

A Logic of Secure Systems and its Application to Trusted Computing

Anupam Datta
danupam@cmu.edu

Jason Franklin
jfrankli@cs.cmu.edu

Deepak Garg
dg@cs.cmu.edu

Dilsun Kaynar
dilsun@cs.cmu.edu

Abstract

We present a logic for reasoning about properties of secure systems. The logic is built around a concurrent programming language with constructs for modeling machines with shared memory, a simple form of access control on memory, machine resets, cryptographic operations, network communication, and dynamically loading and executing unknown (and potentially untrusted) code. The adversary’s capabilities are constrained by the system interface as defined in the programming model (leading to the name CSI-ADVERSARY). We develop a sound proof system for reasoning about programs without explicitly reasoning about adversary actions. We use the logic to characterize trusted computing primitives and prove code integrity and execution integrity properties of two remote attestation protocols. The proofs make precise assumptions needed for the security of these protocols and reveal an insecure interaction between the two protocols.

1 Introduction

Contemporary secure systems are complex and designed to provide subtle security properties in the face of attack. Examples of such systems include hypervisors, virtual machine monitors, security kernels, operating systems, web browsers, and secure co-processor-based systems such as those utilizing the Trusted Computing Group’s Trusted Platform Module (TPM) [38]. In this paper we initiate a program to formally model abstractions of such systems and specify and analyze their security properties in the presence of a general class of adversaries. Specifically, we introduce the Logic of Secure Systems (LS^2) and use it to carry out a detailed analysis of Trusted Computing systems. The logic is built around a programming language for modeling systems and is inspired by a logic for network protocol analysis, Protocol Composition Logic (PCL) [10, 11, 13, 33].

Programming Model. The programming language is designed to be expressive enough to model practical secure systems while still maintaining a sufficiently high level of abstraction to enable simple reasoning. Following PCL, the language includes process calculi and functional constructs for modeling cryptographic operations, straightline code, and network communication. We introduce constructs for modeling machines and shared memory, a simple form of access control on memory, machine resets, and dynamically loading and executing unknown (and potentially untrusted) code. The primitives for reading and writing to memory are inspired by the treatment of memory cells in impure functional languages like Standard ML [26]. We model memory protection, a fundamental building block for secure systems [34], by allowing programs to acquire exclusive-write locks on memory locations. The treatment of dynamically loading and executing unknown code is novel to this work.

While these constructs are the common denominator for many secure systems, including the trusted computing systems examined in this paper, they are by no means sufficient to model all systems of interest. The language, however, is *extensible* in a modular fashion, as we illustrate by extending the core language (presented in Section 2) with a trusted computing subsystem (in Section 3). At a high level, each system component can be viewed as exposing an interface. For example, the interface for memory includes read, write, and reset operations. Adding a new component to the system involves adding operations in the programming language corresponding to the interface exposed by it. Platform Configuration Registers (PCR) in the TPM are an example since they can be modeled as a special form of memory that may be accessed via read, reset, and a new extend operation. Some extensions can have a more global effect on the language semantics. For instance, adding the reset operation to the language affects both how state of local memory and TPM PCRs may be updated.

Interfaces to system components also provide a use-

ful conceptual view of the adversary. Since the capabilities of the adversary are constrained by the system interface, we refer to her as a CSI-ADVERSARY. For example, the adversary can write to unprotected memory locations, but can only update PCRs through the extend operation in its interface. Formally, the adversary may execute any program expressible in our programming model, i.e. the adversary can perform symbolic cryptographic operations, intercept messages on the network, inject messages that it can create, read and write memory locations that are not explicitly locked by another thread, and reset machines. Because of these capabilities, the adversary can launch a broad range of attacks on the network and the local machines including replay attacks, modifying and injecting malicious code on local machines, and exploiting race conditions to compromise systems.

Logic. Security properties of programs are expressed in LS^2 using modal formulas of the form $[P]_I^{t_b, t_e} A$, which means that formula A holds whenever thread I executes exactly the program P in the time interval $(t_b, t_e]$, irrespective of the actions executed concurrently by other threads including the adversary. The thread I identifies the principal executing the program, the machine on which the program is being executed, and includes a unique identifier. The formula A expresses security properties, such as confidentiality, integrity, authentication, as well as code and execution integrity. The logic includes predicates that reflect the programming language constructs for shared memory, memory protection, machine resets and a form of unconditional jump to model branching to dynamically loaded code.

Security properties are established using a proof system for LS^2 . A central design goal that LS^2 achieves (following PCL) is that *the proof system does not mention adversary actions*. Instead, the semantics and soundness of the proof system guarantee that if $[P]_I^{t_b, t_e} A$ is provable, then A holds in all traces in which I completes execution of program P , including those that contain adversarial threads. This implicit treatment of adversaries simplifies proofs significantly. Designing a sound proof system that supports this local style of reasoning, in spite of the global nature of shared memory changes and execution of dynamically loaded unknown code, turned out to be a significant technical challenge.

We formalize local reasoning principles about shared memory with axioms that reason about invariance of values in memory based on local actions of threads that hold locks (see Section 2). This approach is technically similar to concurrent separation logic, whose regions resemble LS^2 's locks [6], but distinct from formal systems

which support global reasoning about concurrent shared memory programs [23]. Our initial idea to reason about execution of dynamically loaded code was to treat the code being branched to as a continuation of the code calling it. However, this approach does not work for the case where the code being branched to is either read from memory or received over the network, because nothing can be determined about the called code by looking at the caller's program. As a result, traditional methods for proving program invariants such as those based on Hoare logic and its extensions [17, 29, 32] do not apply to this setting. Yet this is exactly what we needed to reason in the face of adversaries who can modify or inject code into the system. Our final technical approach for reasoning about execution of dynamically loaded code is based on a program invariance rule, which we elaborate on in Section 2 and illustrate in Section 4.1.

Trusted Computing. We model and analyze two trusted computing protocols that rely on TPMs to provide integrity properties: load-time attestation using a Static Root of Trust for Measurement (SRTM) [40] and late-launch-based attestation using a Dynamic Root of Trust for Measurement (DRTM) [2, 4, 37]. In doing so, we make the following contributions. First, we formalize, using axioms, the behavior of core trusted computing primitives including the TCG's widely-deployed secure co-processor, the Trusted Platform Module (TPM), as well as recently introduced hardware to support the *late launch* of a security kernel in a protected execution environment. Hardware implementations of late launch are publicly available in both AMD's Secure Virtual Machine Architecture (SVM) [4] and Intel's Trusted eXecution Technology (TXT) [2]. These axioms provide a succinct specification of the primitives, which serve as building blocks in the proofs of the protocols (see Section 3).

Second, we formally define and prove code integrity and execution integrity properties of the attestation protocols (Section 4; Theorems 2–4). To the best of our knowledge, these are the first logical security proofs of these protocols.

Finally, the formal proofs yield insights about the security of these protocols. The invariants used in the proofs make precise the properties that the Trusted Computing Base (TCB) must satisfy. In Section 4, we describe these invariants and manually check that an invariant holds on a security kernel implementation used in an attestation protocol. We demonstrate that newly introduced hardware support for late launch actually adversely affects the security of previous generation attestation protocols. We describe an attack that utilizes hard-

ware support for late launch to exploit load-time attestation protocols that measure software starting at system boot. The attack enables an adversary to report false system integrity measurements that are not tied to the actual state of the platform. This attack could be used to exploit Digital Rights Management (DRM) protocols that rely on load-time attestation.

2 Logic of Secure Systems

We introduce the syntax of the Logic of Secure Systems (LS^2) in this section. The next section introduces features of LS^2 that are specific to trusted computing. We restrict technical descriptions to the extent necessary to explain the main concepts and application, and refer the reader to Appendix A for details.

2.1 Programming Model

The programming language definition includes its syntax and operational semantics. The syntax is summarized in Figure 5. The current language includes process calculi and functional constructs for modeling cryptographic operations, straightline code, and network communication among concurrent processes, but does not have conditionals (if...then...else...), returning function calls or loops. Instead, it has a match construct that tests equality of expressions (`match e, e'`) and blocks if the test fails, as well as unconditional jumps to arbitrary code (`jump e`). These constructs are sufficient for applications we have considered so far. In future work, we plan to investigate the technical challenges associated with adding conditionals, returning function calls, and loops to the language. We describe below the core language constructs, the adversary model, and the form of the operational semantics. Examples of programs in the language can be found in Section 4.

Data, agents, and keys. Data is represented in the programming model symbolically as expressions e (also called values). Expressions may be numbers n , identities of agents (principals) \hat{X} , keys K , variables x , pairs (e, e') , signatures using private keys $SIG_K\{e\}$ (denoting the signature on e made using the key K), asymmetric key encryptions $ENC_K\{e\}$, symmetric key encryptions $SYMENC_K\{e\}$, hashes $H(e)$, or code reified as data P . All expressions are assumed to be simply typed (e.g. a pair can be distinguished from a number), but we elide the details of the types. Agents, denoted \hat{X}, \hat{Y} , are users associated with a system on behalf of whom programs execute. Keys are denoted by the letter K . The inverse of

key K is denoted by K^{-1} . We assume that the expression e may be recovered from the signature $SIG_K\{e\}$ if the verification key K is known. We also assume that hashes are confidentiality preserving.

Systems, programs, and actions. A *secure system* is specified as a set of programs P in the programming language. For example, a trusted computing attestation system will contain two programs, one to be executed by the untrusted platform and the other by the remote verifier. Each *program* consists of a number of actions $x := a$ that are executed in a straight line. The name x binds the value returned by the action a , and is used to refer to the value in subsequent actions. Our model of straightline code execution is thus functional. This design choice simplifies reasoning significantly. For some actions such as sending a message, the value returned is meaningless. In such cases we assume that the value returned is the constant 0. A program ends with either an empty action `.`, or one of the special actions `jump e` or `latelaunch .`. The expression `jump e` is described below and `latelaunch .` is covered in the next section. A single executing program is called a *thread* $[P]_I$ (threads are referred to with variables T, S). It contains a program P , and a descriptor I for the thread that is a tuple $\langle \hat{X}, \eta, m \rangle$. \hat{X} is the agent that owns the thread, m is the machine on which the thread is hosted, and η is a unique identifier (akin to a process id). The abstract runtime environment of the language is called a *configuration* C , written $\iota, \sigma, T_1 | \dots | T_n$. It contains all executing threads $(T_1 | \dots | T_n)$, the state of memory on all machines (represented by the map σ), and the state of memory locks held by threads (represented by the map ι).

Cryptography and network primitives. The programming language includes actions for standard operations like signing and signature verification, encryption and decryption (both symmetric and asymmetric), nonce generation, hashing, expression matching, projection from a pair, and evaluation of arbitrary side-effect free functions (`eval f, e`). Threads can communicate with each other using actions to send and receive values over the network. Network communication is untargeted, i.e., any thread may intercept and read any message (dually, a received message could have been sent by any thread). Information being sent over the network may be protected using cryptography, if needed. The treatment of cryptography and network communication follows PCL. The language constructs we present next are new to this work.

