

Formal Verification of Virtualization-Based Trusted Execution Environments

Hasini Witharana¹, Student Member, IEEE, Hansika Weerasena¹, Graduate Student Member, IEEE,
and Prabhat Mishra¹, Fellow, IEEE

Abstract—Trusted execution environments (TEEs) provide a secure environment for computation, ensuring that the code and data inside the TEE are protected with respect to confidentiality and integrity. Virtual machine (VM)-based TEEs extend this concept by utilizing virtualization technology to create isolated execution spaces that can support a complete operating system or specific applications. As the complexity and importance of VM-based TEEs grow, ensuring their reliability and security through formal verification becomes crucial. However, these technologies often operate without formal assurances of their security properties. Our research introduces a formal framework for representing and verifying VM-based TEEs. This approach provides a rigorous foundation for defining and verifying key security attributes for safeguarding execution environments. To demonstrate the applicability of our verification framework, we conduct an analysis of real-world TEE platforms, including Intel’s trust domain extensions (TDX). This work not only emphasizes the necessity of formal verification in enhancing the security of VM-based TEEs but also provides a systematic approach for evaluating the resilience of these platforms against sophisticated adversarial models.

Index Terms—Confidential computing, confidentiality, integrity, property checking, trusted execution environments (TEEs).

I. INTRODUCTION

AS THE nature of computing evolves, ensuring the security and trustworthiness of sensitive data and critical applications has become an important concern. With the rapid growth of cloud computing, edge devices, and the Internet of Things (IoT), the need for robust security measures has never been more critical. Trusted execution environments (TEEs) emerge as a promising solution to improve security by providing a secure environment for the execution of sensitive code with sensitive data and the protection of confidential information [1]. TEEs offer a secure execution environment that is isolated from the rest of the system, safeguarding against various threats, such as malicious software, unauthorized access, and hardware-based attacks.

Manuscript received 2 August 2024; accepted 3 August 2024. This work was supported in part by the Semiconductor Research Corporation (SRC) under Grant 2022-HW-3128. This article was presented at the International Conference on Compilers, Architectures, and Synthesis for Embedded Systems (CASES) 2024 and appeared as part of the ESWEK-TCAD Special Issue. This article was recommended by Associate Editor S. Dailey. (Corresponding author: Hasini Witharana.)

The authors are with the Department of Computer and Information Science and Engineering, University of Florida, Gainesville, FL 32611 USA (e-mail: witharana.hasini@ufl.edu).

Digital Object Identifier 10.1109/TCAD.2024.3443008

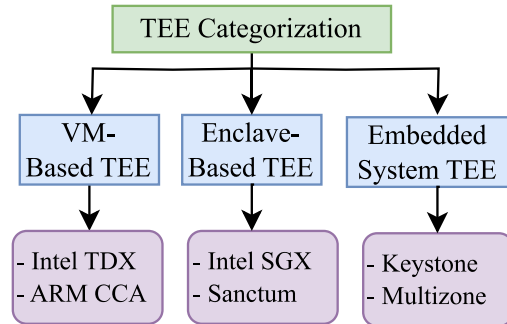


Fig. 1. Categories of TEEs.

TEEs utilize hardware and software components to establish a secure environment where cryptographic operations, key management, and other critical and confidential tasks can be performed with security assurance.

TEEs come in various forms, each tailored to meet specific requirements and challenges. Fig. 1 shows three types of TEEs: 1) enclave-based TEEs (e.g., Intel software guard extensions (SGX) [2], Sanctum [3]); 2) virtual machine (VM)-based TEEs (e.g., Intel trust domain extensions (TDX) [4], AMD secure encrypted virtualization (SEV) [5]), ARM confidential computing architecture (CCA); and 3) TEEs for embedded systems (e.g., Keystone [6]). Enclave-based TEEs leverage hardware-supported isolation to create secure enclaves within a processor. These enclaves are isolated regions of memory resistant to external tampering and surveillance, ensuring the integrity and confidentiality of the code and data. Intel SGX is a prime example of an enclave-based TEE, allowing developers to create secure enclaves for the execution of sensitive operations without revealing the data to the underlying system. VM-based TEEs take advantage of virtualization technologies to create secure execution environments within VMs. Intel TDX and AMD SEV are some examples of VM-based TEEs. They extend security to the virtualization layer by protecting against attacks even in the presence of compromised hypervisors. Embedded system-based TEEs are designed to cater to the unique constraints and requirements of embedded systems and IoT devices. For example, Keystone integrates with RISC-V architectures to provide hardware-enforced memory protection and secure execution environments, making it well-suited for resource-constrained embedded systems.

VM-based TEEs are important in cloud computing due to their ability to offer scalable and flexible security solutions that are well-suited to the dynamic nature of cloud services. Unlike




Fig. 2. Overview of our formal verification framework.

72 enclave-based TEEs, which are designed for securing small
 73 pieces of sensitive code and data within tightly controlled
 74 memory regions, VM-based TEEs can secure entire VMs,
 75 offering a broader and more flexible approach to isolation
 76 and security in cloud environments. In this article, we present
 77 a formal verification framework for VM-based TEEs for
 78 confidentiality and integrity.

79 Fig. 2 presents an overview of our formal verification
 80 framework for security verification of TEE architectures. We
 81 first conduct an abstraction of the TEE architecture that
 82 accurately represents the TEE behavior by following the spec-
 83 ification. This abstraction phase simplifies the specification,
 84 focusing on the essential aspects relevant to confidentiality
 85 and integrity. Next, we develop a formal model for VM-based
 86 TEE architectures based on the abstraction. Then, we derive
 87 properties related to confidentiality and integrity from the TEE
 88 specification. Finally, we perform property checking to verify
 89 whether the TEE formal model satisfies the specified proper-
 90 ties, ensuring that the TEE architecture meets the predefined
 91 security criteria on confidentiality and integrity. Specifically,
 92 this article makes the following major contributions: 1) we
 93 present a comprehensive formal model that defines the security
 94 boundaries of confidential VMs, explicitly considering the
 95 capabilities of adversaries with access to advanced attack
 96 vectors; 2) we formally model confidentiality and integrity
 97 properties, tailoring them for VM-based TEEs; 3) we introduce
 98 a detailed formal model for the Intel TDX architecture,
 99 developed using the Rosette language; and 4) we formally
 100 verify the confidentiality and integrity of Intel TDX for code
 101 and data in use.

102 This article is organized as follows. Section II provides rele-
 103 vant background and surveys related efforts. Section III defines
 104 the formal model for VM and adversary. Section IV pro-
 105 vides formal definitions for both confidentiality and integrity.
 106 Section V conducts a security analysis of Intel TDX for
 107 confidentiality and integrity. Sections VI and VII provide
 108 details of formal modeling of TDX architecture and cache.
 109 Section VIII discusses the results of the formal analysis.
 110 Finally, Section IX concludes this article.

111 II. BACKGROUND AND RELATED WORK

112 This section first provides relevant background on
 113 VM-based TEEs. Next, it surveys related efforts in security
 114 verification of TEE architectures.

115 A. Background: VM-Based TEEs

116 We first introduce VMs. Next, we discuss confidential
 117 VMs. Finally, we provide an overview of Intel TDX, which
 118 implements confidential VM architecture.

119 1) *Virtual Machine*: A VM enables software-based emula-
 120 tion of physical computers. This technology allows for running
 121 an operating system (OS) and applications within an isolated
 122 and encapsulated environment. VMs facilitate multiple OSes
 123 to operate concurrently on the same physical hardware. This is
 124 achieved through virtualization, which significantly enhances
 125 resource utilization, flexibility, and isolation in computing
 126 environments. At the heart of a VM lies the hypervisor, or VM
 127 monitor (VMM), a critical component tasked with managing
 128 and allocating the physical resources among VMs. Hypervisors
 129 come in two varieties: 1) Type 1 (bare-metal), which operates
 130 directly on the host hardware and 2) Type 2 (hosted), which
 131 functions on top of an existing OS.

132 The VMM orchestrates access to hardware components,
 133 such as CPUs, memory, storage, and network interfaces,
 134 enabling the seamless and concurrent operation of multiple
 135 VMs on a single physical machine. The core virtualization
 136 concept: hardware abstraction allows each VM to operate as
 137 though it has its own dedicated hardware. Each VM hosts its
 138 own guest OS, providing an independent operating environ-
 139 ment that interacts with the virtualized hardware, ensuring that
 140 applications run in a manner that is both efficient and isolated
 141 from the host system and other VMs. Even though VMs are
 142 isolated from each other, they are not entirely separate entities
 143 in terms of security. The shared use of the hypervisor, under-
 144 lying hardware, memory subsystem, and other components
 145 of the virtualization stack introduces potential vulnerabilities.
 146 These shared resources can become attack vectors, where a
 147 malicious entity might exploit one VM or host system to gain
 148 unauthorized access to or influence over others or the host
 149 system itself. This inherent risk highlights the critical need for
 150 confidential VMs.

