory. We showed that the concrete Sanctum MMU model refines the abstract MMU. We then used the abstract MMU in modeling the Sanctum Monitor. This simplifies the simulation relation between Sanctum and TAP.

2. Sanctum Simulates TAP: We showed that every Sanctum state and every Sanctum operation, which includes both enclave and adversary operations, can be mapped to a corresponding TAP operation. For this, (i) we constructed a simulation relation between Sanctum states and corresponding TAP states, and (ii) we constructed a *Skolem function* [72] mapping each Sanctum operation to the corresponding TAP operation. We proved that for every pair of states in the simulation relation, every Sanctum operation can be mapped by the Skolem function to a corresponding TAP operation such that the resultant states are also in the simulation relation.

3. Proof of Cache Partitioning: The Sanctum model instantiates the function *pa2set* which maps physical addresses to cache sets. We showed that the Sanctum API's *init_enclave* operation and the definition of *pa2set* together ensure that all Sanctum enclaves' cache sets are partitioned, i.e. Equation 6 is satisfied.

5.3 Intel SGX

Intel Software Guard Extensions extend the Intel architecture to enable execution of enclave programs.

5.3.1 SGX Overview. Unlike Sanctum, enclaves in SGX are implemented in hardware (microcode), which provides an instruction set extension [33] to create enclaves (*ecreate*), enter enclaves (*eenter*), generate attested statements (*ereport*), etc. The SGX processor dedicates a contiguous region of physical memory (called the enclave page cache, or EPC), exclusively usable for storing enclave pages. While this provides confidentiality of enclave's memory, SGX does not protect several side channel leaks such as cache timing or page access patterns. The last level cache is shared between the enclave and the adversary, and adversary memory can map to the same cache set as the enclave's private address, allowing the attacker to perform cache attacks [12, 49, 78, 84]. SGX also allows the OS to examine the page tables that control the enclave's private memory, enabling the OS to read the accessed and dirty bits, thus learning the enclave's memory access pattern at the page-level granularity. The OS also gets notified on page fault exceptions (as part of demand paging), and this is another channel to learn the enclave's page-level memory access patterns [69, 83].

5.3.2 SGX Model. Similar to Sanctum, we use a formal model of the SGX instruction set, following the approach of Moat [71], which models the SGX platform at the level of abstraction presented in the Intel Programmer Reference manual [33]. The model contains ISA-level semantics of the SGX instructions. It includes hardware state such as enclave page cache metadata (EPCM), page tables, private core registers, etc. The model elides details such as attestation, hardware encryption of DRAM pages, and the cryptographic protections (encryption, integrity, and freshness) of demand paging and instead assumes several axioms about these features.

5.3.3 SGX Model Refines TAP. We attempted to prove that SGX refines TAP, in that all SGX traces can be mapped to TAP traces. However, we cannot prove SGX refinement unconditionally. We show that refinement holds only when *priv_mappings* = *false* (see Sec. 4.1.2). This is because SGX implements a mechanism for the attacker (OS) to view page table entries, which contains the accessed and dirty bits. As TAP confidentiality for Adversary *MCP* only holds when *priv_mappings* = *true*, SGX is not secure against *MCP*. Furthermore, the lack of cache partitioning also prevents us from showing that Equation 6 holds, so SGX does not refine TAP instantiated with Adversary *MC*. We *are* able to prove refinement of TAP by SGX for the restricted adversary *M*. This shows SGX provides similar guarantees to Sanctum, except for leakage through the cache and page table side-channels.

6 VERIFICATION RESULTS

This section discusses our models and machine-checked proofs. Our models of the TAP, Intel SGX and MIT Sanctum are constructed using the BoogiePL [7, 24] intermediate verification language. BoogiePL programs can be annotated with assertions, loop invariants and pre-/post-conditions for procedures. The validity of these annotations are checked using the Boogie Verification Condition Generator [7], which in turn uses the Z3 SMT solver [23].

Procedure pre-/post-conditions enable modular verification using Boogie. For example, we specify the behavior of each TAP operation using pre- and post-conditions and verify that the implementation of these procedures satisfies these post-conditions. Once the post-conditions are verified, the TAP's proofs of integrity, confidentiality and secure measurement can soundly ignore the implementation and only reason about the pre- and post-conditions of these operations.

6.1 BoogiePL Model Construction

Description	Size				Verif. Time (s)
	#pr	#fn	#an	#ln	
TAP	22	49	204	1752	5
Integrity	12	13	145	985	26
Measurement	6	3	100	800	6
Confidentiality	8	-	200	1388	194
MMU Model	9	13	68	739	7
MMU Refinement	3	2	38	216	8
Sanctum	23	321	44	780	1
Sanctum Refinement	12	3	94	548	11
SGX	36	113	4	1526	-
SGX Refinement	10	1	38	351	2
Total	141	518	935	9085	260

Table 4: BoogiePL Models and Verification Results

Table 4 shows the approximate size of each of our models. #pr, #fn, #an and #ln refers to the number of procedures, functions, annotations and lines of code, respectively. Annotations refer to the number of loop invariants, assertions, assumptions, pre- and

post-conditions that we manually specify. While Boogie can infer discharge some assertions automatically, we found that we had to manually specify 935 annotations before it accepted our proofs.