Machines and shared memory. Threads can also share data through memory. The programming model contains machines m explicitly. Each machine contains

Expressions/Values	e	$::=$	n \hat{X}, \hat{Y} K K^{-1} x (e, e') $SIG_K\{e\}$ $ENC_K\{e\}$ $SYMENC_K\{e\}$ $H(e)$ P	Number Agent Key Inverse of key K Variable Pair Value e signed by private key K Value e encrypted by public key K Value e encrypted by symmetric key K Hash of e Program reified as data
Machine	m			
Location	l	$::=$	$m.RAM.k \mid m.disk.k \mid m.pcr.k \mid m.dpcr.k$	
Action	a	$::=$	$read\ l$ $write\ l, e$ $extend\ l, e$ $lock\ l$ $unlock\ l$ $send\ e$ $receive$ $sign\ e, K$ $verify\ e, K$ $enc\ e, K$ $dec\ e, K$ $symenc\ e, K$ $symdec\ e, K$ $hash\ e$ $eval\ f, e$ $proj_1\ e$ $proj_2\ e$ $match\ e, e'$ new	Read location l Write e to location l Extend PCR l with e Obtain write lock on location l Release write lock on location l Send e as a message Receive a message Sign e with private key K Check that $e = SIG_K\{e'\}$ Encrypt e with public key K Decrypt e with private key K Encrypt e with symmetric key K Decrypt e with symmetric key K Hash the expression e Evaluate function f with argument e Project the 1st component of a pair Project the 2nd component of a pair Check that $e = e'$ Generate a new nonce
Program	P, Q	$::=$	$\cdot \mid jump\ e \mid late\ launch \mid x := a; P$	
Thread id	I, J	$::=$	$\langle \hat{X}, \eta, m \rangle$	
Thread identifier	η			
Thread	T, S	$::=$	$[P]_I$	
Store	σ	:	Locations \rightarrow Expressions	
Lock map	ι	:	Locations \rightarrow (Thread ids) \cup $\{-\}$	
Configuration	C	$::=$	$\iota, \sigma, T_1 \mid \dots \mid T_n$	

Figure 1: Syntax of the programming language

a number of memory locations l that are shared by all threads running on the machine. Each location is classified as either RAM, persistent store (hard disk), or other special purpose location (such as Platform Configuration Registers that are described in the next section). The machine on which a location exists and the location's type are made explicit in the location's name. For instance, $m.RAM.k$ is the k th RAM location on machine m . The behavior of a location depends on its type. For example, RAM locations are set to a fixed value when a machine resets, whereas persistent locations are not af-

ected by resets. Despite these differences, the prominent characteristics of all locations are that they can be *read* and *written* through actions provided in the programming language, and that they are *shared* by all threads on the machine. Consequently, any thread, including an adversarial thread, has the potential to read or modify any location.

Access control on memory. Shared memory, by its very nature, cannot be used in secure programs unless some access control mechanism enforces the integrity and confidentiality of data written to it. Access control

varies by type of memory and application (e.g., memory segmentation, page table read-only bits, access control lists in file systems, etc). Our programming model provides an abstract form of access control through locks. Any running thread may obtain an exclusive-write lock on any previously unlocked memory location l by executing the action `lock l` . Information on locks held by threads is included in a configuration as a map t from locations to identities of threads that hold locks on them. The semantics of the programming language guarantee that while a lock is held by a thread, no other thread will be able to write the location. A thread may relinquish a lock it holds by executing the action `unlock l` . Locking in this manner may be used to enforce integrity of contents of memory. Similarly, one may add read locks that provide confidentiality of memory contents. Although technically straightforward, read locks are omitted from this paper since we are focusing on integrity properties.

Machine resets. The language allows a machine to be spontaneously reset. There is no specific action that causes a reset. Instead, there is a reduction in the operational semantics that may occur at any time to reset a machine. When this happens, all running threads on the machine are killed, all its RAM and PCR locations are set to a fixed value, and a single new thread is created to reboot the machine. This new thread executes a fixed booting program. We model the reset operation since it has significant security implications for secure systems [8]. In the context of trusted computing, e.g., the fact that a TPM’s Platform Configuration Registers (PCRs) are set to a fixed value is critical in reasoning about the security properties of attestation protocols. In addition, it has been shown that adversaries can launch realistic attacks against trusted computing systems using machine resets [14].

Untrusted code execution. The last salient feature of our programming model is an action `jump e` that dynamically branches to code represented by the expression e . The code e is arbitrary; it may have been read from memory or disk, or even have been received over the network. As a result, it could have come from an adversary. Execution of untrusted code is necessary to model several systems of interest, e.g., trusted computing systems and web browsers.

Adversary Model. We formally model adversaries as extra threads executing concurrently with protocol participants. Such an adversary may contain any number of threads, on any machines, and may execute any program expressible in our programming model. However, the adversary cannot perform operations that are not permitted by the language semantics. For example, the adversary

can neither write to memory locked by another thread, nor can she break cryptography.

Operational semantics. The operational semantics of the language captures how systems execute to produce traces. It is defined using process calculus-style reduction rules that specify how a configuration may transition to another. A *trace* $C_0 \longrightarrow C_1 \dots \longrightarrow C_n$ is a sequence of configurations, such that successive configurations in the sequence can be obtained by applying one reduction rule. A *timed trace* $C_0 \xrightarrow{t_1} C_1 \dots \xrightarrow{t_n} C_n$ associates monotonically increasing time points t_1, \dots, t_n with reductions on a trace. These time points may be drawn from any totally ordered set, such as integers or real numbers.

2.2 Logic

The logic LS^2 is used to specify and reason about properties of secure systems.

Syntax. Figure 2 summarizes LS^2 ’s syntax, including predicates specific to trusted computing that we discuss in the next section. Predicates for representing network communication and cryptographic operations are taken from PCL. Other predicates that capture information about state, unconditional jumps, and resets are new to this work. A significant difference from PCL is that LS^2 incorporates time explicitly in formulas and semantics. All predicates and formulas are interpreted relative to not only a timed trace but also a point of time (modal formulas, described below, are an exception since they are interpreted relative to a timed trace only). In the proof system, time is used to track the relative order of actions on a trace and to specify program invariants.

Action predicates capture actions performed by threads. For instance, $\text{Send}(I, e)$ holds on a trace at time t if thread I executes action `send e` at time t in the trace. $\text{Write}(I, l, e)$ holds on a trace whenever thread I executes `write l, e` . Similarly, we have predicates to capture cryptographic operations. *General predicates* capture other information, including information about the state of the environment. Particularly prominent are the two predicates $\text{Mem}(l, e)$ which holds whenever location l contains value e , and $\text{Jump}(I, e)$ which holds whenever thread I executes `jump e` . Access control on memory is reflected in the logic through three predicates: $\text{Lock}(I, l)$, $\text{Unlock}(I, l)$, and $\text{IsLocked}(l, I)$. The first two of these capture actions: $\text{Lock}(I, l)$ holds on a trace when a thread I obtains an exclusive-write lock on location l , whereas $\text{Unlock}(I, l)$ holds when thread I releases the lock. The third predicate $\text{IsLocked}(l, I)$ captures state: it holds whenever thread I has an exclusive-write lock on location l . As an example, suppose that thread I executes

Action Predicates	R	$::=$	Receive(I, e) Send(I, e) Sign(I, e, K) Verify(I, e, K) Encrypt(I, e, K) Decrypt(I, e, K) SymEncrypt(I, e, K) SymDecrypt(I, e, K) Hash(I, e) Eval(I, f, e, e') Match(I, e, e') New(I, n) Write(I, l, e) Read(I, l, e) Lock(I, l) Unlock(I, l) Extend(I, l, e)
General Predicates	M	$::=$	Mem(l, e) IsLocked(l, I) Reset(m, I) Jump(I, e) LateLaunch(m, I) Contains(e, e') $e = e'$ $t \geq t'$ Honest(\hat{X}, \bar{P})
Formulas	A, B	$::=$	R M \top \perp $A \wedge B$ $A \vee B$ $A \supset B$ $\neg A$ $\forall x.A$ $\exists x.A$ $A @ t$
Modal Formulas	J	$::=$	$[P]_I^{t_b, t_e} A$ $[a]_{I, x}^{t_b, t_e} A$

Figure 2: Syntax of LS^2

an action to obtain the lock on location l at time t and executes another action to release the lock at a later point t' . Then Lock(I, l) will hold exactly at time t , Unlock(I, l) will hold exactly at time t' , and IsLocked(l, I) will hold at all points of time between t and t' . The predicate Reset(m, I) holds at time t if machine m is reset at time t , creating the new thread I to boot it. We define the abbreviations Reset(m) and Jump(I) as $\exists I. \text{Reset}(m, I)$ and $\exists e. \text{Jump}(I, e)$ respectively. Contains(e, e') means that e' is a sub-expression of e . The predicate Honest(\hat{X}, \bar{P}) is described in Section 3.1.

Predicates can be combined using the usual logical connectives: \wedge (conjunction), \vee (disjunction), \supset (implication), and \neg (negation) as well as first-order universal and existential quantifiers that may range over expressions, keys, principals, threads, locations, and time. There is a special formula, $A @ t$, which captures time explicitly in the logic. $A @ t$ means that formula A holds at time t . We often write intervals in the usual mathematical sense; they may take the forms (t_1, t_2) , $[t_1, t_2]$, $(t_1, t_2]$, and $[t_1, t_2)$. For an interval i , we also define the formula A on i as $\forall t. ((t \in i) \supset A @ t)$, where $t \in i$ is the obvious membership predicate. A on i means that A holds at each point in the interval i . This treatment of time in the logic draws ideas from work on hybrid modal logic [5, 12, 31].

Security properties of programs are expressed in LS^2 using one of two forms of modal formulas. The principal of these, $[P]_I^{t_b, t_e} A$, means that formula A holds whenever thread I executes exactly the program P sequentially in the semi-open interval $(t_b, t_e]$. A may mention any variables occurring unbound in P . It usually expresses a safety property about the program P . For example, if P is the client program of a key exchange protocol, A may say that P generated a key after t_b , sent it to a server, and received a confirmation that it was received. Examples of security properties for trusted computing systems can be found in Section 4.

Proof System. Security properties of a program are established using a proof system for LS^2 . This proof system contains some basic rules for reasoning about modal formulas, and a number of axioms that capture in-

tuitive properties of program behavior. Parts of the proof system, particularly the part dealing with cryptographic primitives were easily designed using existing ideas from PCL. As mentioned in the introduction, a central design goal that LS^2 achieves is that the proof system does not mention adversary actions. We elaborate below on the technical approach for designing a sound proof system that supports this local style of reasoning in spite of the global nature of shared memory changes and execution of dynamically loaded code.