151 2) *Confidential Virtual Machines*: Confidential VMs are
 152 designed to protect against threats, including malicious
 153 insiders, compromised hypervisors, and other potential vul-
 154 nerabilities in the virtualization stack. Confidential VMs use
 155 memory encryption with a unique key for each VM to
 156 protect the contents of the VM's memory from unauthorized
 157 access. This ensures confidentiality by ensuring the memory
 158 contents remain inaccessible and secure from external threats
 159 and internal attackers gaining access to physical memory.
 160 Furthermore, they also provide integrity protection mecha-
 161 nisms to verify that data and code have not been tampered
 162 with. The creation of confidential VMs often utilizes hardware-
 163 based security features offering a level of protection that
 164 extends even to the host hypervisor. Confidential VMs often
 165 use secure boot mechanisms and attestation processes. A
 166 secure boot ensures that only authenticated and trusted code
 167 is executed during the VM's startup. This works against
 168 malicious software and rootkits that might attempt to load
 169 during the boot process. Attestation verifies the integrity and
 170 authenticity of the VM for external entities. Specifically,
 171 attestation allows a third party to confirm that the VM is
 172 running the expected software stack.

173 Fig. 3 shows the basic building blocks required for
 174 VM-based TEE architecture. It starts with a secure boot pro-
 175 cess; the system relies on a foundational security mechanism
 176 known as the root of trust (RoT), complemented by the prin-
 177 ciple of a chain of trust. The RoT is pivotal for ensuring that




Fig. 3. Overview of VM-based confidential computing.

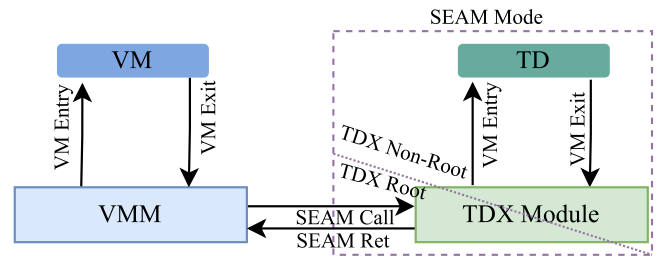


Fig. 4. Overview of Intel TDX architecture.

178 only authenticated and integrity-verified firmware and software
 179 are loaded for execution. It achieves this through the provision
 180 of essential cryptographic functions and services. Initially, the
 181 RoT verifies the integrity and authenticity of the bootloader
 182 and establishes the first link in the chain of trust. Once the
 183 bootloader is authenticated, it securely loads and verifies the
 184 firmware, setting the stage for the execution environment.
 185 Hypervisor can create and manage VMs. In a type 2 hypervisor
 186 configuration, a host OS will work alongside the hypervisor,
 187 while type 1 will have only a hypervisor. A fundamental
 188 element of a VM-based confidential computing framework
 189 is the security monitor. It operates at a low level, closely
 190 interacting with the hypervisor and host OS to monitor and
 191 control access to resources, manage permissions, and ensure
 192 isolation between different confidential VMs. The security
 193 monitor's primary objectives include preventing unauthorized
 194 access to sensitive data, ensuring that software components
 195 cannot interfere with each other maliciously, and enforcing
 196 compliance with security protocols. Intel TDX module [4] is
 197 one of the examples of a secure monitor.

198 3) *Intel TDX*: Intel TDX [4] is an example of a confidential
 199 VM architecture. TDX provides the infrastructure to create
 200 hardware-isolated VMs known as trust domains (TDs), which
 201 are designed to operate securely within a system, separate
 202 from the VMM or hypervisor and any other unrelated software
 203 entities. To protect the confidentiality and integrity of the
 204 code and data within a TD, Intel TDX uses technologies,
 205 such as multikey total memory encryption (MKTME) and
 206 hashing techniques. Fig. 4 shows an overview of the Intel
 207 TDX architecture. Intel TDX is engineered to function within
 208 a secure arbitration mode (SEAM), which is an extension to
 209 the prior VM extension (VMX) architecture. SEAM introduces
 210 a new VMX root operation mode, referred to as SEAM root,
 211 specifically constructed to support CPU-attested modules that
 212 establish TDs for VM guests. The SEAM operation is divided
 213 into two logical modes: 1) TDX nonroot mode for the TD
 214 guest operations and 2) TDX root mode, which is reserved for
 215 host-side activities.

216 B. Related Work

217 This section surveys related efforts, including static analysis,
 218 simulation-based testing, and formal verification.

219 1) *Static Analysis*: Google's security review of Intel TDX
 220 employed static analysis tools to uncover numerous attack
 221 vectors and security issues [7]. Security review discovered
 222 81 potential avenues for attacks, confirmed 10 security flaws,

and made 5 modifications to enhance the code's defense
 mechanisms. The review assessed four components of Intel
 TDX, including the MCHECK mechanism in BIOS, the
 nonpersistent SEAM loader, the persistent SEAM loader, and
 the design of the TDX module. However, this approach did
 not provide formal security guarantees. In fact, this security
 review highlights the need of formal verification.

223
 224
 225
 226
 227
 228
 229 2) *Simulation-Based Validation*: Simulation-based testing
 methodologies have been used to evaluate the security
 of TEEs. Google's examination of AMD SEV technology
 through simulation-based testing uncovered critical vulnera-
 bilities [8]. This hands-on approach allows for a practical
 assessment of TEE security. Simulation-based verification
 faces the exponential input space complexity to cover all
 possible scenarios [9], [10]. Nevertheless, the lack of formal
 security guarantees limits the ability of simulation-based
 testing to guarantee the security properties of TEE systems.

230
 231
 232
 233
 234
 235
 236
 237
 238
 239 3) *Formal Verification*: ProVeriT [11] provides a theorem
 proving solution to formally verify Global Platform TEE
 common criteria. Ma et al. [12] developed a formal model for
 memory isolation that includes a detailed formalization of the
 ARMv8 architecture's hardware components associated with
 memory isolation, as well as the formalization of a TrustZone
 monitor that facilitates switching between secure and non-
 secure worlds. A recent study [13] introduces a verification
 methodology for ARM TrustZone using property checking
 techniques. Sardar et al. [14], [15] formally specifies the
 attestation mechanism using ProVeriT's specification language.
 This work only focuses on the attestation process, whereas
 our work focuses on memory confidentiality and integrity of
 Intel TDX. Ozga [16] presented a methodology for formally
 modeling and proving the security of a security monitor, which
 is a key component of VM-based confidential computing
 systems.

240
 241
 242
 243
 244
 245
 246
 247
 248
 249
 250
 251
 252
 253
 254
 255
 256 While there are promising efforts for formally verifying
 different types of TEE architectures, including verification
 of Intel SGX [17], verification of ARM Trustzone [12],
 and verification of RISC-V based TEEs [16], and existing
 formal verification solutions, cannot be directly applied to
 the TDX architecture due to their inherent differences in the
 implementation of the TEE architecture.

264 III. FORMAL MODELING OF VM AND ADVERSARY

265 In this section, we first define a formal model for VM,
 266 including VM state, VM inputs, and VM outputs. Next, we
 267 define a formal model for the adversary. Throughout this

268 article, the symbol $=$ is used to denote intensional equality,
 269 which asserts that two expressions or variables are equivalent
 270 in all respects, including their state, value, or configuration.
 271 Similarly, the symbol \Leftrightarrow represents an equivalence relation or
 272 extensional equality in our context.

273 A. Formal Model for Virtual Machines

274 A VM is initiated with a specific allocation of resources,
 275 including CPU cores, memory, and storage. The virtualization
 276 platform often uses a unique identifier or configuration snap-
 277 shot of the foundational state, enabling users to guarantee that
 278 the VM was initialized according to the predefined settings.
 279 The VM's configuration includes the boot sequence with OS
 280 image, and the virtual hardware components assigned to the
 281 VM, such as memory size and disk space.

282 *VM*: A user deploys a VM denoted as v . The attributes of
 283 this VM include a unique identifier for the VM ($v.id$), the
 284 VM's OS image ($v.os$), VM's virtual address list ($v.valist$),
 285 VM's memory size ($v.mem$), and VM's data and code pages
 286 ($v.data$). These attributes define the configuration of the VM,
 287 enabling it to perform designated computing tasks within a
 288 virtualized environment.