The rows TAP, MMU Model, Sanctum and SGX correspond to models of the functionality of the TAP, the Sanctum MMU, Sanctum, and SGX, respectively. The other rows correspond to our proofs of SRE and refinement. In total, the models are about 4800 lines of code while the proofs form the remaining \approx 4300 lines of code. A significant part of the effort in developing the proofs was finding the correct inductive invariants to help Boogie prove the properties.

BoogiePL can only verify *safety* properties. But many of our theorems involve *hyperproperties* [19]. We used the self-composition construction [8, 77] to convert these into safety properties. BoogiePL is also incapable of verifying properties involving alternating nested quantifiers, for example, $\forall x. \exists y. \phi(x, y)$. We *Skolemized* [72] such properties to take the form: $\forall x. \forall y. (y = f(x)) \implies \phi(x, y)$; $f(x)$ is called a *Skolem function* and must be manually specified.

6.2 Verification Results

Table 4 lists the total time taken by Boogie/Z3 to check validity of all the manually specified annotations — by verifying the annotations, we omit them from the trusted computing base, which only includes the Boogie/Z3 theorem prover. The verification times for the TAP, MMU Model and Sanctum rows is for proving that procedures in these models satisfy their post conditions, which specify behavior and system invariants. The verification times for the remaining rows is the time taken to prove the SRE properties and refinement. The total computation time in checking validity of the proofs once the correct annotations are specified is only a few minutes.

7 DISCUSSION

We argue that TAP-based verification of enclave platforms is more than just a set of proofs of correctness for our models of Intel SGX and MIT Sanctum. The TAP serves as a specification of primitives for enclave execution, and is designed to be extensible towards additional features (e.g., non-volatile monotonic counters, demand paging) and additional guarantees against sophisticated attackers (e.g., timing attacks). It provides an extensible verification methodology that enables rigorous reasoning about the security properties of *future* enclave platforms. It is also an extensible model for reasoning about enclave programs.

7.1 Implications for Enclave Platforms

The TAP can be beneficial to implementers of enclave platforms in the following ways. First, the TAP can be used as a top-down specification for what operations an enclave platform must support. Implementers can use refinement checks to ensure that the TAP's security guarantees apply to the implementation as well. A important implementation challenge is that security is not compositional: addition of operations to a secure platform can make it insecure. The insecurity of SGX to Adversary MCP [69, 83] stems from demand paging, using which the OS observes the state of enclave's page tables; this feature is *not* present in the current implementation of Sanctum. Suppose that the next version of Sanctum supports demand paging through the use of oblivious RAM to maintain TAP's confidentiality guarantee. Reasoning about these more complex

enclave platforms is infeasible without a verification methodology and TAP-like specification of the platform's primitives.

Second, the TAP can also be used bottom-up, as this paper does, to reason about the security properties of existing platforms. Such analysis exposes differences between various platforms, e.g., leakage through the cache and page tables for SGX. The refinement methodology described in this paper can be used to show that implementations of enclave platforms (e.g., the Sanctum implementation) are refinements of corresponding models (e.g., Sanctum model). Since we have shown that the Sanctum model is a refinement of the TAP, this proof would mean that the security properties of the TAP also hold for the Sanctum *implementation*.

7.2 Implications for Enclave Software

While techniques for verifying security properties of enclave programs [70, 71] have been developed, they rely on models of the underlying enclave platform. We argue that such reasoning is better done using TAP's clean abstraction, which is simpler than the instruction-level model of SGX and API-level model of Sanctum. This simplification makes automated verification more scalable, yet retains soundness because of the refinement checks. It is also a step towards enabling portability among different enclave platforms.

Most importantly, this paper provides a common language for research into security properties of enclaves. While side-channel defenses have been proposed for SGX enclaves [55, 68, 69], the lack of formalization of the enclave's execution, attacker's operations and observations, and the desired confidentiality property makes it infeasible to systematically compare two defenses.

7.3 Limitations

7.3.1 Limitations of the TAP. The TAP does not model the full space of systems for which enclaves are relevant. Instead, it focuses on enclave platforms that provide memory-based isolation. The TAP's current adversary model assumes platform memory (DRAM) is trusted: memory can only be modified via software. Extending the adversary model beyond this to include, for example, memory probing/DMA attacks would require augmenting the TAP's memory operations (fetch/load/store) with encryption and integrity protection. The TAP also does not support demand paging as a naïve implementation is vulnerable to the pigeonhole attack [69, 83]. One path towards confidentiality-preserving demand paging in the TAP would be to use an oblivious-RAM protocol.

The TAP model and proofs are currently limited to concurrent execution on a single-threaded single-core processor. A fully general proof showing that the TAP ensures SRE on multicore systems would require parameterized verification of these properties for an unbounded number of cores, using for example, the CMP method [18]. We would also need to show linearizability of operations in the Sanctum/SGX models [31]. In this context, linearizability means that each operation appears to be atomic and there exists an ordering of operations across all cores that is consistent with their sequential orders of completion in individual cores. Linearizability of a model of SGX was shown by Leslie-Hurd et al. [39]. Extending the TAP to model security of simultaneously multithreaded (SMT) processors would require modeling many more side channels (e.g., branch predictors, shared resources such

as instruction queues and function units) and the development of architectural mechanisms that prevent leakage through these side channels.