We reason about memory locally using axioms that establish invariance of values in memory, using information about locks and actions of threads that hold the locks. These axioms are modular (there is one set of axioms for each type of memory) and extensible (more axioms can be added for new types of memory, as we do for Platform Configuration Registers in Section 3). As examples, the following two axioms are invariance rules for locations of RAM and disk respectively. The first axiom says that if location $m.RAM.k$ (denoting a location with address k in the RAM of machine m) contains value e at time t_b , during the interval (t_b, t_e) thread I has a lock on this location, thread I does not write to the location, and machine m is not reset during the interval, then $m.RAM.k$ must contain the value e throughout the interval (t_b, t_e) . The second axiom is similar, but it applies to locations on disk. In this case, the precondition that machine m not be reset is unnecessary because contents of the disk do not change due to a reset.

$$\begin{aligned}
(\text{MemIR}) \quad & \vdash (\text{Mem}(m.RAM.k, e) @ t_b) \\
& \wedge (\text{IsLocked}(m.RAM.k, I) \text{ on } (t_b, t_e)) \\
& \wedge (\forall e'. \neg \text{Write}(I, m.RAM.k, e') \text{ on } (t_b, t_e)) \\
& \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e)) \\
& \supset (\text{Mem}(m.RAM.k, e) \text{ on } (t_b, t_e))
\end{aligned}$$

$$\begin{aligned}
(\text{MemID}) \quad & \vdash (\text{Mem}(m.disk.k, e) @ t_b) \\
& \wedge (\text{IsLocked}(m.disk.k, I) \text{ on } (t_b, t_e)) \\
& \wedge (\forall e'. \neg \text{Write}(I, m.disk.k, e') \text{ on } (t_b, t_e)) \\
& \supset (\text{Mem}(m.disk.k, e) \text{ on } (t_b, t_e))
\end{aligned}$$

For reasoning about execution of dynamically loaded

code, we introduce the following rule that allows us to combine information about the invariants of a program P with the knowledge that the program was branched to. We define a program invariant as a property that holds whenever any prefix of the sequence of actions of the program executes. The prefixes or initial sequences $IS(P)$ of a program P are formally defined as follows: $IS(\cdot) = \{\cdot\}$, $IS(\text{jump } e) = \{\cdot, \text{jump } e\}$, $IS(\text{latelaunch}) = \{\cdot, \text{latelaunch}\}$, $IS(x := a; P) = \{\cdot\} \cup \{x := a; Q \mid Q \in IS(P)\}$.

$$\frac{\text{For every } Q \text{ in } IS(P) : \vdash [Q]_I^{t_b, t_e} A(t_b, t_e) \quad (t_b, t_e \text{ fresh constants})}{\vdash \text{Jump}(I, P) @ t \supset \forall t'. (t' > t) \supset A(t, t')} \text{Jump}$$

In its premise the rule requires that for every initial sequence Q of P , there be a proof, generic in the constants t_b and t_e , that establishes $A(t_b, t_e)$ given that Q executes in thread I during the interval $(t_b, t_e]$. The conclusion says that if thread I branches to program P at time t (assumption $\text{Jump}(I, P) @ t$), then for any time $t' > t$, $A(t, t')$ must hold. Informally, we may explain the soundness of this rule as follows. If thread I branches to code P at time t , then for any $t' > t$, the thread I must execute some prefix of P in the interval $(t, t']$. Instantiating the premise with this prefix Q , and t, t' for t_b, t_e , we get exactly the desired property $A(t, t')$.

The above rule is central among LS^2 's principles for reasoning about dynamically loaded code, which we believe to be novel. Both a discussion of the novelty and an example of the reasoning principles are postponed to Section 4.1. Whereas their application to reasoning about dynamically loaded code is new, invariants over initial segments of code are not a contribution of this work. PCL uses invariants similar to ours to reason about principals who are executing known pieces of code. LS^2 also uses invariants for many other purposes besides reasoning about jumps, including reasoning about resets. The latter is simpler than reasoning about jumps, because we assume that when a machine is reset, a *fixed* program is started to reboot the machine. The code marked $SRTM(m)$ in Figure 3 is one example of the form this program may have.

Semantics and Soundness. Formulas of LS^2 are interpreted over timed traces obtained from execution of a program in the programming language. The proof system of LS^2 is formally connected to the programming language semantics through a program independent *soundness theorem* which guarantees that any property established in the proof system actually holds over all traces obtainable from the program and any number of

adversarial threads. Let Γ denote a set of formulas, and φ denote a formula or a modal formula. Further, let $\Gamma \vdash \varphi$ denote provability in LS^2 's proof system, and $\Gamma \models \varphi$ denote semantic entailment. Our main technical result for LS^2 is the following soundness theorem.

Theorem 1 (Soundness). *If $\Gamma \vdash \varphi$ then $\Gamma \models \varphi$.*

Proof. See Appendix A. □

3 Modeling Trusted Computing Primitives

This section describes extensions to LS^2 to model and reason about hardware primitives used with protocols specified by the Trusted Computing Group (TCG). These hardware primitives include the TCG's Trusted Platform Module (TPM) and static Platform Configuration Registers (PCRs), as well as the more recent hardware support for late launch and dynamic PCRs as implemented by AMD's Secure Virtual Machine (SVM) extensions [1] and Intel's Trusted eXecution Technology (TXT) [2]. We describe below the hardware primitives and their formalization in LS^2 at a high level. In subsequent sections, we use our formalizations to prove security properties of trusted computing protocols.

3.1 Trusted Platform Module

The Trusted Platform Module (TPM) is a secure co-processor that performs cryptographic operations such as encryption, decryption, and creation and verification of digital signatures. Each TPM includes a unique embedded private key (called the Attestation Identity Key or AIK). The public key corresponding to each AIK is published in a manufacturer-signed certificate. The private component of the AIK is assumed to be protected from compromise by malicious software. As a result, signatures produced by a TPM are guaranteed to be authentic, and unique to the platform on which the TPM resides.

We model relevant aspects of the TPM in LS^2 as follows. The private attestation identity key of the TPM on machine m is modeled as a value in LS^2 , denoted $AIK^{-1}(m)$. Its corresponding public key is denoted $AIK(m)$. The TPM itself is represented as a principal, denoted $AIK\hat{K}(m)$. Of the many programs hardcoded into the TPM, only two are relevant for our purposes. These are idealized by the LS^2 programs marked $TPM_{SRTM}(m)$ and $TPM_{DRTM}(m)$ in Figures 3 and 4 respectively, and are explained in the next section. Both the fact that the TPM executes only one of these programs, and the fact

that the TPM's private key cannot be leaked are modeled in LS^2 by a single predicate:

$$\text{Honest}(AIK(m), \{TPM_{SRTM}(m), TPM_{DRTM}(m)\})$$

This predicate entails (through the rules and axioms of the proof system) that any signature created by the key $AIK^{-1}(m)$ could only have been created in the TPM on machine m . It can also be used to prove invariants about threads which are known to execute on the TPM, using a rule similar to (Jump) that was described in Section 2.2.¹ We emphasize that the predicate mentioned above is not an axiom in LS^2 , since its soundness cannot be established directly. Instead, we always assume it explicitly when we reason about the TPM.

Static PCRs. Static Platform Configuration Registers (PCRs) are protected registers contained in every TPM. From our perspective, the relevant property of PCRs is that their contents can only be modified in two ways: (a) by resetting the machine on which the TPM resides; this sets all the static PCRs to a special value that we denote symbolically using the name *sinit* (*sinit* is zero on most platforms), and (b) through a special TPM interface `extend`, which takes two arguments: a PCR to modify, and a value v that is appended to the PCR. Since each PCR is of a fixed length but may be asked to store arbitrarily many values, `extend` replaces the current value of the PCR with a hash of the concatenation of its current value and a hash of v . In pseudocode, the effect of extending PCR p with value v may be described as the assignment $p \leftarrow H(p \parallel H(v))$, where \parallel denotes concatenation and H denotes a hash function. More generally, if the values extended into a PCR after a reset are v_1, \dots, v_n in sequence, its contents will be $H(\dots(H(sinit \parallel H(v_1)) \parallel H(v_2)) \dots \parallel H(v_n))$. We use the notation $seq(sinit, v_1, \dots, v_n)$ to denote this value. A common use for PCRs is to extend integrity measurements of program code into them during the boot process, then to have the TPM sign them with its AIK, and to submit this signed aggregate to a remote party as evidence that the values were generated in sequence on the machine.

We model PCRs as a special class of memory in LS^2 . The k th static PCR on machine m is denoted $m.pcr.k$. PCRs can be read using the usual read action in LS^2 's programming language, and they can be locked for access control, but the usual write action does not apply to them. Instead, the `extend` program is modeled as

¹The predicate `Honest` is adapted from a predicate of the same name in PCL. PCL's predicate is slightly weaker since it lacks the second argument, but the reasoning principles associated with the two are similar.

a primitive action in the programming language. It has exactly the effect described in the previous paragraph. Properties of PCRs are captured through axioms in LS^2 . For example, the following axiom models the fact that *sinit* is written to every PCR when a machine is reset. In words, it states that if machine m is reset at time t , then any PCR k on m contains value *sinit* at time t .

$$(\text{MemPR}) \quad \vdash (\text{Reset}(m) @ t) \supset (\text{Mem}(m.pcr.k, sinit) @ t)$$

Several other important properties of PCRs arise as a consequence of their restricted interface. First, if a PCR contains *sinit* at time t , then the machine m on which it resides must have been reset *most recently* at some time t' since a reset is the only way to put *sinit* into a PCR. This is captured by the following axiom:

$$\begin{aligned} (\text{PCR2}) \quad & \vdash (\text{Mem}(m.pcr.k, sinit) @ t) \\ & \supset (\exists t'. (t' \leq t) \wedge (\text{Reset}(m) @ t') \\ & \quad \wedge (\neg \text{Reset}(m) \text{ on } (t', t])) \end{aligned}$$

Second, if a PCR contains $seq(sinit, v_1, \dots, v_n)$ at time t , it must also have contained $seq(sinit, v_1, \dots, v_{n-1})$ at some prior time t' , without any reset in the interim. Thus the contents of a PCR are witness to every extension performed on it since its last reset. Formally, this property is captured in LS^2 by the following axiom:

$$\begin{aligned} (\text{PCR1}) \quad & \vdash (\text{Mem}(m.pcr.k, seq(sinit, v_1, \dots, v_n)) @ t) \\ & \supset (\exists t'. (t' < t) \\ & \quad \wedge (\text{Mem}(m.pcr.k, seq(sinit, v_1, \dots, v_{n-1})) @ t') \\ & \quad \wedge (\neg \text{Reset}(m) \text{ on } (t', t])) \quad (n \geq 1) \end{aligned}$$

In many cases of interest, we need to prove that the value in a PCR does not change over a period of time. To this end, we introduce an invariance axiom for PCRs, similar to axioms (MemIR) and (MemID) from Section 2.2. The modular design of the logic eases the introduction of this axiom.

$$\begin{aligned} (\text{MemIP}) \quad & \vdash (\text{Mem}(m.pcr.k, e) @ t_b) \\ & \quad \wedge (\text{IsLocked}(m.pcr.k, I) \text{ on } (t_b, t_e)) \\ & \quad \wedge (\forall e'. \neg \text{Extend}(I, m.pcr.k, e') \text{ on } (t_b, t_e)) \\ & \quad \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e)) \\ & \quad \supset (\text{Mem}(m.pcr.k, e) \text{ on } (t_b, t_e)) \end{aligned}$$

3.2 Late Launch and Dynamic PCRs

Another hardware feature available in trusted computing platforms is late launch. Late launch provides the ability to measure and invoke a program, typically a security kernel or Virtual Machine Monitor (VMM), in a protected environment. Upon receiving a late launch instruction (`SKINIT` on the AMD SVM and `SENDER` on

the Intel TXT), the processor switches from the currently executing operating system to a Dynamic Root of Trust for Measurement (DRTM) from which it is possible to later resume the suspended operating system. The program to be executed in a late launch session is specified by providing the physical address of the Secure Loader Block (SLB). When a late launch is performed, interrupts are disabled, direct memory access (DMA) is disabled to all physical memory pages containing the SLB and debugging access is disabled. The processor then jumps to the code in the SLB. This code may load other code. In addition to providing a protected environment, a special set of PCRs called *dynamic* PCRs are reset with a special value that we call *dinit* symbolically and the code in the SLB is hashed and extended into the dynamic PCR 17 (*dinit* is distinct from *sinit*). The dynamic PCRs can then be extended with other values, and the contents of the PCRs, signed by the TPM’s key AIK, can be submitted as evidence that a late launch was performed.