289 *VM State*: At any given moment, the host machine exists in
 290 a certain state (m). The state of the VM, $S_v(m)$, can be seen as
 291 a specific instance of the overall system state, capturing key
 292 operational data. This includes the virtual memory mapping
 293 $Vmem : Va \rightarrow W$, which represents a function from virtual
 294 addresses (Va) to their corresponding values in machine words
 295 (W); a set of general-purpose registers $regs : \mathbb{N} \rightarrow W$ indexed
 296 by natural numbers; the program counter $pc : Va$, indicating the
 297 VM's current execution point; and the VM's attributes, which
 298 are established at the VM's creation and remain unchanged
 299 during its operation. The initial state $init_v$ defines the starting
 300 condition of the VM's memory ($Vmem$) at the time of its
 301 instantiation. For simplicity, $init_v(S_v(m_v))$ denotes that $S_v(m_v)$
 302 is in its initial state, as configured before any operations have
 303 been executed within the VM.

304 *VM Inputs*: The inputs to the VM, $I_v(m) = (I_v^D(m) + I_v^R(m))$,
 305 can be categorized into two main types: 1) external inputs
 306 and 2) internal inputs. External Inputs ($I_v^D(m)$) can change the
 307 state of a VM, such as initializing it to runnable. These inputs
 308 are received from the external environment and, given the VM
 309 operates in a potentially hostile environment, may come from
 310 sources under adversarial control. However, they consist of
 311 a predefined set of instructions, making the impact of these
 312 inputs deterministic. On the other hand, internal inputs ($I_v^R(m)$)
 313 run inside the VM, such as code and data, where the impact
 314 of the internal inputs can be variable.

315 *VM Outputs*: The outputs of the VM, $O_v(m)$, project the
 316 machine state of the VM. VM output can have both encrypted
 317 data in memory and decrypted data inside the processor.
 318 Specifically, $O_v(m)$ focuses on memory elements while data
 319 is in use.

320 *VM Execution*: The execution of a VM is modeled as a
 321 deterministic process with respect to input $I_v(m)$, where the
 322 next state of the VM, is a function of its current state, $S_v(m)$,
 323 and its inputs, $I_v(m)$. We assume that one virtual CPU for

each VM and single thread is used per applications in the 324
 VM. We also assume that code and data running inside the 325
 VM is not malicious. Therefore, it is safe to assume that 326
 $I_v^R(m)$ does not lead to nondeterministic process. Given the 327
 deterministic assumption, the VM's execution at any step can 328
 be defined by the transitive closure of the transition relation 329
 $m_i \rightsquigarrow m_j$, indicating that the VM can transition from state m_i 330
 to state m_j based on the operational semantics of its instruction 331
 set. This transition relation implies a set of all possible 332
 states that can be reached from a given state, directly or 333
 indirectly, through multiple steps or transitions. It essentially 334
 expands the basic transition relation to include not just direct 335
 successors but all reachable states. This transition process 336
 involves first identifying the next instruction to execute based 337
 on the current state of the virtual memory ($Vmem$) and the 338
 program counter (pc). Following this, the identified instruction 339
 is executed, which may involve bitvector operations, memory 340
 accesses, and interactions with VM-specific primitives for 341
 security, randomness, and I/O operations. 342

B. Formal Model for Adversary 343

A confidential VM operates under the assumption of a 344
 privileged adversary who has compromised all software layers 345
 except for the confidential VM platform itself (security moni- 346
 tor). This section defines the adversary's potential actions and 347
 their implications for the VM. 348

Adversary State: The privileged adversary is capable of 349
 pausing the VM at any moment, executing arbitrary instruc- 350
 tions that can modify the adversary's state ($A_v(m)$), the VM's 351
 inputs ($I_v(m)$), and can initiate or terminate VM instances. We 352
 show the adversary's influence through the *attack* relation over 353
 pairs of states: $(m_1, m_2) \in attack$ if the attacker can transition 354
 the system's state from m_1 to m_2 . A key constraint is that 355
 the *attack* operation cannot alter the confidential VM state, 356
 ensuring $S_v(m_1) = S_v(m_2)$. This maintains the integrity of the 357
 VM despite the adversary's actions. 358

Here, *attack* is a subset of the transition relation \rightsquigarrow , indi- 359
 cating that an adversary's actions are confined to utilizing the 360
 platform's instructions to alter the system's state. Furthermore, 361
 the *attack* relation is reflexive, denoting that the adversary 362
 might choose not to alter the state: $\forall m.(m, m) \in attack$. This 363
 model allows the adversary to operate concurrently with the 364
 VM, with the capability to modify the machine's state before 365
 the VM's launch and to alter the VM's initial state. 366

Adversary Monitoring: In a confidential VM environ- 367
 ment, untrusted software, including potential adversaries, may 368
 observe aspects of the VM's execution. These observations are 369
 contingent on the confidentiality protections enforced by the 370
 VM platform. While explicit outputs are invariably observable, 371
 adversaries might also detect patterns through indirect means, 372
 such as side channels, including memory access patterns and 373
 computational timing. 374

The capability for an adversary to make observations is 375
 formalized through the execution of arbitrary instructions or 376
 the utilization of platform primitives. These actions allow 377
 the adversary to monitor the effects of their operations 378
 on the VM's state. Let $monitor_v(m)$ denote the result of 379

an observation for the machine state m . For instance, an attacker that only observes outputs enjoys the monitor function $\text{monitor}_v(m) \doteq O_v(m)$. Observations by an adversary may include explicit data produced by the VM's computational results intended for external consumption. Also the observations may include indirect information that can be inferred from the VM's operation, such as timing information, power consumption patterns, or memory access patterns. These observations require more sophisticated analysis and may reveal sensitive information without direct access to the VM's data.

VM Execution With an Attacker: An execution trace of the VM is an unbounded-length sequence of states, denoted by $\sigma = (m_0, m_1, \dots, m_n)$, satisfying the condition $\forall i. m_i \rightsquigarrow m_{i+1}$; here, $\sigma[i]$ refers to the i th element of the trace. Considering the ability of the attacker to pause and resume the VM at any time, we define the VM's execution as the sequence of states from σ where the VM is actively executing.

To identify when the VM is executing, we use the function $\text{curr}(m)$ to denote the current mode of the platform, with $\text{curr}(m) = v$ if the platform is executing the VM (v) in state m . Using this function, we can extract the steps in σ where the VM is executing, resulting in a subsequence $(m'_0, m'_1, \dots, m'_m)$ where $\text{init}(S_v(m'_0)) \wedge \forall i. \text{curr}(m'_i) = v$. This subsequence represents the VM's execution trace, including inputs, execution states, and outputs at each step. Given the VM's execution trace, the attacker may perform attack actions between any two consecutive steps, represented as $\forall i. (m'_i, m'_{i+1}) \in \text{attack}$. This action effectively introduces uncertainty in the VM's state and inputs, providing the VM with potentially fresh inputs at each step.

The semantics of a VM, denoted by $[v]$, is defined as the set of all possible finite or infinite execution traces, capturing every possible input sequence. Formally

$$[v] = \{(I_v(m'_0), S_v(m'_0), O_v(m'_0)), \dots | \text{init}(S_v(m_0))\}.$$

This model accounts for all potential input sequences because the VM may receive any value of input at any step. Furthermore, $[v]$ is prefix-closed, acknowledging that the attacker can pause and terminate the VM's execution at any time. The determinism of the VM's program means that a specific sequence of inputs uniquely identifies a trace from $[v]$ and determines the expected execution trace under that sequence of inputs.

Table I provides a summary of notation used in defining formal models for both VMs and adversary.

IV. FORMAL MODELING OF CONFIDENTIALITY AND INTEGRITY PROPERTIES

In this section, we provide the formal definition for confidentiality and integrity properties with respect to VM-based trusted execution.

Let $\lambda(v)$ denote the measurement of a VM instance v , computed upon its launch. This measurement process guarantees

$$\begin{aligned} &\forall m_1, m_2. \text{init}_{v_1}(S_{v_1}(m_1)) \wedge \text{init}_{v_2}(S_{v_2}(m_2)) \\ &\Rightarrow \lambda(v_1) = \lambda(v_2) \\ &\Leftrightarrow S_{v_1}(m_1) = S_{v_2}(m_2). \end{aligned}$$

TABLE I
TABLE OF NOTATIONS FOR DEFINING FORMAL MODELS FOR VMs AND ADVERSARY

Symbol	Description
v	A virtual machine instance.
$v.id$	Unique identifier for the VM.
$v.os$	The operating system image used by the VM.
$v.valist$	List of virtual addresses assigned to the VM.
$v.mem$	The allocated memory size for the VM.
$v.data$	The data and code pages within the VM.
$S_v(m)$	The state of the VM at a given moment m .
$I_v(m)$	Inputs to the VM at moment m , comprising external ($I_v^D(m)$) and internal ($I_v^R(m)$) inputs.
$O_v(m)$	Outputs from the VM at moment m .
\rightsquigarrow	The transition relation for the VM's execution.
$A_v(m)$	The state of the adversary with respect to VM v at moment m .
$\text{monitor}_v(m)$	The result of an adversary's observation at machine state m .
σ	An execution trace of the VM.
$\text{curr}(m)$	A function denoting the current execution mode of the platform at state m .
$[v]$	The semantics of a VM, representing the set of all possible execution traces.