Finally, the TAP does not model the cryptography in the attestation operation. Instead, we model the hash function using uninterpreted functions and assume the properties of collision resistance and second pre-image resistance.

7.3.2 Limitations of the Sanctum/SGX Models. Our SGX model does not include demand paging, memory encryption or memory integrity protection and assumes correctness of SGX measurement. Unlike the TAP, the Sanctum implementation's measurement is not computed on a single snapshot of state, but is instead updated incrementally during enclave creation. A proof that this incremental measurement is identical to the TAP's measurement is ongoing. Like the TAP, our Sanctum model also uses uninterpreted functions to model the cryptography in the measurement operation.

8 RELATED WORK

Secure Processors: There have been several commercial deployments of secure processors. ARM TrustZone [2] implements a secure and normal mode of execution to effectively create a single privileged enclave in an isolated address space. TPM+TXT [29] enables attestation of the platform's state, but includes all privileged software layers in the trusted computing base. Most recently, Intel SGX [3, 20, 32, 47] implements unprivileged enclaves protecting the integrity of memory and enclave state against software adversaries and certain physical access attacks. Academic work seeking to improve the security of aspects of conventional processors is also abundant [25, 43, 44].

Several clean-slate academic projects have been seeking to build a trusted system. The XOM [42] architecture introduced the concept of isolated software containers managed by untrusted host software, and employed encryption and HMAC to protect DRAM. Aegis [76] showed how a security kernel could measure and sign enclaves, and employed counter-mode memory encryption and a Merkle tree to guarantee freshness of data in DRAM. Bastion [16] encrypts and HMACs DRAM and employs a trusted (and authenticated as part of remote attestation) hypervisor which is invoked at each TLB miss to check address translation against a policy. Fides [74] uses a small trusted hypervisor to provide security guarantees for so-called protected modules; protected modules offer very similar programming abstractions and security guarantees as enclaves. Sancus [54] builds upon Fides, and proposes hardware extensions that ensure security of protected modules in the context of resource-constrained embedded processors. Ascend [27] and Phantom [46] ensure privacy and integrity of all CPU memory accesses through a hardware ORAM primitive.

Attacks on Secure Processors: Secure systems often expose complex threat models in practice, leading to vulnerabilities in the application layers. Side channel observations, such as attacks observing cache timing, are known to compromise cryptographic keys used by numerous cryptosystems [11, 14, 15, 37]. Attacks exploiting other shared resources exist as well, such as those observing a core's branch history buffers [38]. These attacks are viable at any privilege, separated by arbitrary protection boundaries [34, 45, 56]. Similar

attacks apply on trusted hardware, as shown on SGX with attacks observing shared caches [12, 49, 66], and shared page tables [69, 83]. **Formal Models/Verification of Enclave Platforms:** A formal cryptographic model of SGX's anonymous attestation scheme is presented in [58]. Barbosa et al. [6] present a formal analysis of cryptographic protocols for secure outsourced computation using enclave platforms. In contrast to these efforts, which reason about cryptographic protocols in the context of enclaves, our work formalizes the security guarantees of enclave platforms themselves. Patrignani et al. [60, 61] develop abstract trace semantics for low-level code (not including side-channels) on protected module architectures – Fides [74] and Sancus [54] – in order to build secure compilers for these platforms [59]. This is complementary to our work which focuses on formalizing the security guarantees of enclave platforms in presence of a precisely modelled adversary (including side-channel observations).

While our work analyzes models of enclave platforms, verifying actual implementations remains an important challenge. One aspect of this was studied by Leslie-Hurd et al. [39]. They verified linearizability of concurrent updates to shared SGX data structures. **Non-Interference and Hyperproperties:** The security properties of secure measurement, integrity, and confidentiality are formulated as 2-safety observational determinism properties [48, 62], which is a restricted class of hyperproperties [19]. SRE relates closely to the notion of non-interference introduced by Goguen and Meseguer [28], and separability proposed by Rushby [63]. Our confidentiality definition is an adaptation of standard non-interference to the enclave execution model. Separability provides isolation from other programs on the system, but it is too strict for practical applications as it forbids communication between programs and assumes the absence of covert channels; it also does not consider privileged software adversaries as the formalism assumes a safe underlying hardware-software system.

Security Type Systems: A large body of work has studied type systems that enforce information-flow security [17, 22, 65, 73, 80]. Recent examples for hardware design are [40, 41, 85]. SecVerilog [85] extends the Verilog hardware description language with dependent type annotations that define a security policy statically verified by the SecVerilog compiler. One could conceivably implement SGX or Sanctum processors using SecVerilog, thus guaranteeing the implementation does not have unsafe information flow. However, these techniques reason about the security policy at the level of individual signals in the hardware. A higher level of abstraction (like TAP) is needed for reasoning about enclave software. Software fault isolation prevents an untrusted software module from arbitrarily tampering the system's state, and instead restrict all communication to occur via a narrow interface. RockSalt [50] uses Coq to reason about an x86 processor model and guarantee software fault isolation of an untrusted module.