We formally model late launch by adding a new action `lateLaunch` to LS^2 ’s programming language. This action can be executed by any thread. The operational semantics of the language are extended to ensure that whenever `lateLaunch` executes a new thread I is created with a special program $LL(m)$, which extends the SLB into a dynamic PCR and branches to it. This program is shown in Figure 4. Protection of I is modeled using locks – when started, I is given locks to all dynamic PCRs on the machine m it uses. I may subsequently acquire more locks to protect itself. In the logic, the implicit locking of dynamic PCRs is captured by the following axiom, which means that if some thread executes `lateLaunch` on machine m at time t , creating the thread I , then I has a lock on any dynamic PCR on m at time t . $m.dpcr.k$ denotes the k th dynamic PCR on machine m .

$$\begin{aligned} (\text{LockLL}) \quad &\vdash (\text{LateLaunch}(m, I) @ t) \\ &\supset (\text{IsLocked}(m.dpcr.k, I) @ t) \end{aligned}$$

Dynamic PCRs have properties very similar to static PCRs. For example, the following axiom, similar to (MemPR) described above, means that *dinit* is written to every dynamic PCR when a late launch happens.

$$\begin{aligned} (\text{MemLL}) \quad &\vdash (\text{LateLaunch}(m, I) @ t) \\ &\supset (\text{Mem}(m.dpcr.k, dinit) @ t) \end{aligned}$$

Axioms corresponding to (PCR1) and (PCR2) are also sound for dynamic PCRs. The difference is that `Reset` and `sinit` must be replaced by `LateLaunch` and *dinit* respectively. The following axiom is used to prove invari-

ance properties of dynamic PCRs.

$$\begin{aligned} (\text{MemIdP}) \quad &\vdash (\text{Mem}(m.dpcr.k, e) @ t_b) \\ &\wedge (\text{IsLocked}(m.dpcr.k, I) \text{ on } (t_b, t_e)) \\ &\wedge (\forall e'. \neg \text{Extend}(I, m.dpcr.k, e') \text{ on } (t_b, t_e)) \\ &\wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e)) \\ &\wedge (\neg \exists I. \text{LateLaunch}(m, I) \text{ on } (t_b, t_e)) \\ &\supset (\text{Mem}(m.dpcr.k, e) \text{ on } (t_b, t_e)) \end{aligned}$$

4 Trusted Computing Protocols

We analyze two trusted computing protocols that rely on TPMs to provide integrity properties: load-time attestation using an SRTM and late-launch-based attestation using a DRTM. In an attestation protocol, a platform utilizes a TPM to attest to platform state by performing two steps: integrity measurement and integrity reporting. Integrity measurement consists of collecting cryptographic hashes of local software events such as program loading. Integrity reporting consists of transmitting collected measurements in a signed aggregate to an external verifier. The external verifier may then use the measurements to make trust decisions. We first analyze an attestation protocol using a Static Root of Trust for Measurement (SRTM), then we consider an attestation protocol utilizing hardware support for late launch and a Dynamic Root of Trust for Measurement (DRTM). We simplify both protocols by assuming the AIK has been certified as authentic by a manufacturer certificate and by verifying a fixed sequence of system integrity measurements.

4.1 Attestation Using a Static Root of Trust

We start by performing an analysis of a load-time attestation protocol using an SRTM. The security skeleton of the protocol is specified in Figure 3. A security skeleton retains only relevant actions, in this case, actions performing integrity measurement and reporting. The SRTM protocol is composed of code that performs measurement followed by code that performs integrity reporting. We analyze the components separately.

4.1.1 Integrity Measurement

In the SRTM protocol, integrity measurement starts after a machine reset. The programs marked $SRTM(m)$, $BL(m)$, and $OS(m)$ in Figure 3 represent those portions of the SRTM, boot loader, and operating system that participate in the measurement process. The $SRTM(m)$ program is always the first program invoked when a machine reboots. It first reads the boot loader’s code b

$SRTM(m) \equiv$ $b = \text{read } m.bl_Loc;$
 $\text{extend } m.pcr.s, b;$
 $\text{jump } b$

$BL(m) \equiv$ $o = \text{read } m.os_Loc;$
 $\text{extend } m.pcr.s, o;$
 $\text{jump } o$

$OS(m) \equiv$ $a = \text{read } m.app_Loc;$
 $\text{extend } m.pcr.s, a;$
 $\text{jump } a$

$APP(m) \equiv \dots$

$TPM_{SRTM}(m) \equiv$ $w = \text{read } m.pcr.s;$
 $r = \text{sign}(PCR(s), w), AIK^{-1}(m);$
 $\text{send } r$

$Verifier(m) \equiv$ $sig = \text{receive};$
 $v = \text{verify } sig, AIK(m);$
 $\text{match } v, (PCR(s),$
 $\quad seq(sinit, BL(m), OS(m), APP(m)))$

Figure 3: Security Skeleton for SRTM Attestation Protocol

from the fixed disk address $m.bl_Loc$, then measures the code by extending it into a static PCR $m.pcr.s$ (which in this case stores all measurements), and then branches to the boot loader by executing the instruction $\text{jump } b$. The boot loader ($BL(m)$) in turn reads the operating system's code o from a fixed location $m.os_Loc$, extends it into PCR $m.pcr.s$, and branches to it. The operating system ($OS(m)$) performs similar actions with the application's code a . The application ($APP(m)$) may perform any actions. In practice, the sequence of measurement and loading may continue beyond the first application but we have chosen to terminate it here because extending the chain further does not lead to any new insights about the security of the system.

Security Property. We summarize the integrity measurement security property as follows: if $m.pcr.s$ is protected while a machine boots, and the contents of $m.pcr.s$ are $seq(sinit, BL(m), OS(m), APP(m))$, then the initial software loaded on machine m since its last reboot was $BL(m)$ followed by $OS(m)$. We now state this property formally. We define the formulas $\text{ProtectedSRTM}(m)$

and $\text{MeasuredBoot}_{SRTM}(m, t)$ as follows.

$\text{ProtectedSRTM}(m) =$
 $\forall t, I. (\text{Reset}(m, I) @ t) \supset (\text{IsLocked}(m.pcr.s, I) @ t)$

$\text{MeasuredBoot}_{SRTM}(m, t) =$
 $\exists t_T. \exists t_B. \exists t_O. \exists J. (t_T < t_B < t_O < t) \wedge$
 $(\text{Reset}(m, J) @ t_T) \wedge (\text{Jump}(J, BL(m)) @ t_B) \wedge$
 $(\text{Jump}(J, OS(m)) @ t_O) \wedge (\neg \text{Reset}(m) \text{ on } (t_T, t])$
 $(\neg \text{Jump}(J) \text{ on } (t_T, t_B)) \wedge (\neg \text{Jump}(J) \text{ on } (t_B, t_O))$

$\text{ProtectedSRTM}(m)$ means that any thread I created to boot machine m after a reset obtains an exclusive-write lock on $m.pcr.s$. $\text{MeasuredBoot}_{SRTM}(m, t)$ identifies software events on m such as the boot loader and operating system being branched to before time t . It comprises four facts: (1) There exists a time t_T before t at which m was reset, creating a thread J to boot the machine ($\text{Reset}(m, J) @ t_T$), (2) This thread J branched to the programs $BL(m)$ and $OS(m)$ at later time points t_B and t_O ($\text{Jump}(J, BL(m)) @ t_B$ and $\text{Jump}(J, OS(m)) @ t_O$), (3) J did not make any other jumps in the interim ($\neg \text{Jump}(J) \text{ on } (t_T, t_B)$) and ($\neg \text{Jump}(J) \text{ on } (t_B, t_O)$), and (4) Machine m was not reset between t_T and t ($\neg \text{Reset}(m) \text{ on } (t_T, t]$). Equivalently, after its last reboot before time t , the first programs loaded on m were $BL(m)$ and $OS(m)$. We believe this is a natural property to expect from a system integrity measurement protocol.

The following theorem formalizes our security property. It states that under the assumptions that $m.pcr.s$ is protected during booting, and that $m.pcr.s$ contains $seq(sinit, BL(m), OS(m), APP(m))$ at time t , it is guaranteed that the boot loader and operating system used to boot the machine are $BL(m)$ and $OS(m)$ respectively.

Theorem 2 (Security of Integrity Measurement). *The following is provable in LS^2 :*

$\text{ProtectedSRTM}(m) \vdash$
 $\text{Mem}(m.pcr.s, seq(sinit, BL(m), OS(m), APP(m))) @ t$
 $\supset \text{MeasuredBoot}_{SRTM}(m, t)$

We refer the reader to Appendix B for a detailed proof of this theorem. Major steps in the proofs are discussed below to illustrate novel reasoning principles in LS^2 . All programs mentioned below refer to Figure 3.

- (1) Using axioms (PCR1) and (PCR2) in succession on the antecedent $\text{Mem}(m.pcr.s, seq(sinit, BL(m), OS(m), APP(m))) @ t$, we show that all sub-sequences of $seq(sinit, BL(m), OS(m), APP(m))$ must have appeared in $m.pcr.s$ at times earlier than t , and that

machine m must have been reset at some time t_T , creating a thread J to boot it. Formally, we obtain

$$\begin{aligned} & \exists t_T, t_1, t_2, t_3, J. (t_T \leq t_1 < t_2 < t_3 < t) \\ & \wedge (\text{Mem}(m.pcr.s, \text{seq}(\text{sinit}, BL(m), OS(m))) @ t_3) \\ & \wedge (\text{Mem}(m.pcr.s, \text{seq}(\text{sinit}, BL(m))) @ t_2) \\ & \wedge (\text{Mem}(m.pcr.s, \text{sinit}) @ t_1) \\ & \wedge (\text{Reset}(m, J) @ t_T) \\ & \wedge (\neg \text{Reset}(m) \text{ on } (t_T, t]) \end{aligned}$$

- (2) Since m was reset creating thread J (second from last conjunct above), it follows in our model that the thread J above must have started with the program $SRTM(m)$. (We have omitted a description of the rules that force this to be the case.) Thus, we would like to proceed by proving an invariant of $SRTM(m)$. However, we can say *nothing* about the program b loaded at the end of $SRTM(m)$. This is because b is read from a memory location $m.bl_Loc$, which could potentially have been written by an adversarial thread earlier. Fortunately, the extension of b into $m.pcr.s$ in the second line of $SRTM(m)$ lets us proceed. Precisely, this extension along with some basic properties of PCRs lets us prove the following property that is *parametric* (universally quantified) in the code b . t_T and J were obtained in property (1).

$$\begin{aligned} & \forall t', b, o. \\ & (((\text{Mem}(m.pcr.s, \text{seq}(\text{sinit}, b, o)) @ t') \\ & \wedge (t_T < t' \leq t)) \\ & \supset \exists t_B. ((t_T < t_B < t') \wedge (\text{Jump}(J, b) @ t_B))) \\ & \wedge (\text{IsLocked}(m.pcr.s, J) @ t_B) \end{aligned}$$

This property means that if at any time t' between t_T and t , $m.pcr.s$ contained $\text{seq}(\text{sinit}, b, o)$, then thread J must have branched to b at some time t_B between t_T and t' , and that J must hold a lock on $m.pcr.s$ at t_B . Informally this holds because the action immediately following the extension in $SRTM(m)$ is jump b , so if there is a further extension with o , jump b must have happened in the interim. The assumption $\text{ProtectedSRTM}(m)$ is used to rule out the possibility that a thread other than J extended o into $m.pcr.s$ before jump b happened, and to show that J holds the lock on $m.pcr.s$ at t_B .