This measurement process involves computing a cryptographic hash of the VM's initial content and configuration, providing a unique identity for the VM that serves as the basis for authenticating its legitimacy. This hash serves as a fingerprint of the VM at a particular point in time. The measurement process asserts that any two VM instances with the same measurement must have identical initial states, ensuring that any deviation from the expected VM program is detectable by the user. This assertion is based on the cryptographic property of collision resistance, which implies that it is computationally infeasible to find two distinct inputs (in this case, VM states) that result in the same hash output.

A. Confidentiality

Confidentiality ensures that a privileged software attacker cannot distinguish between the executions of two VMs, except for what is revealed through observable outputs. An attacker cannot gain information about the VM's execution state or internal processes beyond what is explicitly allowed through the monitoring function, denoted as monitor . This function provides all observations, including initial configurations, outputs to non-VM memory, and any potential side channel leakages. To formally assert the confidentiality guarantee, we propose the following:

$$\begin{aligned} &\forall \sigma_1, \sigma_2. \left(A_{v_1}(\sigma_1[0]) = A_{v_2}(\sigma_2[0]) \wedge \right. \\ &\forall i. (\text{curr}(\sigma_1[i]) = \text{curr}(\sigma_2[i]) \wedge I_{v_1}(\sigma_1[i]) = I_{v_2}(\sigma_2[i])) \wedge \\ &\forall i. (\text{curr}(\sigma_1[i]) = v) \Rightarrow \\ &\left. \text{monitor}_{v_1}(\sigma_1[i+1]) = \text{monitor}_{v_2}(\sigma_2[i+1]) \right) \\ &\Rightarrow \left(\forall i. A_{v_1}(\sigma_1[i]) = A_{v_2}(\sigma_2[i]) \right). \end{aligned}$$

This formulation implies that for any two traces, σ_1 and σ_2 , that exhibit equivalent attacker operations and observations (as permitted by monitor) but may differ in their private VM




Fig. 5. Overview of confidentiality property.




Fig. 6. Overview of integrity property.

states and internal executions, the observable outcome to the attacker must be identical. Here, $\sigma[i]$ means the i th index of the execution trace. The input that is responsible for i th state in the trace is denoted as $I(\sigma[i])$ and the corresponding output is denoted as $O(\sigma[i])$. By adhering to this model, a VM platform ensures that all potential traces of VM execution, which may yield the same observable outputs but originate from distinct internal states, remain indistinguishable to an external observer. This guarantees that the VM's confidentiality is preserved, preventing attackers from leveraging observable information to infer sensitive internal states or execution paths.

Fig. 5 shows the confidentiality property. Let us assume that the two traces start with an equivalent state and differ from state m_1 . This is because the two VMs can perform different computations. Adversary monitoring is assumed to be the same in both traces. Also, adversary actions are assumed to be the same in both traces. This should lead to the adversary state being identical in each step. The confidentiality property implies that the adversary state only depends on the adversary's actions and the initial state. Therefore, whatever the VM state is, it should not affect the adversary state. This shows that the adversary can only know information through the monitor function and not more.

B. Integrity

The integrity property states that the execution trace of the VM is solely determined by the sequence of inputs, independent of any interference by privileged software attackers beyond the provision of inputs. The integrity of a VM execution ensures that the VM's operational sequence and its resultant states and outputs are determined solely by its sequence of inputs. Operations by an attacker, such as manipulation of I/O peripherals or execution of privileged instructions, should not deviate the VM's execution from its intended path. The integrity property can be formalized as

$$\begin{aligned} & \forall \sigma_1, \sigma_2. \left(S_V(\sigma_1[0]) = S_V(\sigma_2[0]) \wedge \right. \\ & \quad \forall i. (\text{curr}(\sigma_1[i]) = V) \Leftrightarrow (\text{curr}(\sigma_2[i]) = V) \wedge \\ & \quad \left. \forall i. (\text{curr}(\sigma_1[i]) = V) \Rightarrow I_V(\sigma_1[i]) = I_V(\sigma_2[i]) \right) \\ & \Rightarrow \left(\forall i. S_V(\sigma_1[i]) = S_V(\sigma_2[i]) \wedge O_V(\sigma_1[i]) = O_V(\sigma_2[i]) \right). \end{aligned}$$

This states that if two execution traces, σ_1 and σ_2 , begin with identical initial states and receive the same sequence of inputs, then despite any differences in the attacker's operations

across the traces, the VM's state transitions and outputs will remain consistent across both traces.

This integrity model emphasizes a crucial aspect of VM security: the system's ability to maintain a predictable and reliable execution path, even when under adversarial influence. It ensures that the VM's computation integrity is preserved, thereby guaranteeing that the execution outcomes are solely the result of the provided inputs and the VM's deterministic behavior. The determinism and equivalence of VM execution can be formalized as

$$\begin{aligned} & \forall \sigma_1, \sigma_2. \left(S_{v1}(\sigma_1[0]) = S_{v2}(\sigma_2[0]) \wedge \right. \\ & \quad \forall i. (\text{curr}(\sigma_1[i]) = v1) \Leftrightarrow (\text{curr}(\sigma_2[i]) = v2) \wedge \\ & \quad \left. \forall i. (\text{curr}(\sigma_1[i]) = v1) \Rightarrow I_{v1}(\sigma_1[i]) = I_{v2}(\sigma_2[i]) \right) \\ & \Rightarrow \left(\forall i. S_{v1}(\sigma_1[i]) = S_{v2}(\sigma_2[i]) \wedge O_{v1}(\sigma_1[i]) = O_{v2}(\sigma_2[i]) \right). \end{aligned}$$

This formalism establishes that if two VM instances start with the same initial state and receive identical input sequences, then their execution traces, including state transitions and outputs, will be equivalent. This equivalence emphasizes the determinism property of the VM platform's execution model, ensuring that VM programs operate predictably and securely even in the presence of potential attackers.

Fig. 6 shows the integrity property. Actions by the adversary are marked as A_1 and A_2 . Let us assume that the VM's inputs and actions remain consistent across both traces. Similarly, the initial conditions of the VMs are identical. The adversary's actions are specified through the *attack* function, allowing for possible variations between the traces. The integrity verification necessitates demonstrating that the state and outputs of the VM remain unchanged in spite of these differences. The assumption that the adversary operates for an equal number of steps in both traces does not limit their capability, as any attack necessitating a variable number of steps across traces can be replicated within this model by extending the shorter trace of the adversary with a series of nonoperative steps. According to this theorem, under the specified assumptions, the state and outputs of the VM at every step are guaranteed to be the same across both traces.

V. ANALYSIS OF TDX ARCHITECTURE

This section provides a security analysis for data confidentiality and integrity of Intel TDX [4] architecture, which is an example of a VM-based TEE architecture.

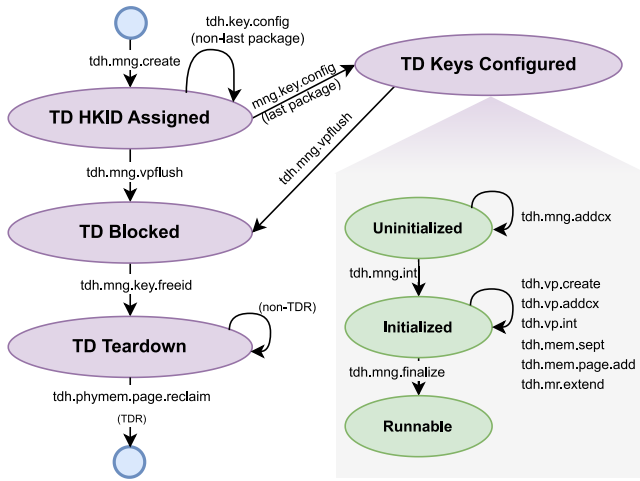


Fig. 7. TD life cycle state diagram with HKID states.