Machine-Checked Proofs of Systems: We perform machine-checked proofs in this work, and similar efforts have verified other classes of systems. The seL4 project [36, 51] proves isolation and information flow enforcement in the seL4 microkernel using the Isabelle/HOL proof assistant [53]. The Ironclad project [30] built a fully verified software stack (including an OS, device drivers, and cryptographic libraries) from the ground-up. The miTLS project [10]

is building a verified reference implementation of TLS which complements our work nicely – enclaves indubitably require TLS channels to communicate with other enclaves and clients. Vijayaraghavan et al. [79] used the Coq proof assistant [9] to verify correctness of a multiprocessor directory-based cache coherence protocol. While our Boogie [7, 24] proofs do involve manual effort, we contend that they are more automated than their hypothetical counterparts in systems such as Isabelle and Coq.

9 CONCLUSION

This paper introduced a framework and methodology to reason about the security guarantees provided by enclave platforms. We introduced the Trusted Abstract Platform (TAP), and performed proofs demonstrating that TAP satisfies the three properties required for secure remote execution (SRE): secure measurement, integrity and confidentiality. We then presented machine-checked proofs stating that models of Intel SGX and Sanctum are refinements of TAP under certain adversarial conditions. Therefore, these platforms also satisfy the properties required for SRE. Overall, this paper took a step towards a unified, extensible framework for reasoning about enclave programs and platforms.

ACKNOWLEDGMENTS

Funding for this research was partially provided by the National Science Foundation under grants CNS-1413920 and CNS-1528108, by the Intel ADEPT Center, by Delta Electronics, and by TerraSwarm, one of six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA.