- (3) We instantiate the property in (2), choosing $b = BL(m)$, $o = OS(m)$, and $t' = t_3$ (t_3 was obtained in (1)). Eliminating the antecedents of the implication using facts from (1), we obtain:

$$\begin{aligned} & \exists t_B. ((t_T < t_B < t_3) \wedge (\text{Jump}(J, BL(m)) @ t_B) \\ & \wedge (\text{IsLocked}(m.pcr.s, J) @ t_B)) \end{aligned}$$

- (4) From (3) we know $\text{Jump}(J, BL(m)) @ t_B$. Next we use the (Jump) rule from Section 2.2. In the premise we show that $\forall t_b, t_e. \forall Q \in IS(BL(m)). \vdash [Q]_J^{t_b, t_e} A(t_b, t_e)$ for a suitable invariant $A(t_b, t_e)$, whose details we omit here (see Appendix B for details). The main difficulty here is similar to that in (2): we do not know what o in the program of $BL(m)$ may be. Again, the invariant we prove is parametric in o . Using the (Jump) rule, we obtain the following property.

$$\begin{aligned} & \forall t', o, a. \\ & (((\text{Mem}(m.pcr.s, \text{seq}(\text{sinit}, BL(m), o, a)) @ t') \\ & \wedge (t_B < t' \leq t)) \\ & \supset \exists t_O. ((t_B < t_O < t) \wedge (\text{Jump}(J, o) @ t_O))) \end{aligned}$$

This property is very similar to that in (2), except that it follows from an invariant of $BL(m)$, not $SRTM(m)$. The fact $\text{IsLocked}(m.pcr.s, J) @ t_B$ from (3) is needed to rule out the possibility that a thread other than J extended a into $m.pcr.s$.

- (5) We instantiate the property in (4), choosing $o = OS(m)$, $a = APP(m)$ and $t' = t$. Combining with facts from (1), we obtain:

$$\exists t_O. ((t_B < t_O < t) \wedge (\text{Jump}(J, OS(m)) @ t_O))$$

The facts $(\text{Reset}(m, J) @ t_T)$, $(\neg \text{Reset}(m) \text{ on } (t_T, t])$, $(\text{Jump}(J, BL(m)) @ t_B)$, and $(\text{Jump}(J, OS(m)) @ t_O)$ in (1), (3), and (5) establish part of $\text{MeasuredBoots}_{SRTM}(m, t)$. The remaining part follows from a similar analysis with slightly stronger invariants in (2) and (4).

The hardest part in designing LS^2 's proof system was coming up with sound principles for reasoning about dynamically branching to unknown code that are illustrated above, and in particular, the (Jump) rule. Although the final design is simple to use, it was not obvious at first. We believe that this method for reasoning about branching to completely unknown code (like b and o) is new to this work. Prior work on reasoning about dynamically loaded code, usually based on higher-order extensions of Hoare logic [7, 20, 27, 28, 39], assumes that at least the invariants of the code being branched to are known at the point of branch in the program. In our setting, this assumption is unrealistic because we allow executable code to either be obtained over the network or be read from memory, and hence, potentially, to come from an adversary.

4.1.2 Insights From Analysis

A number of insights follow from the analysis. These insights include highlighting an unexpected property, clar-

ifying assumptions on the TCB, and identifying program invariants required for security.

Property Excludes Last Jump. A key insight from the analysis is that the integrity measurement protocol does not provide sufficient evidence to deduce that the last program in a chain of measurements is actually executed. For example, an adversary can reboot the platform after $OS(m)$ extends $APP(m)$, but before it is jumped to. Alternatively, a race may occur between two application-level processes whereby the OS extends the first into $m.pcr.s$ and then the other process reads the value in $m.pcr.s$ before the first process is branched to.

TCB Assumptions. The value of $m.pcr.s$ does not guarantee that the measured software was also executed unless it is also guaranteed that no other process had write access to $m.pcr.s$. If the latter assumption fails, an attack exists: a malicious process may extend a piece of code into $m.pcr.s$ without executing it. This assumption usually holds in practice because booting is generally single threaded, but may fail if for example a malicious thread executes on another processor core concurrently with the measurement thread. Formally, this shows up as the ProtectedSRTM(m) formula, which is a necessary assumption for the proof.

Program Invariants. To establish Theorem 2, we prove program invariants for the $SRTM(m)$ and $BL(m)$ programs. These invariants provide a specification of the properties that an SRTM and a boot loader program must satisfy to be secure in an integrity measurement protocol, i.e. the assumptions about the TCB. The $SRTM(m)$ invariant states that there exists a time point t' and thread J such that J branches to the boot loader b , J does not branch to any program at any time point before t' , $m.pcr.s$ contains the hashed value of the boot loader b , and $m.pcr.s$ is locked by J at t' . The invariant of $BL(m)$ states that there exists a time point t and thread J such that J branches to program code o only after the entire program code o has been measured into $m.pcr.s$. Kauer [18] performed a manual source code audit of a number of TPM-enabled boot loaders to check the informal security condition that “no code...is executed but not hashed.” Our invariant on the boot loader BL was developed independently during the course of proving the above theorem and is a formal specification of this condition. We envisage that these invariants can be used to derive properties to automatically check against implementations of TCB components, thereby providing greater assurance that the trusted components are trustworthy.

4.1.3 Integrity Reporting

After the integrity measurement protocol loads the PCRs with measurements, the measurements can be used by the TPM to attest to the identify of the software loaded on the local platform. This protocol, called integrity reporting, involves two participants. One of the participants is the remote party itself, called the verifier. Its code is marked $Verifier(m)$ in Figure 3. The other participant in the protocol is the TPM of machine m in the role of $TPM_{SRTM}(m)$. This code is also shown in the same figure.

The integrity reporting protocol contains two steps. In the first, the TPM on machine m reads the contents w of $m.pcr.s$, signs them and an identifier (denoted $PCR(s)$) that uniquely identifies $m.pcr.s$ with its embedded private key $AIK^{-1}(m)$ and sends the signed aggregate to the remote verifier. In the second step, the remote verifier verifies this signature with the known public key $AIK(m)$, and checks that the contents of the signature match the pair of $PCR(s)$ and $seq(sinit, BL(m), OS(m), APP(m))$.

Security Properties. The security properties of integrity reporting are formalized by the following two LS^2 formulas, which we call J_1 and J_2 respectively.

$$[Verifier(m)]_V^{t_b, t_e} \exists t. (t < t_e) \wedge (Mem(m.pcr.s, seq(sinit, BL(m), OS(m), APP(m))) @ t)$$

$$[Verifier(m)]_V^{t_b, t_e} \exists t. (t < t_e) \wedge MeasuredBoot_{SRTM}(m, t)$$

The first property (J_1) states that if the code $Verifier(m)$ is executed successfully between the time points t_b and t_e , then there must be a time t before t_e at which $m.pcr.s$ contained $seq(sinit, BL(m), OS(m), APP(m))$. The second property (J_2) means that a remote verifier can identify the boot loader and operating system that were loaded on m at some time prior to t_e .

To prove these properties, we require two new assumptions, which we combine in the set Γ_{SRTM} below. The first of these assumptions states that the remote verifier is distinct from the TPM. This assumption is needed to distinguish protocol participants, and is true in practice. The second assumption is the honesty assumption for the TPM from Section 3.1 that guarantees that the TPM’s signature cannot be forged, and that the TPM always executes only specified programs.

$$\Gamma_{SRTM} = \{ \hat{V} \neq AIK(m), \text{Honest}(AIK(m), \{TPM_{SRTM}(m), TPM_{DRTM}(m)\}) \}$$

Theorem 3 (Security of Integrity Reporting). *The following are provable in LS^2 ’s proof system:*

(1) $\Gamma_{SRTM} \vdash J_1$

(2) $\Gamma_{SRTM, ProtectedSRTM(m)} \vdash J_2$

The proof of (1) critically relies on the assumption $Honest(A\hat{I}K(m), \{TPM_{SRTM}(m), TPM_{DRTM}(m)\})$ to establish both that the TPM on m actually produced a signature in the past, and that the value signed by the TPM was actually read from $m.pcr.s$. The latter follows from knowledge of the programs that the TPM may be executing. (2) follows from (1) and Theorem 2. See Appendix B for details of the proofs.

4.1.4 Insights From Analysis

The security analysis lead to a number of insights including highlighting weaknesses in the property provided by the protocol and identifying program invariants required for security.

Staleness of Measurements. A key insight from the analysis is that after executing the integrity reporting protocol, the verifier has no knowledge of how recent the time of measurement t is in comparison to t_e , the time the verifier's execution finished. This staleness of measurements is inherent in the protocol: it is possible to reboot the machine with a different boot sequence after sending the signature to the remote verifier, as is known from prior work [14]. Formally, one can only prove that $m.pcr.s$ contained the reported measurements at time t , but not after.

Program Invariants. In the process of proving the above theorem, we prove a program invariant for the roles of the TPM (i.e., $TPM_{SRTM}(m)$ and $TPM_{DRTM}(m)$). This invariant provides a specification of the properties that a TPM's signing role must satisfy. In particular, the invariant requires that if the TPM returns a value then the value is a signature over the value stored in $m.pcr.s$ and that the TPM does not write to any memory locations. The latter constraint is necessary to prevent previously measured code from being modified after being measured.

4.2 Attestation Using a Dynamic Root of Trust

We perform an analysis of DRTM attestation using our model of hardware support for late launch. We jointly analyze the protocol code that performs integrity measurement and reporting.

4.2.1 DRTM Protocol

We describe the security skeleton of the DRTM attestation protocol in Figure 4. The DRTM protocol is a four agent protocol. The processes are: (1) $OS(m)$, executed by the machine itself (called \hat{m}), that receives a nonce from the remote verifier, and performs a late launch. (2) $LL(m)$, executed by the hardware platform, that reads the binary of the program $P(m)$ from the secure loader block (SLB), and measures then branches to $P(m)$, (3) $P(m)$ that measures the nonce, evaluates the function f on input 0 (the function f and its input may be changed depending on application), and extends a distinguished string EOL into $m.dpcr.k$ to signify the end of the late launch session. (4) $TPM_{DRTM}(m)$, executed by the TPM of m , that signs the dynamic PCR $m.dpcr.k$, and sends it to the verifier. (5) $Verifier(m)$, executed by a remote verifier, that generates and sends a nonce, receives signed measurements, verifies the signature, and checks that the measurements match the sequence $(dinit, P(m), n, EOL)$.