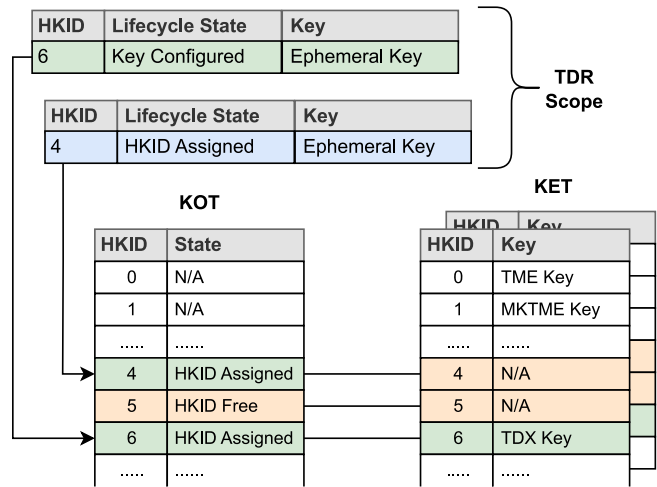


Fig. 8. Key management with TDR, KOT, and KET.

548 A. Intel TDX: Ensuring Data Confidentiality

549 Intel TDX safeguards the confidentiality of TD data across
 550 memory, the processor, and the bus by encrypting data dur-
 551 ing transmission from the processor back to memory. This
 552 encryption uses the MKTME system, employing AES-XTS
 553 with 128-bit encryption for each cache line. The unique keys
 554 for each TD, identifiable through a host KeyID (HKID), are
 555 generated and managed securely, with encryption keys stored
 556 internally and not disclosed to unauthorized entities.

557 Upon activation of TDX, physical memory is partitioned
 558 into secure (private) and normal (shared) regions, with the
 559 former designated for sensitive TD data and the latter for
 560 interactions with nontrusted entities. The allocation to either
 561 region is determined by the state of the highest order bit of
 562 the guest physical address (GPA), ensuring a clear separation
 563 and safeguarding of confidential data.

564 The life cycle of a TD includes several key states, from
 565 creation and key configuration to potential blocking and
 566 eventual teardown, each facilitated by specific API calls.
 567 This process begins with the creation of a new TD and the
 568 generation of an ephemeral key, followed by its configuration
 569 and operational management through the key encryption table
 570 (KET) and KeyID ownership table (KOT), ensuring secure and
 571 efficient key management throughout the TD’s existence.

572 In this section, we briefly describe the functionality of
 573 different HKID states in Fig. 7.

- 574 1) *HKID Assigned State*: Achievable through the
 575 *tdh_mng_create* API, this state marks the initialization
 576 of a new TD. Initially, the hypervisor ensures that any
 577 changes in the cache related to the TD’s physical pages
 578 are committed. Following this, it establishes the TD
 579 root (TDR) and creates a unique, temporary key for the
 580 TD. An HKID is generated and recorded in the KOT
 581 for each involved package.
- 582 2) *Keys Configured State*: This state occupies the majority
 583 of a TD’s operational lifespan. The TD’s temporary
 584 key is set up in the KET, with a secondary state
 585 machine managing the TD’s activities. The TD transi-
 586 tions through several substates: from “uninitialized” to

“initialized,” and finally to “runnable.” To incorporate
 the necessary TD control extension (TDCX) pages, the
tdh_mng_addcx API is utilized. The TD’s state within
 the TDR is initialized using *tdh_mng_init*, and achieving
 the runnable state is finalized with *tdh_mng_finalize*.

- 3) *Blocked State*: Any interruptions or faults prompt the
 TD to move into the blocked state, during which
 access to the TD’s private memory is suspended, and
 related caches are cleared. The *tdh_mng_vpflushdone*
 API checks for the complete flushing of cache lines
 associated with the TD’s address or HKID.
- 4) *Teardown State*: In this final phase, the host’s
 VMM reclaims the HKID and clears both the
 translation lookaside buffer (TLB) and cache. It
 proceeds to remove all private and control pages
 of the TD through *tdh_phymem_page_reclaim*, with
tdh_phymem_page_wbinvd being employed to ensure
 any modified cache lines are flushed.

These states highlight the dynamic and secure manage-
 ment of TDs within the TDX framework, emphasizing data
 confidentiality through stringent key control and memory
 encryption practices, as shown in Fig. 8.

609 B. Intel TDX: Guaranteeing Memory Integrity

610 Intel TDX maintains memory integrity via a dual approach,
 611 incorporating a TD owner bit and a message authentication
 612 code (MAC), both embedded within ECC memory. A 128-bit
 613 MAC key is created during system initialization, with a 28-bit
 614 MAC generated for each memory write. This MAC, alongside
 615 the TD owner bit, aids in verifying data integrity during reads,
 616 with discrepancies indicating potential integrity breaches and
 617 resulting in the marking of compromised cache lines.

618 The TD owner bit serves as a gatekeeper, controlling access
 619 based on whether a physical address is associated with a
 620 private HKID. This mechanism ensures that only authorized
 621 SEAM mode operations can access secured memory segments,
 622 with all other requests being denied and returned as null,
 623 thereby preserving the integrity of sensitive data.

```

(define sec_ept (make-hash))
(define-struct sec_ept_entry (hpa gpa_shared state))
(define/contract sec_ept-contract
  (hash/c integer? sec_ept_entry? #:flat? #t)
  sec_ept)

```

Listing 1. EPT.

```

(define KET (make-hash))
(define/contract KET-contract
  (hash/c integer? bitvector? #:flat? #t)
  KET)

(define KOT (make-hash))
(define/contract KOT-contract
  (hash/c integer? integer? #:flat? #t)
  KOT)

```

Listing 2. KET and key ownership table.

```

(define (TDH_MNG_CREATE hpa HKID)
  (define hkid_state (hash-ref KOT HKID #f))
  (define page_state (hash-ref PAMT hpa #f))
  (when (and (is_hkid_private HKID)
             (or (equal? hkid_state #f)
                 (equal? hkid_state HKID_FREE))
             (or (equal? page_state #f)
                 (equal? page_state PT_NDA))))
  (begin
    (hash-set! KOT HKID HKID_ASSIGNED)
    (hash-set! PAMT hpa
              (make-PAMT_entry PT_TDR 0 0))
    (make-TDR #f #f 0 0
              HKID_ASSIGNED HKID 0 #f #f)))

```

Listing 3. TDH_MNG_CREATE ABI.

624 Furthermore, any attempt to write to a protected memory
625 segment outside of SEAM mode triggers the reset of the cor-
626 responding TD owner bit, marking the segment as poisoned.
627 This serves as a critical fail-safe, triggering a TD exit and,
628 if necessary, transitioning the TD to a fatal state for security,
629 thereby emphasizing the robust measures in place to maintain
630 memory integrity within the TDX architecture.

631 VI. FORMAL MODELING OF INTEL TDX ARCHITECTURE

632 The formal model is developed following the Intel TDX
633 module specification. We model the Intel TDX architec-
634 ture using Rosette [18] to enable symbolic simulation of
635 the complex mechanisms, including TDX tables, application
636 binary interfaces (ABIs), and other configurations essential
637 for ensuring memory confidentiality and integrity within TDs.
638 This section details the formal modeling, highlighting only the
639 key components that form the backbone of TDX security.

640 A. Defining TDX Tables

641 The TDX specification has various tables, each serving a
642 unique purpose in the security architecture. Among these, the
643 extended page table (EPT), KET, and KOT are foundational
644 elements for confidentiality.

645 1) *Extended Page Table*: The EPT maps GPAs to host
646 physical addresses (HPAs) and maintains the page state,
647 incorporating a shared bit to differentiate between secure and
648 shared memory spaces. This mapping is crucial for memory
649 isolation and confidentiality.

650 Listing 1 shows a hash table `sec_ept` created using `make-`
651 `hash`, which serves as a repository for managing EPT entries,
652 and a custom-defined structure `sec_ept_entry`, which keep
653 track of the essential attributes of each entry, including HPA,
654 GPA shared status (`gpa_shared`), and the entry's current
655 state (`state`). This approach enables efficient tracking and
656 manipulation of memory addresses between the host and VMs,
657 facilitating a streamlined mechanism to oversee the shared or
658 exclusive access to physical memory resources.

659 2) *Key Encryption Table and KeyID Ownership Table*: The
660 KET (Listing 2) associates each TD's ephemeral encryption
661 key with its corresponding HKID, playing a pivotal role in
662 encrypting memory access and safeguarding data in transit.
663 The KOT, on the other hand, tracks the lifecycle state of each
664 HKID, ensuring proper key management and assignment. Two
665 hash table structures are used to represent the two tables.

666 B. Trust Domain Management

667 The management of TDs includes TD creation, key config-
668 uration, handling exceptions, and teardown. We implemented
669 structures to facilitate this, the TDR and TD control structure

(TDCS), which maintain state and control information for
each TD.