REFERENCES

- [1] Lenovo ThinkPad System Management Mode Arbitrary Code Execution 0day Exploit. Available at <https://github.com/Cr4sh/ThinkPwn.git>.
- [2] T. Alves and D. Felton. TrustZone: Integrated Hardware and Software Security. *Information Quarterly*, 3(4):18–24, 2004.
- [3] I. Anati, S. Gueron, S. P. Johnson, and V. R. Scarlata. Innovative Technology for CPU Based Attestation and Sealing. In *Proceedings of the 2nd International Workshop on Hardware and Architectural Support for Security and Privacy, HASP*, volume 13, 2013.
- [4] K. Asanovic, R. Avizienis, J. Bachrach, S. Beamer, D. Biancolin, C. Celio, H. Cook, D. Dabbert, J. Hauser, A. Izraelevitz, et al. The Rocket Chip Generator. *EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2016-17*, 2016.
- [5] K. Asanović and D. A. Patterson. Instruction Sets Should Be Free: The Case For RISC-V. *EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2014-146*, 2014.
- [6] M. Barbosa, B. Portela, G. Scerri, and B. Warinschi. Foundations of Hardware-Based Attested Computation and Application to SGX. In *IEEE European Symposium on Security and Privacy, EuroS&P 2016, Saarbrücken, Germany, March 21-24, 2016*, pages 245–260, 2016.
- [7] M. Barnett, B. E. Chang, R. DeLine, B. Jacobs, and K. R. M. Leino. Boogie: A Modular Reusable Verifier for Object-Oriented Programs. In *FMCO '05, LNCS 4111*, pages 364–387, 2005.
- [8] G. Barthe, P. R. D'Argenio, and T. Rezk. Secure information flow by self-composition. *Mathematical Structures in Computer Science*, 21(6):1207–1252, 2011.
- [9] Y. Bertot and P. Castéran. *Interactive Theorem Proving and Program Development: Coq'Art: The Calculus of Inductive Constructions*. Springer Science & Business Media, 2013.
- [10] K. Bhargavan, C. Fournet, and M. Kohlweiss. miTLS: Verifying Protocol Implementations against Real-World Attacks. *IEEE Security & Privacy*, 14(6):18–25, 2016.
- [11] J. Bonneau and I. Mironov. *Cache-Collision Timing Attacks Against AES*, pages 201–215. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006.
- [12] F. Brasser, U. Müller, A. Dmitrienko, K. Kostianen, S. Capkun, and A. Sadeghi. Software Grand Exposure: SGX Cache Attacks Are Practical. *CoRR*, abs/1702.07521, 2017.
- [13] M. C. Browne, E. M. Clarke, and O. Grumberg. Characterizing Finite Kripke Structures in Propositional Temporal Logic. *Theoretical Computer Science*, 59:115–131, 1988.
- [14] B. B. Brumley and N. Tuveri. Remote Timing Attacks Are Still Practical. In *Proceedings of the 16th European Conference on Research in Computer Security, ESORICS'11*, pages 355–371, Berlin, Heidelberg, 2011. Springer-Verlag.
- [15] D. Brumley and D. Boneh. Remote Timing Attacks Are Practical. In *Proceedings of the 12th Conference on USENIX Security Symposium - Volume 12, SSYM'03*, pages 1–1, Berkeley, CA, USA, 2003. USENIX Association.
- [16] D. Champagne and R. B. Lee. Scalable architectural support for trusted software. In *High Performance Computer Architecture (HPCA), 2010 IEEE 16th International Symposium on*, pages 1–12. IEEE, 2010.
- [17] A. Chaudhuri. Language-based security on Android. In *Proceedings of the 2009 Workshop on Programming Languages and Analysis for Security, PLAS 2009, Dublin, Ireland, 15-21 June, 2009*, pages 1–7, 2009.
- [18] C.-T. Chou, P. K. Mannava, and S. Park. A simple method for parameterized verification of cache coherence protocols. In A. J. Hu and A. K. Martin, editors, *Proceedings of the 5th International Conference on Formal Methods in Computer-Aided Design*, pages 382–398, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.
- [19] M. R. Clarkson and F. B. Schneider. Hyperproperties. *Journal of Computer Security*, 18(6):1157–1210, Sept. 2010.
- [20] V. Costan and S. Devadas. Intel SGX Explained. Cryptology ePrint Archive, Report 2016/086, 2016. <http://eprint.iacr.org/2016/086>.
- [21] V. Costan, I. Lebedev, and S. Devadas. Sanctum: Minimal Hardware Extensions for Strong Software Isolation. In *25th USENIX Security Symposium (USENIX Security 16)*, pages 857–874, Austin, TX, 2016. USENIX Association.
- [22] A. Datta, J. Franklin, D. Garg, and D. Kaynar. A Logic of Secure Systems and Its Application to Trusted Computing. In *Proceedings of the 2009 30th IEEE Symposium on Security and Privacy, SP '09*, pages 221–236, Washington, DC, USA, 2009. IEEE Computer Society.
- [23] L. de Moura and N. Bjørner. Z3: An Efficient SMT Solver. In *TACAS '08*, pages 337–340, 2008.
- [24] R. DeLine and K. R. M. Leino. BoogiePL: A typed procedural language for checking object-oriented programs. Technical Report MSR-TR-2005-70, Microsoft Research, 2005.
- [25] L. Domnitsier, A. Jaleel, J. Loew, N. Abu-Ghazaleh, and D. Ponomarev. Non-monopolizable caches: Low-complexity mitigation of cache side channel attacks. *Transactions on Architecture and Code Optimization (TACO)*, 2012.
- [26] S. Embleton, S. Sparks, and C. C. Zou. SMM rootkit: a new breed of OS independent malware. *Security and Communication Networks*, 6(12):1590–1605, 2013.
- [27] C. W. Fletcher, M. v. Dijk, and S. Devadas. A Secure Processor Architecture for Encrypted Computation on Untrusted Programs. In *Proceedings of the Seventh ACM Workshop on Scalable Trusted Computing*, pages 3–8. ACM, 2012.
- [28] J. A. Goguen and J. Meseguer. Security Policies and Security Models. In *1982 IEEE Symposium on Security and Privacy, Oakland, CA, USA, April 26-28, 1982*, pages 11–20, 1982.
- [29] D. Grawrock. *Dynamics of a Trusted Platform: A building block approach*. Intel Press, 2009.
- [30] C. Hawblitzel, J. Howell, J. R. Lorch, A. Narayan, B. Parno, D. Zhang, and B. Zill. Ironclad Apps: End-to-End Security via Automated Full-System Verification. In *Proceedings of the 11th USENIX conference on Operating Systems Design and Implementation*, pages 165–181, 2014.
- [31] M. P. Herlihy and J. M. Wing. Linearizability: A Correctness Condition for Concurrent Objects. *ACM Transactions on Programming Languages and Systems*, 12(3):463–492, July 1990.
- [32] M. Hoekstra, R. Lal, P. Pappachan, V. Phegade, and J. Del Cuvillo. Using Innovative Instructions to Create Trustworthy Software Solutions. In *Proceedings of the 2nd International Workshop on Hardware and Architectural Support for Security and Privacy, HASP*, volume 13, 2013.
- [33] Intel Software Guard Extensions Programming Reference. Available at <https://software.intel.com/sites/default/files/329298-001.pdf>.
- [34] G. Irazoqui, T. Eisenbarth, and B. Sunar. S&A: A Shared Cache Attack That Works across Cores and Defies VM Sandboxing – and Its Application to AES. In *IEEE Symposium on Security and Privacy*, pages 591–604, May 2015.
- [35] Joanna Rutkowska. Red Pill... or how to detect VMM using (almost) one CPU instruction. <https://github.com/Cr4sh/ThinkPwn.git>.
- [36] G. Klein, K. Elphinstone, G. Heiser, J. Andronick, D. Cock, P. Derrin, D. Elkaduwe, K. Engelhardt, R. Kolanski, M. Norrish, T. Sewell, H. Tuch, and S. Winwood. seL4: Formal Verification of an OS Kernel. In *Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles, SOSP '09*, pages 207–220, New York, USA, 2009.
- [37] P. C. Kocher. Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. In *Proceedings of the 16th Annual International Cryptology Conference on Advances in Cryptology, CRYPTO '96*, pages 104–113, London, UK, 1996. Springer-Verlag.
- [38] S. Lee, M. Shih, P. Gera, T. Kim, H. Kim, and M. Peinado. Inferring Fine-grained Control Flow Inside SGX Enclaves with Branch Shadowing. *CoRR*, abs/1611.06952, 2016.