Security Property. We summarize the DRTM security property as follows: if the verifier is not the TPM, the TPM does not leak its signing key, and the TPM executes only the processes $TPM_{DRTM}(m)$ and $TPM_{SRTM}(m)$, then the remote verifier is guaranteed that J performed a single late launch on machine m at some time t_L , J branched to $P(m)$ only once at t_C , J evaluated f once at t_E (and this happened after the verifier generated the nonce), J extended EOL into $m.dpcr.k$ at some time t_X , and $m.d.pcr.k$ was locked for the thread J from t_L to t_X . We formalize this security property called J_{DRTM} below.

$$\begin{aligned}
[Verifier(m)]_V^{t_b, t_e} \quad & \exists J, t_X, t_E, t_N, t_L, t_C, n. \\
& \wedge (t_L < t_C < t_E < t_X < t_e) \\
& \wedge (t_b < t_N < t_E) \\
& \wedge (New(V, n) @ t_N) \\
& \wedge (LateLaunch(m, J) @ t_L) \\
& \wedge (\neg LateLaunch(m) \text{ on } (t_L, t_X]) \\
& \wedge (\neg Reset(m) \text{ on } (t_L, t_X]) \\
& \wedge (Jump(J, P(m)) @ t_C) \\
& \wedge (\neg Jump(J) \text{ on } (t_L, t_C)) \\
& \wedge (Eval(J, f) @ t_E) \\
& \wedge (Extend(J, m.dpcr.k, EOL) @ t_X) \\
& \wedge (\neg Eval(J, f) \text{ on } (t_C, t_E)) \\
& \wedge (\neg Eval(J, f) \text{ on } (t_E, t_X)) \\
& \wedge (IsLocked(m.dpcr.k, J) \text{ on } (t_L, t_X])
\end{aligned}$$

In order to prove the property, we have to make the

$OS(m)$	\equiv	$n' = \text{receive};$ $\text{write } m.\text{nonce}, n';$ late_launch
$LL(m)$	\equiv	$P = \text{read } m.\text{SLB};$ $\text{extend } m.\text{dpcr.k}, P;$ $\text{jump } P$
$P(m)$	\equiv	$n'' = \text{read } m.\text{nonce};$ $\text{extend } m.\text{dpcr.k}, n'';$ $\text{eval } f, 0;$ $\text{extend } m.\text{dpcr.k}, \text{EOL}$
$TPM_{DRTM}(m)$	\equiv	$w = \text{read } m.\text{dpcr.k};$ $r = \text{sign}(dPCR(k), w), AIK^{-1}(m);$ $\text{send } r$
$Verifier(m)$	\equiv	$n = \text{new};$ $\text{send } n;$ $\text{sig} = \text{receive};$ $v = \text{verify } \text{sig}, AIK(m);$ $\text{match } v, (dPCR(k),$ $\quad \text{seq}(dinit, P(m), n, \text{EOL}))$

Figure 4: Security Skeleton for DRTM Attestation Protocol

following assumptions.

$$\Gamma_{DRTM} = \{ \hat{V} \neq AIK(m), \text{Honest}(AIK(m), \{TPM_{SRTM}(m), TPM_{DRTM}(m)\}) \}$$

We also made the same assumptions in the SRTM protocol ($\Gamma_{SRTM} = \Gamma_{DRTM}$). We prove the following theorem:

Theorem 4 (Security of DRTM). *The following is provable in LS^2 : $\Gamma_{DRTM} \vdash J_{DRTM}$*

Proof. See Appendix C. \square

As in the SRTM protocol, the security of the DRTM protocol relies on PCRs being append-only and write-protected in memory. In addition, the DRTM protocol relies on (1) write locks on all dynamic PCRs that are provided by the late launch and (2) a dynamic reset of $m.\text{dpcr.k}$, to reset the values in the dynamic PCRs to *dinit* and signal that $P(m)$ was executed with the protections provided by late launch.

4.2.2 Insights From Analysis

The security analysis lead to a number of insights including revealing an insecure protocol interaction between

the DRTM and SRTM attestation protocols, highlighting differences with the SRTM protocol, and identifying program invariants required for DRTM security that we subsequently used to manually audit a security kernel implementation.

Insecure Protocol Interaction. In extending LS^2 to model DRTM, we discovered that adding late launch required us to weaken some axioms related to reasoning about invariance of values in memory in order to retain soundness in the proof system. With these weaker axioms, we were unable to prove the safety property of the SRTM protocol. Soon after, we realized that SRTM’s safety property can actually be violated using `lateLaunch`. Specifically, during the execution of the SRTM protocol, a late launch instruction may be issued by another thread before $OS(m)$ has been extended into $m.\text{pcr.s}$. The invoked program may then extend the code of the programs $OS(m)$ and $APP(m)$ into $m.\text{pcr.s}$ without executing them, and send signed measurements to the remote verifier. Since the contents of $m.\text{pcr.s}$ would be the sequence $\text{seq}(sinit, BL(m), OS(m), APP(m))$, the remote verifier would believe incorrectly that $OS(m)$ was executed and the SRTM protocol would fail to provide its expected integrity property. This vulnerability can be countered if the program loaded in a DRTM session were unable to change the contents of $m.\text{pcr.s}$ if SRTM were executing in parallel. In the final design of our formal model, we force this to be the case by letting the thread booting a machine to retain an exclusive-write lock on $m.\text{pcr.s}$ even in the face of a concurrent late launch, thus allowing a proof of correctness of SRTM.

Late launch also opens the possibility of a code modification attack on SRTM. Specifically, after the code of a program such as $BL(m)$ or $OS(m)$ has been extended into $m.\text{pcr.s}$ in SRTM, a concurrent thread may invoke a DRTM session and change the code in memory before it is executed. Any subsequent attestation of integrity of the loaded code to a remote party would then be incorrect. Our model prevents this attack by assuming that code measured in PCRs during SRTM cannot be modified in memory.

Comparison to SRTM. The property provided by the DRTM protocol is stronger than the SRTM protocol for a number of reasons:

Fewer Assumptions. The proof of security for the DRTM protocol does not rely on the $\text{ProtectedSRTM}(m)$ assumption that static PCRs are locked. Instead the `lateLaunch` action locks all dynamic PCRs. If the machine M is a multi-processor or multi-core machine that is capable of running mul-

multiple threads in parallel, the locks on the dynamic PCRs will prevent attacks where malicious threads running concurrently with the measurement thread extend additional programs into $m.dpcr.k$ in an attempt to attest to their execution within a late launch session.

Smaller TCB. The security proof of the DRTM does not reason about the measurements of the BIOS, boot loader, or operating system stored in the static PCRs (e.g., $m.pcr.s$), indicating that the security of the DRTM protocol does not depend on these large software components. This considerably reduces the trusted computing base to just $P(m)$ and $LL(m)$ and opens up the possibility of verifying that the TCB satisfies the required program invariants.

Execution Integrity. Unlike the SRTM protocol that does not provide sufficient evidence to deduce that the last program in a sequence of measurements is branched to, the J_{DRTM} property states that all programs measured during the protected session were executed. The property goes further to state that the programs completed execution. Specifically, the end of session measurement EOL proves that $P(m)$ executes to completion.

Program Invariants. In the process of proving the above theorem, we prove program invariants for the roles of the TPM (i.e., $TPM_{SRTM}(m)$ and $TPM_{DRTM}(m)$), and the programs $LL(m)$ and $P(m)$. These invariants specify the properties that $TPM_{SRTM}(m)$, $TPM_{DRTM}(m)$, $LL(m)$, and $P(m)$ must satisfy for the DRTM protocol to be secure. The invariant over the roles of the TPM is similar to the TPM's role invariant used for SRTM. The invariant for $LL(m)$ states that the code must maintain a lock on $m.dpcr.k$ and measure then branch to the program $P(m)$. The invariant for $P(m)$ is shown below. The invariant states that if there are no resets or late launches on m from t_b to t_e , $m.dpcr.k$ is locked at t_b and $m.dpcr.k$ contains the sequence $seq(dinit, P(m))$ at t_b and later contains $seq(dinit, P(m), x, EOL)$, then there exists a thread J such that J extended a value x (e.g., a nonce) into $m.dpcr.k$, then evaluated f , then extended the end of session symbol EOL , and that each action was performed once, in the order specified, and $m.dpcr.k$ was locked from t_b to t_X .

$$\begin{aligned}
& [Q]_J^{t_b, t_e} \forall t, x. ((\neg \text{Reset}(m) \text{ on } (t_b, t_e]) \\
& \quad \wedge (\neg \text{LateLaunch}(m) \text{ on } (t_b, t_e]) \\
& \quad \wedge (\text{Mem}(m.dpcr.k, seq(dinit, P(m))) @ t_b) \\
& \quad \wedge (\text{IsLocked}(m.dpcr.k, J) @ t_b) \\
& \quad \wedge (t_b < t \leq t_e) \\
& \quad \wedge (\text{Mem}(m.dpcr.k, seq(dinit, P(m), x, EOL)) @ t)) \\
& \quad \supset \exists t_n, t_E, t_X. ((t_b < t_n < t_E < t_X < t) \\
& \quad \wedge (\text{Extend}(J, m.dpcr.k, x) @ t_n) \\
& \quad \wedge (\text{Extend}(J, m.dpcr.k, EOL) @ t_X) \\
& \quad \wedge (\text{Eval}(J, f) @ t_E) \\
& \quad \wedge (\neg \text{Eval}(J, f) \text{ on } (t_b, t_E)) \\
& \quad \wedge (\neg \text{Eval}(J, f) \text{ on } (t_E, t_X)) \\
& \quad \wedge (\text{IsLocked}(m.dpcr.k, J) \text{ on } (t_b, t_X)))
\end{aligned}$$

Manual Audit of DRTM Implementation. To check that the invariants required by our security analysis are correct, we performed a manual source code audit of the Flicker implementation of the DRTM protocol [24]. We checked that Flicker's security kernel implementation, represented by our program $P(m)$, respects the invariant above. We were able to quickly extract the security skeleton of the security kernel from Flicker's approximately 250 lines of C code. To verify that the skeleton respects the exact invariant from our security proof, we checked that instructions were present to evaluate the function f , that the EOL marker was subsequently extended into $m.dpcr.k$, and that each of the instructions would only be executed once on all code paths. In several cases, we matched multiple C instructions to a single action since the instructions are a refinement of the action. For example, the extension of EOL consists of two instructions, a `memset` to create the sequence of characters corresponding to an EOL and a call to a wrapper for the extend instruction. The entire manual process of extracting the security skeleton and auditing the invariant took less than one hour for an individual with no previous experience with the Flicker security kernel. Although we did not formally verify the property, one interesting direction for future work is to use these invariants to derive refined invariants to check on the implementation, possibly using software model checking techniques.

5 Related Work

LS^2 draws on certain conceptual ideas from PCL [11], in particular, the local reasoning style by which security properties of protocols are proved without explicitly reasoning about adversary actions. In PCL, global security properties are derived by combining properties achieved

by individual protocol steps with invariants proved by induction over the protocol programs executed by honest parties. LS^2 supports this form of reasoning for a much richer language that includes not only network communication and cryptography as in PCL, but also shared memory, memory protection, machine resets, and dynamically loaded unknown pieces of code. The insights on which the new proof rules are based are described in Section 2.2. The technical definition of LS^2 also differs significantly from PCL: instead of associating pre-conditions and post-conditions with all actions in a process (as PCL does), we model time explicitly, and associate monotonically increasing time points with events on a trace. The presence of explicit time allows us to express invariants about memory; for instance, we may express in LS^2 that a memory location contains the same value throughout the interval $[t_1, t_2]$. Explicit time is also used to reason about the relative order of events. Whereas explicit use of time may appear to be low-level and cumbersome for practical use, the proof system for LS^2 actually uses time in a very limited way that is quite close to temporal logics such as LTL [30]. Indeed, it seems plausible to rework the proof system in this paper using operators of LTL in place of explicit time. However, we refrain from doing so because we believe that a model of real time may be needed to analyze some systems of interest (e.g., [19, 35, 36]).