670
671
672 1) *TD Creation*: TDs are instantiated and assigned unique
673 HKIDs, with their ephemeral keys generated and stored
674 securely. This process involves interactions with the physical
675 address metadata table (PAMT) for memory allocation and the
676 cache to ensure confidentiality during TD operations. Listing 3
677 shows the ABI for `TDH_MNG_CREATE`, which manages
678 creating and initializing a transactional data handler (TDH)
679 by mapping `hpa` to HKID. Initially, it checks the current state
680 of the given HKID and the `hpa` in two hash tables (KOT
681 for HKIDs and PAMT for physical addresses) to determine
682 if the HKID is private, not yet assigned, or if the hardware
683 page is not yet allocated or is in a nondisclosure agreement
684 state (`PT_NDA`). If these conditions are met, the function
685 proceeds to mark the HKID as assigned (`HKID_ASSIGNED`)
686 in the KOT table and creates a new PAMT entry with initial
687 parameters. Finally, it initializes a new TDR with default or
688 initial values. This setup indicates a mechanism for managing
689 access and operations on hardware resources, ensuring data
690 privacy and integrity through proper handling of hardware keys
691 and memory pages.

692 2) *Key Configuration*: `TDH_MNG_KEY_CONFIG`
693 (Listing 4) is designed to configure keys for a TDR associated
694 with a specific `hpa`. It begins by retrieving the current entry
695 for the given PAMT and determining its state, specifically
696 checking if it matches the expected type for transactional data
697 records (`PT_TDR`). It then checks the `tdr` for a fatal error
698 condition (`td_fatal`) and its lifecycle state to ensure it is in
699 the `HKID_ASSIGNED` state, indicating that a HKID has been
700 assigned but not yet configured with keys.

701 If the page is correctly prepared for transactional data,
702 no fatal errors are present, and the TDH is ready for key
703 configuration, the function proceeds to update KET with the


```

(define (TDH_MNG_KEY_CONFIG hpa tdr)
  (define page_entry (hash-ref PAMT hpa #f))
  (define page_state
    (if (PAMT_entry? page_entry)
        (PAMT_entry-PAGE_TYPE page_entry) #f))
  (define td_fatal (TDR-FATAL tdr))
  (define hkid_state (TDR-LIFECYCLE_STATE tdr))
  (when (and (equal? page_state PT_TDR)
             (not td_fatal) (equal? hkid_state HKID_ASSIGNED)))
  (begin
    (hash-set! KET (TDR-HKID tdr) key_val)
    (struct-copy TDR tdr
      [LIFECYCLE_STATE KEYS_CONFIGURED])))

```

Listing 4. TDH_MNG_KEY_CONFIG ABI.

```

(define (TDH_MNG_VPFLUSH pa tdr)
  (define page_entry (hash-ref PAMT pa #f))
  (define page_state
    (if (PAMT_entry? page_entry)
        (PAMT_entry-PAGE_TYPE page_entry) #f))
  (define td_state (TDR-LIFECYCLE_STATE tdr))
  (define hkid_state (hash-ref KOT (TDR-HKID tdr)))
  (if (and (equal? page_state PT_TDR)
           (or (equal? td_state TD_HKID_ASSIGNED)
               (equal? td_state TD_KEYS_CONFIGURED))
           (equal? hkid_state HKID_ASSIGNED)))
      (begin
        (flush_cache (TDR-HKID tdr))
        (hash-set! KOT (TDR-HKID tdr)
          HKID_FLUSHED)(struct-copy TDR tdr
            [LIFECYCLE_STATE TD_BLOCKED]))#f))

```

Listing 5. TDH_MNG_VPFLUSH ABI.

```

(define (TDH_MNG_KEY_FREEID pa tdr)
  (define page_entry (hash-ref PAMT pa #f))
  (define page_state
    (if (PAMT_entry? page_entry)
        (PAMT_entry-PAGE_TYPE page_entry)
        #f))
  (define hkid_state (TDR-LIFECYCLE_STATE tdr))
  (define hkid (TDR-HKID tdr))
  (define kot_entry (hash-ref KOT hkid #f))
  (when (and (equal? page_state PT_TDR)
             (equal? hkid_state TD_BLOCKED)
             (equal? kot_entry HKID_FLUSHED)))
  (begin (hash-set! KOT hkid HKID_FREE)
    (struct-copy TDR tdr
      [LIFECYCLE_STATE TD_TEARDOWN]
      [HKID 0])))

```

Listing 6. TDH_MNG_KEY_FREEID ABI.

state (*TD_BLOCKED*), and that the HKID has been flushed (*HKID_FLUSHED*). This ensures the function operates under safe conditions where the data associated with the *tdr* and HKID has been securely managed and is ready for cleanup. Upon confirming these prerequisites, the function sets the HKID’s state to *HKID_FREE* in the KOT, marking it available for future assignments. Additionally, it updates the *tdr* to reflect a teardown lifecycle state (*TD_TEARDOWN*) and resets the HKID within the *tdr*, effectively clearing the association and preparing the system for new transactions. This process is essential for the secure and efficient reuse of hardware resources, ensuring that data integrity and confidentiality are maintained throughout the lifecycle of a TD.

hardware key ID extracted from TDR and sets a new key (*key_val*). It then updates the *tdr* structure itself to reflect that the keys have been configured, changing its lifecycle state to *KEYS_CONFIGURED*.

3) *Handling Interrupts and Exceptions*: Handling of interrupts and exceptions is crucial for the secure and stable operation of TDs. This involves saving the current TD state, scrubbing the VCPU state, and executing a cache flush to maintain data integrity.

The function *TDH_MNG_VPFLUSH* (Listing 5) is responsible for securely flushing a virtual page from the cache. The function starts by looking up the state of the page associated with the given physical address in PAMT. It assesses the lifecycle state of the *tdr* to ensure it is either in the *HKID_ASSIGNED* or *KEYS_CONFIGURED* state, and verifies that the HKID related to the *tdr* is marked as assigned in KOT. If these conditions are met, the *tdr* is in an appropriate state for flushing. This is critical for maintaining data consistency and security, ensuring that no sensitive data remains in the cache that could be accessed inappropriately. After flushing the cache, the function updates the state of the HKID in the KOT to *HKID_FLUSHED*, indicating that the flush operation has been completed. Finally, it updates the lifecycle state of the *tdr* to *TD_BLOCKED*, indicating that the *tdr* is in a state where it cannot perform regular operations.

4) *TD Teardown: TDH_MNG_KEY_FREEID* (Listing 6) function is designed for releasing or freeing HKIDs that are no longer in use.

It first verifies that the specified physical address is associated with a page prepared, that the TDR is in a blocked

VII. FORMAL MODELING OF CACHE FOR TDX SYSTEMS

When the cache is unencrypted, the data stored within remains in plaintext, posing significant risks to both data integrity and confidentiality. Such a scenario lays the groundwork for multiple security vulnerabilities, as unauthorized access to this unencrypted data can lead to information leakage or manipulation. In this section, we evaluate how the Intel TDX module addresses these critical security concerns within the context of a shared cache environment. This ensures that even in a shared cache scenario, where multiple processes or VMs might access the same physical cache resources, data remains secure, isolated, and impervious to unauthorized access or tampering, thereby upholding the highest standards of integrity and confidentiality. We model and formally evaluate two distinct cache types, each capable of enhancing security in TDX environments independently: 1) HKID-tagged cache and 2) TD-owner-bit cache.

A. Basic Cache Structure and Initialization

The cache model in Listing 7 is designed to simulate a 4-way associative cache, a common setup in modern computing systems. This setup is characterized by a finite number of cache sets, each with multiple ways to store data. The model initializes hash maps to track the validity, tag, data, and HKID of each cache line, providing a foundational structure for simulating cache operations.

The *init-cache* function (Listing 8) populates these structures, initially setting all cache lines to an invalid state and

```
(define kmax-cache-set-index-t 256)
(define kmax-cache-way-index-t 4)

(define cache-valid-map (make-hash))
(define cache-tag-map (make-hash))
(define cache-data-map (make-hash))
(define cache-hkid-map (make-hash))
```

Listing 7. Cache configuration.

```
(define (init-cache)
  (for ([i (in-range kmax-cache-set-index-t)])
    (for ([j (in-range kmax-cache-way-index-t)])
      (let ([key (cons i j)])
        (hash-set! cache-valid-map key #false)
        (hash-set! cache-data-map key 0)
        (hash-set! cache-tag-map key 0)
        (hash-set! cache-hkid-map key 0))))))
```

Listing 8. Cache initialization.

```
(define (query-cache pa repl-way hkid)
  (define set (paddr2set pa))
  (define tag (paddr2tag pa))
  (define hit-way
    (for/or ([way (in-range kmax-cache-way-index-t)])
      (let ([key (cons set way)])
        (and (hash-ref cache-valid-map key #false)
              (= (hash-ref cache-tag-map key) tag)
              (= (hash-ref cache-hkid-map key) hkid)
              way))))))
  (if hit-way
    (let ([key (cons set hit-way)])
      (values #true hit-way (hash-ref
                            cache-data-map key))))
    (begin
      (let ([key (cons set repl-way)])
        (hash-set! cache-valid-map key #true)
        (hash-set! cache-tag-map key tag)
        (hash-set! cache-data-map key 0)
        (hash-set! cache-hkid-map key hkid)
        (values #false repl-way 0))))))
```

Listing 9. HKID tagged cache model.