- 2016.
- [39] R. Leslie-Hurd, D. Caspi, and M. Fernandez. Verifying Linearizability of Intel® Software Guard Extensions. In *Computer Aided Verification - 27th International Conference, CAV 2015, San Francisco, CA, USA, July 18–24, 2015, Proceedings, Part II*, pages 144–160, 2015.
- [40] X. Li, V. Kashyap, J. K. Oberg, M. Tiwari, V. R. Rajarathinam, R. Kastner, T. Sherwood, B. Hardekopf, and F. T. Chong. Sapper: A Language for Hardware-Level Security Policy Enforcement. In *Architectural Support for Programming Languages and Operating Systems, ASPLOS '14, Salt Lake City, UT, USA, March 1–5, 2014*, pages 97–112, 2014.
- [41] X. Li, M. Tiwari, J. Oberg, V. Kashyap, F. T. Chong, T. Sherwood, and B. Hardekopf. Caisson: A Hardware Description Language for Secure Information Flow. In *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2011, San Jose, CA, USA, June 4–8, 2011*, pages 109–120, 2011.
- [42] D. Lie, C. Thekkath, M. Mitchell, P. Lincoln, D. Boneh, J. Mitchell, and M. Horowitz. Architectural support for copy and tamper resistant software. *ACM SIGPLAN Notices*, 35(11):168–177, 2000.
- [43] F. Liu, Q. Ge, Y. Yarom, F. McKeen, C. Rozas, G. Heiser, and R. B. Lee. CATalyst: Defeating Last-Level Cache Side Channel Attacks in Cloud Computing. In *2016 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Mar 2016.
- [44] F. Liu and R. B. Lee. Random Fill Cache Architecture. In *2014 47th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, IEEE, 2014.
- [45] F. Liu, Y. Yarom, Q. Ge, G. Heiser, and R. B. Lee. Last-Level Cache Side-Channel Attacks Are Practical. In *Proceedings of the 2015 IEEE Symposium on Security and Privacy*, pages 605–622, Washington, DC, USA, 2015. IEEE Computer Society.
- [46] M. Maas, E. Love, E. Stefanov, M. Tiwari, E. Shi, K. Asanovic, J. Kubiawicz, and D. Song. Phantom: Practical oblivious computation in a secure processor. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 311–324. ACM, 2013.
- [47] F. McKeen, I. Alexandrovich, A. Berenzon, C. V. Rozas, H. Shafi, V. Shanbhogue, and U. R. Savagaonkar. Innovative Instructions and Software Model for Isolated Execution. *HASP*, 13:10, 2013.
- [48] J. Mclean. Proving Noninterference and Functional Correctness Using Traces. *Journal of Computer Security*, 1:37–58, 1992.
- [49] A. Moghimi, G. Irazoqui, and T. Eisenbarth. CacheZoom: How SGX Amplifies The Power of Cache Attacks. *CoRR*, abs/1703.06986, 2017.
- [50] G. Morrisett, G. Tan, J. Tassarotti, J. Tristan, and E. Gan. RockSalt: better, faster, stronger SFI for the x86. In *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '12, Beijing, China - June 11 - 16, 2012*, pages 395–404, 2012.
- [51] T. Murray, D. Matchuk, M. Brassil, P. Gammie, T. Bourke, S. Seefried, C. Lewis, X. Gao, and G. Klein. seL4: From General Purpose to a Proof of Information Flow Enforcement. In *Security and Privacy (SP), 2013 IEEE Symposium on*, pages 415–429. IEEE, 2013.
- [52] M. Neugschwandtner, C. Platzer, P. M. Comparetti, and U. Bayer. dAnubis - Dynamic Device Driver Analysis Based on Virtual Machine Introspection. In *Detection of Intrusions and Malware, and Vulnerability Assessment, 7th International Conference, DIMVA 2010, Bonn, Germany, July 8–9, 2010. Proceedings*, pages 41–60, 2010.
- [53] T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL: A Proof Assistant for Higher-Order Logic*, volume 2283. Springer Science & Business Media, 2002.
- [54] J. Noorman, P. Agten, W. Daniels, R. Strackx, A. Van Herreweghe, C. Huygens, B. Preneel, I. Verbauwhede, and F. Piessens. Sancus: Low-cost Trustworthy Extensible Networked Devices with a Zero-software Trusted Computing Base. In *Proceedings of the 22nd USENIX Conference on Security, SEC'13*, pages 479–494, Berkeley, CA, USA, 2013. USENIX Association.
- [55] O. Ohrimenko, F. Schuster, C. Fournet, A. Mehta, S. Nowozin, K. Vaswani, and M. Costa. Oblivious Multi-Party Machine Learning on Trusted Processors. In *25th USENIX Security Symposium (USENIX Security 16)*, pages 619–636, Austin, TX, 2016. USENIX Association.
- [56] Y. Oren, V. P. Kemerlis, S. Sethumadhavan, and A. D. Keromytis. The Spy in the Sandbox - Practical Cache Attacks in Javascript. *CoRR*, abs/1502.07373, 2015.
- [57] B. Parno, J. R. Lorch, J. R. Douceur, J. Mickens, and J. M. McCune. Memoir: Practical State Continuity for Protected Modules. In *Proceedings of the 2011 IEEE Symposium on Security and Privacy, SP '11*, pages 379–394, Washington, DC, USA, 2011. IEEE Computer Society.
- [58] R. Pass, E. Shi, and F. Tramèr. Formal Abstractions for Attested Execution Secure Processors. *IACR Cryptology ePrint Archive*, 2016:1027, 2016.
- [59] M. Patrignani, P. Agten, R. Strackx, B. Jacobs, D. Clarke, and F. Piessens. Secure Compilation to Protected Module Architectures. *ACM Trans. Program. Lang. Syst.*, 37(2):6:1–6:50, 2015.
- [60] M. Patrignani and D. Clarke. Fully abstract trace semantics for low-level isolation mechanisms. In *Symposium on Applied Computing, SAC 2014, Gyeongju, Republic of Korea - March 24 - 28, 2014*, pages 1562–1569, 2014.
- [61] M. Patrignani and D. Clarke. Fully abstract trace semantics for protected module architectures. *Computer Languages, Systems & Structures*, 42:22–45, 2015.