LS^2 also shares some features with other logics of programs [6, 16, 17]. Hoare logic and dynamic logic focus on sequential imperative programs, and do not consider concurrency, network communication and adversaries. LS^2 's abstract locks are similar to regions that are used to reason about synchronized access to memory in concurrent separation logic [6]. However, the two primitives differ in application. Whereas we use locks to enforce integrity of data stored in memory, regions are intended to prevent race conditions. Another key difference between concurrent separation logic and LS^2 is that the former does not consider network communication. Furthermore, concurrent separation logic and other approaches for verifying concurrent systems [21] typically do not consider an adversary model. An adversary could be encoded as a regular program in these approaches, but then proving invariants would involve an induction over the steps of the honest parties programs and the attacker.

Prior proposals for reasoning about dynamically loaded code use higher-order extensions of Hoare logic [7, 20, 27, 28, 39]. However, they are restricted to reasoning about sequential programs only and require that invariants of code being called be known in the program at the point of the call. LS^2 's method addresses the

problem of reasoning about dynamically loaded code in the more general context of concurrent program execution where one thread is allowed to modify code that is loaded by another. As illustrated in Section 4.1, using the (Jump) rule, evidence that *some* code executed can be combined with separate evidence about the identity of the code to reason precisely about the effects of the jump. Such reasoning is essential in some applications including trusted computing, and is impossible in all prior work known to us.

There have been several previous analyses of trusted computing. Abadi and Wobber used an authorization logic to describe the basic ideas of NGSCB, the predecessor to the TCG [3]. Their formalization documents and clarifies basic NGSCB concepts rather than proving specific properties of systems utilizing a TPM. Chen et al. developed a formal logic tailored to the analysis of a remote attestation protocol and suggested improvements [9]. Unlike LS^2 , these logics are not tied to the execution semantics of the protocols. Gurgens et al. used a model checker to analyze the security of several TCG protocols [15]. Millen et al. employed a model checker to understand the role and trust relationships of a system performing a remote attestation protocol [25]. Our analysis with LS^2 is a complementary approach: It proves security properties even for an infinite number of simultaneous invocations of attestation protocols, but with a more abstract model of the TPM's primitives. LS^2 is designed to be a more general logic with TCG protocols providing one set of applications. Lin [22] used a theorem prover and model finder to analyze the security of the TPM against invalid sequences of API calls.

6 Conclusion

In this paper, we presented LS^2 and used it to carry out a substantial case study of trusted computing attestation protocols. The design of LS^2 was conceptually and technically challenging. Specifically, it was difficult to define a realistic adversary model and formulate sound reasoning principles for dynamically loaded unknown (and untrusted) code. The proof system was designed to support reasoning at a high level of abstraction. This was particularly useful in the case studies where the proofs yielded many insights about the security of trusted computing systems.

In future work, we will build upon this work to model and analyze security properties of web browsers, security hypervisors and virtual machine monitors. We also plan to develop further principles for modeling and reasoning

about security at the level of system interfaces, in particular, to support richer access control models and system composition and refinement.

Acknowledgments. The authors would like to thank Michael Hicks, Jonathan McCune, and the anonymous reviewers for their helpful comments and suggestions. This work was partially supported by the U.S. Army Research Office contract on Perpetually Available and Secure Information Systems (DAAD19-02-1-0389) to CMU’s CyLab, the NSF Science and Technology Center TRUST, and the NSF CyberTrust grant “Realizing Verifiable Security Properties on Untrusted Computing Platforms”. Jason Franklin is supported in part by an NSF Graduate Research Fellowship.

References

- [1] Secure virtual machine architecture reference manual. AMD Corp., May 2005.
- [2] Intel Trusted Execution Technology: Software Development Guide. Document Number: 315168-005, June 2008.
- [3] Martn Abadi and Ted Wobber. A logical account of NGSCB. In *Proceedings of Formal Techniques for Networked and Distributed Systems*, 2004.
- [4] Advanced Micro Devices. AMD64 virtualization: Secure virtual machine architecture reference manual. AMD Publication no. 33047 rev. 3.01, May 2005.
- [5] Torben Braüner and Valeria de Paiva. Towards constructive hybrid logic. In *Electronic Proceedings of Methods for Modalities 3 (M4M3)*, 2003.
- [6] Stephen Brookes. A semantics for concurrent separation logic. In *Proceedings of 15th International Conference on Concurrency Theory*, 2004.
- [7] Hongxu Cai, Zhong Shao, and Alexander Vaynberg. Certified self-modifying code. In *ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 66–77, New York, NY, USA, 2007. ACM.
- [8] Ellick M. Chan, Jeffrey C. Carlyle, Francis M. David, Reza Farivar, and Roy. H. Campbell. Boot-Jacker: Compromising computers using forced restarts. In *Proceedings of 15th ACM Conference on Computer and Communications Security*, 2008.
- [9] Shuyi Chen, Yingyou Wen, and Hong Zhao. Formal analysis of secure bootstrap in trusted computing. In *Proceedings of 4th International Conference on Autonomic and Trusted Computing*, 2007.
- [10] Anupam Datta, Ante Derek, John C. Mitchell, and Dusko Pavlovic. A derivation system and compositional logic for security protocols. *Journal of Computer Security*, 13(3):423–482, 2005.
- [11] Anupam Datta, Ante Derek, John C. Mitchell, and Arnab Roy. Protocol Composition Logic (PCL). *Electr. Notes Theor. Comput. Sci.*, 172:311–358, 2007.
- [12] Henry DeYoung, Deepak Garg, and Frank Pfening. An authorization logic with explicit time. In *Proceedings of the 21st IEEE Computer Security Foundations Symposium (CSF-21)*, June 2008.
- [13] Nancy Durgin, John C. Mitchell, and Dusko Pavlovic. A compositional logic for proving security properties of protocols. *Journal of Computer Security*, 11:677–721, 2003.
- [14] Scott Garriss, Ramón Cáceres, Stefan Berger, Reiner Sailer, Leendert van Doorn, and Xiaolan Zhang. Towards trustworthy kiosk computing. In *Workshop on Mobile Computing Systems and Applications*, February 2006.
- [15] Sigrid Gurgens, Carsten Rudolph, Dirk Scheuermann, Marion Atts, and Rainer Plaga. Security evaluation of scenarios based on the TCG’s TPM specification. In *Proceedings of 12th European Symposium On Research In Computer Security*, 2007.
- [16] David Harel, Dexter Kozen, and Jerzy Tiuryn. *Dynamic Logic*. Foundations of Computing. MIT Press, 2000.
- [17] C. A. R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12(10):576–580, 1969.
- [18] B. Kauer. OSLO: Improving the security of trusted computing. In *Proceedings of the USENIX Security Symposium*, August 2007.
- [19] Rick Kennell and Leah H. Jamieson. Establishing the genuinity of remote computer systems. In *Proceedings of the 2003 USENIX Security Symposium*, August 2003.

- [20] Neel Krishnaswami. Separation logic for a higher-order typed language, 2006. In *Workshop on Semantics, Program Analysis and Computing Environments for Memory Management, SPACE06*.
- [21] Leslie Lamport. The temporal logic of actions. *ACM Transactions on Programming Languages and Systems*, 16(3), May 1994.
- [22] Amerson H. Lin. Automated analysis of security apis. Master’s thesis, Massachusetts Institute of Technology, 2005.
- [23] Z. Manna and A. Pnueli. *Temporal Verification of Reactive Systems: Safety*. Springer-Verlag, 1995.
- [24] Jonathan M. McCune, Bryan Parno, Adrian Perrig, Michael K. Reiter, and Hiroshi Isozaki. Flicker: An execution infrastructure for tcb minimization. In *Proceedings of the ACM European Conference in Computer Systems (EuroSys)*, April 2008.
- [25] Jon Millen, Joshua Guttman, John Ramsdell, Justin Sheehy, and Brian Sniffen. Analysis of a measured launch. Technical report, The MITRE Corporation, 2007.
- [26] Robin Milner, Mads Tofte, and Robert Harper. *The Definition of Standard ML*. MIT Press, Cambridge, MA, USA, 1990.
- [27] Aleksandar Nanevski, Greg Morrisett, and Lars Birkedal. Hoare type theory, polymorphism and separation. *Journal of Functional Programming*, 18(5&6):865–911, 2008.
- [28] Zhaozhong Ni and Zhong Shao. Certified assembly programming with embedded code pointers. In *33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 320–333, New York, NY, USA, 2006. ACM.
- [29] Peter W. O’Hearn, John C. Reynolds, and Hongseok Yang. Local reasoning about programs that alter data structures. In *Proceedings of the 15th International Workshop on Computer Science Logic*, pages 1–19, London, UK, 2001. Springer-Verlag.
- [30] A. Pnueli. The temporal logic of programs. In *Proceedings of 19th Annual Symposium on Foundations on Computer Science*, 1977.
- [31] Jason Reed. Hybridizing a logical framework. In *International Workshop on Hybrid Logic 2006 (HyLo 2006)*, *Electronic Notes in Computer Science*, August 2006.
- [32] John C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Proceeding of the 17th Annual IEEE Symposium on Logic in Computer Science (LICS)*, pages 55–74. IEEE Computer Society, 2002.
- [33] Arnab Roy, Anupam Datta, Ante Derek, John C. Mitchell, and Jean-Pierre Seifert. Secrecy analysis in protocol composition logic. *Formal Logical Methods for System Security and Correctness*, 2008.
- [34] J. Saltzer and M. Schroeder. The protection of information in computer systems. *Proceedings of the IEEE*, 63(9):1278–1308, September 1975.
- [35] Arvind Seshadri, Mark Luk, Elaine Shi, Adrian Perrig, Leendert van Doorn, and Pradeep Khosla. Pioneer: Verifying code integrity and enforcing untampered code execution on legacy platforms. In *Proceedings of ACM Symposium on Operating Systems Principles (SOSP)*, October 2005.
- [36] Arvind Seshadri, Adrian Perrig, Leendert van Doorn, and Pradeep Khosla. SWATT: Software-based attestation for embedded devices. In *Proceedings of the IEEE Symposium on Security and Privacy*, May 2004.
- [37] TCG. PC client specific TPM interface specification (TIS). Version 1.2, Revision 1.00, July 2005.
- [38] Trusted Computing Group (TCG). <https://www.trustedcomputinggroup.org/>, 2008.
- [39] Hayo Thielecke. Frame rules from answer types for code pointers. In *33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 309–319, New York, NY, USA, 2006. ACM.
- [40] Trusted Computing Group. TCG Specification Architecture Overview, Specification Revision 1.4. https://www.trustedcomputinggroup.org/groups/TCG_1.4_Architecture_Overview.pdf, August 2007.