774 assigning default values to the tags, data, and HKIDs. This
775 ensures a clean state from which cache operations can start.

776 B. HKID Tagged Cache

777 Integrating HKID into the cache model adds a layer of
778 security by ensuring that cache lines are accessible only by the
779 appropriate TD. Listing 9 shows our modeling of HKID tagged
780 cache. This approach leverages HKIDs to tag cache lines, thus
781 facilitating the validation of access requests based on the TD’s
782 identity. HKID tagging is implemented by extending the cache
783 model to include a mapping of cache lines to HKIDs. This
784 extension allows the cache to check not only the validity and
785 tag match for cache hits but also the HKID, ensuring that only
786 requests from the owning TD can access the cached data.

787 C. TD Owner Bit

788 By using a TD owner bit in access control, TDX enforces
789 strict access policies, allowing only SEAM mode processes
790 to read secure cache lines, thereby significantly mitigating
791 the risk of unauthorized data access. Listing 10 shows our
792 modeling of TD owner bit cache. This cache management

```
(define (query-cache pa repl-way hkid seam-mode?)
  (define set (paddr2set pa))
  (define tag (paddr2tag pa))
  (define hit-way
    (for/or ([way (in-range kmax-cache-way-index-t)])
      (let ([key (cons set way)])
        (and (hash-ref cache-valid-map key #false)
              (= (hash-ref cache-tag-map key) tag)
              (= (hash-ref cache-hkid-map key) hkid)
              (or seam-mode? (not (hash-ref
                                    cache-td-owner-map key))))
              way))))))
  (if hit-way
    (let ([key (cons set hit-way)])
      (if (or seam-mode? (not (hash-ref
                              cache-td-owner-map key)))
          (values #true hit-way (hash-ref
                                cache-data-map key))
          (values #true hit-way 0)))
      (begin
        (let ([key (cons set repl-way)])
          (hash-set! cache-valid-map key #true)
          (hash-set! cache-tag-map key tag)
          (hash-set! cache-data-map key 0)
          (hash-set! cache-hkid-map key hkid)
          (hash-set! cache-td-owner-map key
                    (if (> hkid 0) #true #false)))
          (values #false repl-way 0))))))
```

Listing 10. TD owner bit cache model.

strategy not only enhances data security by providing fine- 793
grained access control based on hardware-level identifiers but 794
also introduces a flexible framework for managing cache data 795
across multiple TDs. 796

VIII. EXPERIMENTS 797

This section demonstrates the effectiveness of our proposed 798
VM-based TEE verification framework to verify the Intel TDX 799
module. First, we describe our formal verification setup. Next, 800
we present the formal verification results of our framework. 801

A. Experimental Setup 802

Our model for Intel TDX and caches and properties for 803
formal security verification is constructed using Rosette [18] 804
formal verification language. We used the publicly avail- 805
able specification as well as implementation for the TDX 806
module [4] to derive the formal model. The Rosette model 807
has assertions, symbolic variables, and solver-aided functions. 808
The correctness of these elements is verified using Rosette’s 809
symbolic execution engine, which internally uses Z3 [19] SMT 810
solver to check the feasibility of paths and the satisfaction 811
of constraints. We ran our experiments on Intel i7-5500U @ 812
3.0GHz CPU with 16GB RAM machine. We have developed 813
15 confidentiality properties and 9 integrity properties for two 814
cache models: 1) HKID-tagged and 2) TD-owner-bit-tagged 815
cache. 816

B. Generation of Confidentiality Properties 817

The properties for confidentiality and integrity are defined 818
based on the threat model outlined in the TDX specification. 819
We have developed 15 confidentiality properties. 820

- 1) cP_1 : Assert that any GPA mapped to a specific HPA 821
within the secure EPT maintains confidentiality, mean- 822
ing no other GPA can map to this HPA. 823

- 824 2) cP_2 : Assert that the ephemeral encryption key associated
825 with a specific HKID in the KET table remains confi-
826 dential and is not leaked or accessible to unauthorized
827 entities.
- 828 3) cP_3 : Assert that once an HKID is assigned, its state
829 remains confidential and accurately reflects its assigned
830 status within the KOT table.
- 831 4) cP_4 : Assert that the lifecycle state of a TDR remains
832 confidential and can only be one of the predefined
833 states (INIT, FATAL, RUNNING), safeguarding the state
834 transitions from unauthorized access.
- 835 5) cP_5 : Assert that the page state of any entry in the
836 secure EPT is limited to predefined states, protecting the
837 confidentiality of page mappings.
- 838 6) cP_6 : Assert that the key configuration state for a given
839 HKID remains confidential and accurately reflects the
840 $TD_KEYS_CONFIGURED$ state, protecting the key con-
841 figuration status from unauthorized changes.
- 842 7) cP_7 : Assert that the finalization status of a TDCS
843 remains confidential and is always set to true after
844 finalizing.
- 845 8) cP_8 : Assert that the confidentiality of the shared bit
846 status for any given entry in the secure EPT.
- 847 9) cP_9 : Assert that the package configuration bitmap of a
848 TDR remains confidential, ensuring that the configura-
849 tion details are protected from unauthorized disclosure.
- 850 10) cP_{10} : Assert that the association between a VCPU and
851 its corresponding HKID is kept confidential.
- 852 11) cP_{11} : Assert that querying the cache with an incorrect
853 HKID results in a cache miss.
- 854 12) cP_{12} : Assert that after querying with the correct HKID,
855 a subsequent query with the same HKID and address
856 results in a cache hit.
- 857 13) cP_{13} : Assert that querying with a different HKID
858 (assuming unauthorized access) after a cache line is
859 populated does not provide access to the data.
- 860 14) cP_{14} : Assert that after updating a cache line with a new
861 HKID and data, the previous HKID no longer has access.
- 862 15) cP_{15} : Assert that once data is written to a cache line, it
863 remains unchanged unless explicitly modified through a
864 valid cache update.

865 C. Generation of Integrity Properties

866 We have developed nine integrity properties.

- 867 1) iP_1 : Assert that the integrity of the EPT mappings by
868 asserting that any entry mapping a GPA to a HPA cannot
869 be in a “blocked” state.
- 870 2) iP_2 : Assert that the integrity of the TDR lifecycle by
871 asserting that once a TDR is finalized, its lifecycle state
872 cannot be “INIT” or “FATAL.”
- 873 3) iP_3 : Assert that the integrity of key state transitions
874 within the TDX module by providing consistent transi-
875 tions for a HKID based on its current state. Specifically,
876 it asserts that an HKID assigned state can only move
877 to the keys configured state, a keys configured state can
878 transition to either blocked or teardown, and a blocked
879 state can only move to teardown.

```
(define-symbolic pa1 pa2 bv28)
(define-symbolic repl-way1 repl-way2 integer?)
(define-symbolic hkid1 hkid2 integer?)

(define-values (hit1 way1 data1)
  (query-cache pa1 repl-way hkid1))
(define-values (hit2 way2 data2)
  (query-cache pa1 repl-way hkid2))

(define confidentiality-assertion
  (assert (or (not (and hit1 hit2)) (= hkid1 hkid2))))

(define result (solve (confidentiality-assertion)))
(displayln result)
```

Listing 11. Sample confidentiality assertion.