- [62] A. W. Roscoe. CSP and determinism in security modelling. In *Proceedings of the 1995 IEEE Symposium on Security and Privacy, Oakland, California, USA, May 8–10, 1995*, pages 114–127, 1995.
- [63] J. M. Rushby. Proof of separability: A verification technique for a class of a security kernels. In *International Symposium on Programming, 5th Colloquium, Torino, Italy, April 6–8, 1982, Proceedings*, pages 352–367, 1982.
- [64] J. Rutkowska. Security challenges in virtualized environments.
- [65] A. Sabelfeld and A. C. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 21(1):5–19, 2003.
- [66] M. Schwarz, S. Weiser, D. Gruss, C. Maurice, and S. Mangard. Malware Guard Extension: Using SGX to Conceal Cache Attacks. *CoRR*, abs/1702.08719, 2017.
- [67] J. Seo, B. Lee, S. Kim, M.-W. Shih, I. Shin, D. Han, and T. Kim. SGX-Shield: Enabling address space layout randomization for SGX programs. In *23rd Annual Network and Distributed System Security Symposium, NDSS 2017, San Diego, California, USA, February 26–March 1, 2017*, 2017.
- [68] M.-W. Shih, S. Lee, T. Kim, and M. Peinado. T-SGX: Eradicating Controlled-Channel Attacks Against Enclave Programs. In *Proceedings of the 2017 Annual Network and Distributed System Security Symposium (NDSS)*, San Diego, CA, Feb. 2017.
- [69] S. Shinde, Z. L. Chua, V. Narayanan, and P. Saxena. Preventing Page Faults from Telling Your Secrets. In *Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security, AsiaCCS 2016, Xi'an, China, May 30 - June 3, 2016*, pages 317–328, 2016.
- [70] R. Sinha, M. Costa, A. Lal, N. P. Lopes, S. K. Rajamani, S. A. Seshia, and K. Vaswani. A Design and Verification Methodology for Secure Isolated Regions. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2016, Santa Barbara, CA, USA, June 13–17, 2016*, pages 665–681, 2016.
- [71] R. Sinha, S. K. Rajamani, S. A. Seshia, and K. Vaswani. Moat: Verifying Confidentiality of Enclave Programs. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12–6, 2015*, pages 1169–1184, 2015.
- [72] T. Skolem. Logico-combinatorial investigations in the satisfiability or provability of mathematical propositions: a simplified proof of a theorem by L. Löwenheim and generalizations of the theorem. *From Frege to Gödel. A Source Book in Mathematical Logic, 1879–1931*, pages 252–263, 1967.
- [73] G. Smith and D. M. Volpano. Secure Information Flow in a Multi-Threaded Imperative Language. In *POPL '98, Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Diego, CA, USA, January 19–21, 1998*, pages 355–364, 1998.
- [74] R. Strackx and F. Piessens. Fides: Selectively Hardening Software Application Components Against Kernel-level or Process-level Malware. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security, CCS '12*, pages 2–13, New York, NY, USA, 2012. ACM.
- [75] P. Subramanyan, R. Sinha, I. Lebedev, S. Devadas, and S. A. Seshia. Models and Proofs for the Trusted Abstract Platform (TAP), Intel SGX and MIT Sanctum. <https://github.com/0tcb/TAP>.
- [76] G. E. Suh, D. Clarke, B. Gassend, M. Van Dijk, and S. Devadas. AEGIS: architecture for tamper-evident and tamper-resistant processing. In *Proceedings of the 17th annual international conference on Supercomputing*, pages 160–171. ACM, 2003.
- [77] T. Terauchi and A. Aiken. Secure Information Flow as a Safety Problem. In *Static Analysis Symposium (SAS '05)*, LNCS 3672, pages 352–367, 2005.
- [78] E. Tromer, D. A. Osvik, and A. Shamir. Efficient Cache Attacks on AES, and Countermeasures. *J. Cryptology*, 23(1):37–71, 2010.
- [79] M. Vijayaraghavan, A. Chlipala, Arvind, and N. Dave. Modular Deductive Verification of Multiprocessor Hardware Designs. In *Computer Aided Verification - 27th International Conference, CAV 2015, San Francisco, CA, USA, July 18–24, 2015, Proceedings, Part II*, pages 109–127, 2015.
- [80] D. Volpano, C. Irvine, and G. Smith. A Sound Type System for Secure Flow Analysis. *Journal of Computer Security*, 4(2-3):167–187, Jan. 1996.
- [81] A. Waterman, Y. Lee, R. Avizienis, D. A. Patterson, and K. Asanović. The RISC-V Instruction Set Manual Volume II: Privileged Architecture Version 1.9.1. Technical Report UCB/EECS-2016-161, EECS Department, University of California, Berkeley, Nov 2016.
- [82] A. Waterman, Y. Lee, D. A. Patterson, and K. Asanović. The RISC-V Instruction Set Manual, Volume I: User-Level ISA, Version 2.0. Technical Report UCB/EECS-2014-54, EECS Department, University of California, Berkeley, May 2014.
- [83] Y. Xu, W. Cui, and M. Peinado. Controlled-Channel Attacks: Deterministic Side Channels for Untrusted Operating Systems. In *2015 IEEE Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17–21, 2015*, pages 640–656, 2015.
- [84] Y. Yarom and K. Falkner. FLUSH+RELOAD: A High Resolution, Low Noise, L3 Cache Side-Channel Attack. In *Proceedings of the 23rd USENIX Security Symposium, San Diego, CA, USA, August 20–22, 2014*, pages 719–732, 2014.
- [85] D. Zhang, Y. Wang, G. E. Suh, and A. C. Myers. A Hardware Design Language for Timing-Sensitive Information-Flow Security. In *Proceedings of the Twentieth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS '15, Istanbul, Turkey, March 14–18, 2015*, pages 503–516, 2015.

Table 5: Glossary of Symbols

Symbol	Description
\mathbb{N}	Natural numbers.
\mathbb{B}	Booleans ($\mathbb{B} = \{\text{true}, \text{false}\}$).
$\lambda x. \text{expr}$	Function with argument x ; computes expr .
$m[i]$	Element i in map m .
$\text{ITE}(c, x, y)$	If-then-else operator: evaluates to x if c is true, y otherwise.
rec.fld	Field fld in record rec .
VA	Virtual addresses.
PA	Physical addresses.
ACL	Permissions for virtual address (Readable/Writeable/eXecutable etc.)
Tag	Cache tags.
\mathcal{E}_{id}	Type of enclave “ids” (integer/pointer).
$\mathcal{E}_{\mathcal{M}}$	Enclave metadata type.
Σ	Set of all TAP states.
\leadsto	TAP transition relation.
init	TAP initial state.
$\sigma, \sigma_0, \sigma_1, \dots$	TAP states. ($\sigma \in \Sigma$).
$\sigma_i \leadsto \sigma_j$	TAP steps from state σ_i to σ_j .
π, π_0, π_1, \dots	Traces of the TAP.
$\pi[0]$	Initial state of trace π .
$\pi[i]$	i th state of in trace π .
$\sigma(\text{expr})$	Expression expr evaluated in state σ .
e, e_1, e_2, \dots	Enclave programs.
$E_e(\sigma)$	Enclave e 's state when in platform state σ .
$I_e(\sigma)$	Enclave e 's input when in platform state σ .
$I_e^R(\sigma)$	Randomness component of enclave e 's input.
$I_e^U(\sigma)$	Untrusted component of enclave e 's input.
$O_e(\sigma)$	Enclave e 's output when in platform state σ .
$A_e(\sigma)$	Adversary's state when in platform state σ .
$I^P(\sigma)$	Non-deterministic component of platform state.
$\mu(e)$	Measurement of enclave e .
$\llbracket e \rrbracket$	Set of all traces of e .
$m[v]_{\text{PA}}$	Physical address for virtual address v in virtual to physical mapping m .
$\text{valid}(m[v])$	Is virtual address v is mapped to some physical address in the virtual to physical mapping m ?
$\text{executable}(m[v])$	Does virtual address v has execute permissions in the virtual to physical mapping m .
Σ_L	Set of all implementation states.
\leadsto_L	Implementation transition relation.
init_L	Implementation initial state.
s, s_0, s_1, \dots	Implementation states. ($s \in \Sigma_L$).
$s_i \leadsto_L s_j$	Implementation steps from state s_i to s_j .
$m_1 m_2$	m_1 concatenated with m_2 .
PK_k	Public key k .
SK_k	Secret key k .
$\{m\}_{SK_k}$	Msg m encrypted/signed with key SK_k .