- 4) iP_4 : Assert that if a cache entry is marked as valid, it
880 must have a corresponding tag and data in the cache. 881
- 5) iP_5 : Asserts that within a single set in a set-associative
882 cache, all valid entries must have unique tags. 883
- 6) iP_6 : Asserts that each valid cache entry, the correspond-
884 ing HKID is correctly mapped to the same set and way
885 in the cache-hkid-map. 886
- 7) iP_7 : Assert that any cache entry from SEAM mode
887 marked as valid has a corresponding and correct TD
888 owner bit set. 889
- 8) iP_8 : Assert that for any two valid cache entries in
890 the same set but in different ways, their tags must be
891 different. 892
- 9) iP_9 : Assert that if two cache entries have the same
893 tag and are valid, they must have the same TD owner
894 bit. 895

Listing 11 shows sample confidentiality property (cP_{13})
896 of a cache system through symbolic execution. Initially,
897 symbolic variables $pa1$ and $pa2$ with a bit-vector size of 28
898 ($bv28$) representing physical addresses are initialized. Then
899 symbolic integers $repl-way1$, $reply-way2$, $hkid1$, and $hkid2$ are
900 introduced, with the latter two representing replacement cache
901 element. Sample assertion queries the cache twice, using the
902 same physical address ($pa1$) but different key identifiers ($hkid1$
903 and $hkid2$), and stores the results (hit flags, ways, and data)
904 in ($hit1$, $way1$, $data1$) and ($hit2$, $way2$, $data2$), respectively.
905 The confidentiality assertion checks if, $hkid1$ and $hkid2$ are
906 different, both cannot have cache hits for the same physical
907 address that could violate confidentiality. The assertion is
908 then solved, and the result is displayed, indicating whether
909 the cache system maintains confidentiality across the given
910 symbolic inputs. 911

D. Verification Results

Table II provides a summary of the verification outcomes
913 for three distinct models: 1) the TDX module; 2) HKID-
914 tagged-cache; and 3) TD-owner-bit-cache. It details the
915 number of lines of code in each model, with the TDX
916 module being the largest at 400 lines, and the HKID-tagged-
917 cache the smallest at 100 lines. The table also indicates
918 the number of confidentiality and integrity properties verified
919 for each module. Verification time, measured in seconds,
920 showcases the efficiency of the verification process for each
921

TABLE II
ROSETTE MODELS AND VERIFICATION RESULTS

Description	# Lines	Confidentiality	Integrity	Verification Time (s)
TDX Module	400	10 properties ($cP_1 - cP_{10}$)	3 properties ($iP_1 - iP_3$)	12.86
HKID-tagged cache	100	5 properties ($cP_{11} - cP_{15}$)	3 properties ($iP_4 - iP_6$)	5.37
TD-Owner-bit cache	150	5 properties ($cP_{11} - cP_{15}$)	6 properties ($iP_4 - iP_9$)	7.92
Total	650	15	9	26.15

922 model, demonstrating the practicality and scalability of the
923 verification process in evaluating the reliability and robustness
924 of the models.

925 Our work can be extended to other VM-based solutions.
926 For example, if we consider AMD SEV, which uses x86
927 architecture similar to Intel TDX, the changes needed are
928 minimal. Similarly, AMD SEV uses an VM address space
929 identifier (ASID) to uniquely identify the VM addresses,
930 which is similar to HKID used by Intel TDX. To extend our
931 formal model to AMD SEV, we need to model the ASID,
932 but most of the formal model of Intel TDX module can be
933 reused.

934 IX. CONCLUSION

935 This article has presented a comprehensive framework for
936 the formal verification of VM-based TEEs, addressing the
937 critical need for robust security mechanisms in the face
938 of evolving threats. We have developed a formalization of
939 confidentiality and integrity for confidential VMs, proposing
940 a secure and verifiable model in the context of powerful
941 adversaries. Our contributions, including the formalization of
942 a confidential VM, the establishment of formal definitions for
943 confidentiality and integrity within VM-based TEEs, and the
944 development of a refinement-based methodology, underline the
945 importance and effectiveness of formal verification in ensuring
946 the security of VM-based TEEs. Our experimental results
947 demonstrate the applicability and resilience of our framework
948 to analyze sophisticated attack scenarios, highlighting its
949 potential to significantly enhance the security posture. By
950 proving the confidentiality and integrity guarantees of the Intel
951 TDX platform through machine-checked proofs, we not only
952 validate our approach but also pave the way for future research
953 in securing virtualized TEE environments.

REFERENCES

- 954
- [1] H. Witharana, D. Chatterjee, and P. Mishra, "Verifying memory confi- 955
956 dentiality and integrity of Intel TDX trusted execution environments,"
957 in *Proc. IEEE Int. Symp. Hardw. Orient. Security Trust (HOST)*, 2024,
958 pp. 44–54.
 - [2] "Intel software guard extensions (SGX)," [Online]. Available: [https:// 959
960 www.intel.com/content/www/us/en/developer/tools/software-guard-
961 extensions/overview.html](https://www.intel.com/content/www/us/en/developer/tools/software-guard-extensions/overview.html)
 - [3] V. Costan, I. A. Lebedev, and S. Devadas, "Sanctum: Minimal hardware 962
963 extensions for strong software isolation," in *Proc. USENIX Secur. Symp.*,
964 2016, pp. 857–874.
 - [4] "Intel trust domain extensions (TDX)." [Online]. Available: [https:// 965
966 www.intel.com/content/www/us/en/developer/articles/technical/intel-
967 trust-domain-extensions.html](https://www.intel.com/content/www/us/en/developer/articles/technical/intel-trust-domain-extensions.html).
 - [5] "AMD secure encrypted virtualization (SEV)." 2023. [Online]. 968
969 Available: <https://developer.amd.com/sev/>
 - [6] D. Lee, D. Kohlbrenner, S. Shinde, K. Asanović, and D. Song, 970
971 "Keystone: An open framework for architecting trusted execution envi-
972 ronments," in *Proc. 15th Eur. Conf. Comput. Syst.*, 2020, pp. 1–16.
 - [7] (Intel, Santa Clara, CA, USA). *Intel Trust Domain Extensions (TDX)* 973
974 *Security Review*. 2023. [Online]. Available: [https://services.google.com/ 975
976 fh/files/misc/intel_tdx_-_full_report_041423.pdf](https://services.google.com/fh/files/misc/intel_tdx_-_full_report_041423.pdf)
 - [8] "AMD secure processor for confidential computing 977
978 security review." 2023. [Online]. Available:
979 [https://storage.googleapis.com/gweb-uniblog-publish- 979
980 prod/documents/AMD_GPZ-Technical_Report_FINAL_05_2022.pdf](https://storage.googleapis.com/gweb-uniblog-publish-prod/documents/AMD_GPZ-Technical_Report_FINAL_05_2022.pdf)
 - [9] H. Witharana, Y. Lyu, and P. Mishra, "Directed test generation for 980
981 activation of security assertions in RTL models," *ACM Trans. Design*
982 *Autom. Electron. Syst.*, vol. 26, no. 4, pp. 1–28, 2021.
 - [10] H. Witharana, Y. Lyu, S. Charles, and P. Mishra, "A survey on assertion- 983
984 based hardware verification," *ACM Comput. Surveys*, vol. 54, no. 11,
985 pp. 1–33, 2022.
 - [11] J. Hu et al., "ProveriT: A parameterized, composable, and verified model 986
987 of TEE protection profile," *IEEE Trans. Depend. Secure Comput.*, early
988 access, Mar. 11, 2024, doi: [0.1109/TDSC.2024.3375311](https://doi.org/10.1109/TDSC.2024.3375311).
 - [12] Y. Ma, Q. Zhang, S. Zhao, G. Wang, X. Li, and Z. Shi, 989
990 "Formal verification of memory isolation for the trustzone-based
991 TEE," in *Proc. 27th Asia-Pacific Softw. Eng. Conf. (APSEC)*, 2020,
992 pp. 149–158.
 - [13] H. Sun and H. Lei, "A design and verification methodology for 993
994 a trustzone trusted execution environment," *IEEE Access*, vol. 8,
995 pp. 33870–33883, 2020.
 - [14] M. U. Sardar, S. Musaev, and C. Fetzer, "Demystifying attestation in 996
997 Intel trust domain extensions via formal verification," *IEEE Access*,
998 vol. 9, pp. 83067–83079, 2021.
 - [15] M. U. Sardar, T. Fossati, and S. Frost, "Comprehensive specification 999
1000 and formal analysis of attestation mechanisms in confidential comput-
1001 ing," in *Proc. ICE*, 2023, pp. 90–91.
 - [16] W. Ozga, "Towards a formally verified security monitor for VM-based 1002
1003 confidential computing," in *Proc. 12th Int. Workshop Hardw. Archit.*
1004 *Support Secur. Privacy*, 2023, pp. 73–81.
 - [17] P. Subramanyan, R. Sinha, I. Lebedev, S. Devadas, and S. A. Seshia, 1005
1006 "A formal foundation for secure remote execution of enclaves," in *Proc.*
1007 *ACM SIGSAC Conf. Comput. Commun. Secur.*, 2017, pp. 2435–2450.
 - [18] "Rosette language guide." [Online]. Available: [https://docs.racket- 1008
1009 lang.org/rosette-guide/index.html](https://docs.racket-lang.org/rosette-guide/index.html)
 - [19] L. De Moura and N. Bjørner, "Z3: An efficient SMT solver," in *Proc.* 1010
1011 *Int. Conf. Tools Algorithms Construct. Anal. Syst.*, 2008, pp. 337–340.