Expressions	e	$::=$	n	Number
			\hat{X}, \hat{Y}	Agent
			K	Key
			K^{-1}	Inverse of key K
			x	Variable
			(e, e')	Pair
			$SIG_K\{e\}$	Value e signed by private key K
			$ENC_K\{e\}$	Value e encrypted by public key K
			$SYMENC_K\{e\}$	Value e encrypted by symmetric key K
			$H(e)$	Hash of e
			P	Program reified as data
Machine	m			
Location	l			
Action	a	$::=$	read l	Read location l
			write l, e	Write e to location l
			extend l, e	Extend PCR l with e
			lock l	Obtain write lock on location l
			unlock l	Release write lock on location l
			send e	Send e as a message
			receive	Receive a message
			sign e, K^{-1}	Sign e with private key K^{-1}
			verify e, K	Check that $e = SIG_{K^{-1}}\{e'\}$
			enc e, K	Encrypt e with public key K
			dec e, K^{-1}	Decrypt e with private key K^{-1}
			symenc e, K	Encrypt e with symmetric key K
			symdec e, K	Decrypt e with symmetric key K
			hash e	Hash the expression e
			eval f, e	Evaluate function f with argument e
			proj₁ e	Project the 1st component of a pair
			proj₂ e	Project the 2nd component of a pair
			match e, e'	Check that $e = e'$
			new	Generate a new nonce
Program	P, Q	$::=$	$\cdot \mid \text{jump } e \mid \text{late_launch} \mid x := a; P$	
Thread id	I, J	$::=$	$\langle \hat{X}, \eta, m \rangle$	
Thread identifier	η			
Thread	T, S	$::=$	$[P]_I$	
Store	σ	$:$	Locations \rightarrow Expressions	
Lock map	ι	$:$	Locations \rightarrow (Thread ids) $\cup \{-\}$	
Configuration	C	$::=$	$\iota, \sigma, T_1 \mid \dots \mid T_n$	

Figure 5: Syntax of the programming language

A Description of LS^2

This appendix summarizes LS^2 's programming language and logic discussed in Section 2, including extensions to trusted computing discussed in Section 3. We start by elaborating the programming language and its semantics, and follow with the logic, its semantics, and proof system. We conclude with the soundness theorem.

A.1 Programming Language

Our language for specifying systems descends from the corresponding language in PCL and extends the latter with constructs for reading, writing, and protecting memory. Its syntax is summarized in Figure 5.

We assume an algebra of expressions (denoted e). Expressions may be numbers n , identities of agents (principals)

\hat{X} , keys K , variables x , pairs (e, e') , signatures using private keys $SIG_K\{e\}$ (denoting the signature on e made using the key K), asymmetric key encryptions $ENC_K\{e\}$, symmetric key encryptions $SYMENC_K\{e\}$, hashes $H(e)$, or code reified as data P . Expressions are assumed to be typed (for e.g., a pair can be distinguished from a number), but we elide the details of the types. We assume that the expression e may be recovered from the signature $SIG_K\{e\}$ if the verification key corresponding to K is known. We also assume that hashes are confidentiality preserving. If needed this may be implemented by combining hashes with encryption.

Agents and keys. Agents, denoted \hat{X}, \hat{Y} , are principals associated with a system (e.g., users). Keys are denoted by the letter K . The inverse of key K is denoted by K^{-1} . As a convention, we use the notation K^{-1} for private keys and the notation K for public keys. If K is a private key, we write \hat{K} to denote the agent who owns it. If K is a public key, we write \hat{K} to denote the agent who owns the corresponding private key. By definition, $\hat{K} = K^{-1}$.

Machines and locations. Machines (denoted m) are the sites of program execution, and the sites that hold locations of memory that are classified into several mutually exclusive categories: RAM, disk, PCRs (also called static PCRs), and dynamic PCRs. Locations are either denoted by the generic letter l , or explicitly by writing the machine, category, and the address of the location, separated by dots. For example, $m.RAM.s$ denotes the s th RAM location on machine m . RAM is replaced by $disk$, pcr , $dpcr$ to represent locations on disk, in static PCRs, and in dynamic PCRs. We also use the wildcard $*$ to represent an arbitrary address or category. For example, $m.*.*$ represents any location on machine m . \bar{l} denotes a set of locations, $locs(m)$ denotes the set of all locations on machine m , and $locs(m, disk)$ denotes the set of all disk locations on machine m . The function $machine(l)$ returns the machine on which location l exists.

Actions and programs. Actions, denoted a , perform specific functions. Allowed actions with their intuitive meanings are listed in Figure 5. Salient among these are actions for reading and writing locations (`read` and `write`), obtaining and releasing locks on locations (`lock` and `unlock`), and extending values into PCRs (`extend`). Actions that change locations may apply to only certain kinds of locations. For example, `write` only applies to RAM and disk locations, whereas `extend` only applies to PCRs. (Such restriction are captured in the operational semantics of the language.) The actions `send` and `receive` perform network communication, which is undirected; an expression sent by any program may be received by any other program. Action `eval` f, e evaluates any side-effect free function f with argument e and returns the result.

If an action (such as signature verification) fails, we assume that the thread executing the action blocks forever. A successfully executed action always returns an expression. For example, the action `receive` returns the expression obtained by synchronizing with another thread, and the action `verify` returns the message contained in the signature it verifies. The expressions returned by some actions like `write`, `send`, and `extend` are unimportant. We assume that these actions always return the constant 0.

A program (denoted P, Q) either terminates (\cdot), or ends by dynamically branching to another program whose code is contained in the expression e (`jump` e), enters a late launch session (`late_launch`), or executes an action a and continues with the program P ($x := a; P$). In the latter case, x is a name for the expression returned by the action a that may be used to refer to the expression in P . The scope of x is P . Variables may be α -varied. In this sense, our treatment of variables is functional, not imperative. We write $P(e/x)$ to denote the program P with e substituted for x . This substitution avoids capture of bound variables through tacit renaming.

Threads and Thread ids. A thread T is a sequentially executing program. Formally, it is a pair containing a program and an identity I , written $[P]_I$. The identity (id) I is a three tuple $\langle \hat{X}, \eta, m \rangle$. \hat{X} is the agent who owns the thread, η is a unique identifier for the thread (akin to a process id), and m is the machine on which the thread executes. For $I = \langle \hat{X}, \eta, m \rangle$, we define $\hat{I} = \hat{X}$ and $machine(I) = m$.

Configurations. A configuration C is the collection of all threads executing concurrently on all machines. Concurrent threads are separated by the symbol $|$, which is assumed to be commutative and associative. In addition to threads, a configuration also contains a *store* σ , which is a map from the set of all locations to the values that they contain, and a *lock map* ι that maps each location to the id of the thread that has a write lock on it. If no thread has a lock on location l , then $\iota(l) = _$. We often write $\sigma[l \mapsto e]$ to denote the map σ augmented with the mapping of l to e . $\iota[l \mapsto I]$ is defined similarly. We assume implicitly that all threads in a configuration are closed, i.e., they do not contain any free variables.

(sign)	$[x := \text{sign } e, K^{-1}; P]_I \longrightarrow [P(\text{SIG}_{K^{-1}}\{e\}/x)]_I$	
(verify)	$[x := \text{verify } \text{SIG}_{K^{-1}}\{e\}, K; P]_I \longrightarrow [P(e/x)]_I$	
(enc)	$[x := \text{enc } e, K; P]_I \longrightarrow [P(\text{ENC}_K\{e\}/x)]_I$	
(dec)	$[x := \text{dec } \text{ENC}_K\{e\}, K^{-1}; P]_I \longrightarrow [P(e/x)]_I$	
(symenc)	$[x := \text{symenc } e, K; P]_I \longrightarrow [P(\text{SYMENC}_K\{e\}/x)]_I$	
(symdec)	$[x := \text{symdec } \text{SYMENC}_K\{e\}, K; P]_I \longrightarrow [P(e/x)]_I$	
(hash)	$[x := \text{hash } e; P]_I \longrightarrow [P(H(e)/x)]_I$	
(eval)	$[x := \text{eval } f, e; P]_I \longrightarrow [P(f(e)/x)]_I$	
(proj1)	$[x := \text{proj}_1(e_1, e_2); P]_I \longrightarrow [P(e_1/x)]_I$	
(proj2)	$[x := \text{proj}_2(e_1, e_2); P]_I \longrightarrow [P(e_2/x)]_I$	
(match)	$[x := \text{match } e, e; P]_I \longrightarrow [P(0/x)]_I$	
(new)	$[x := \text{new}; P]_I \longrightarrow [P(n/x)]_I$	(n fresh)
(read*)	$\sigma, [x := \text{read } l; P]_I \longrightarrow \sigma, [P(e/x)]_I$	($\sigma(l) = e$)
(write)	$\iota, \sigma[l \mapsto e'], [x := \text{write } l, e; P]_I \longrightarrow \iota, \sigma[l \mapsto e], [P(0/x)]_I$	($\text{machine}(I) = m, (l = m.\text{RAM}.* \text{ or } l = m.\text{disk}.*), \iota(l) \in \{I, _ \}$)
(extend)	$\iota, \sigma[l \mapsto \text{seq}(e_1, \dots, e_n)], [x := \text{extend } l, e; P]_I \longrightarrow \iota, \sigma[l \mapsto \text{seq}(e_1, \dots, e_n, e)], [P(0/x)]_I$	($\text{machine}(I) = m, (l = m.\text{pcr}.* \text{ or } l = m.\text{dpcr}.*), \iota(l) \in \{I, _ \}$)
(lock*)	$\iota[l \mapsto _], [x := \text{lock } l; P]_I \longrightarrow \iota[l \mapsto I], [P(0/x)]_I$	
(unlock*)	$\iota[l \mapsto I], [x := \text{unlock } l; P]_I \longrightarrow \iota[l \mapsto _], [P(0/x)]_I$	
(comm)	$[x := \text{send } e; P]_I \mid [y := \text{receive}; Q]_{I'} \longrightarrow [P(0/x)]_I \mid [Q(e/y)]_{I'}$	
(reset)	$\iota, \sigma, T_1 \mid \dots \mid T_n \longrightarrow \iota[\text{locs}(m) \mapsto _], \sigma[(\text{locs}(m) - \text{locs}(m, \text{disk})) \mapsto \text{sinit}], (T_1 \mid \dots \mid T_n) - \{m\} \mid [\text{SRTM}(m)]_I$	($I = \langle \hat{m}, \eta, m \rangle, \eta$ fresh)
(jump)	$[\text{jump } P]_I \longrightarrow [P]_I$	
(llaunch*)	$\iota, \sigma, [\text{late_launch}]_I \longrightarrow \iota[\text{locs}(m, \text{dpcr}) \mapsto J], \sigma[\text{locs}(m, \text{dpcr}) \mapsto \text{dinit}], [\text{LL}(m)]_J$	($\text{machine}(I) = m, J = \langle \hat{m}, \eta, m \rangle, \eta$ fresh)

* Side Condition: $\text{machine}(I) = \text{machine}(I)$

Figure 6: Reduction Rules of the Process Calculus

A.2 Operational Semantics of the Language

The operational semantics of the programming language are defined by reduction rules on configurations, summarized in Figure 6. In each rule, we include only the relevant parts of configurations. Parts not shown in a rule remain unchanged in the reduction.

Rules (sign)–(new) represent internal reductions of a thread. In the rule (new), the value n is a fresh symbolic constant that has never occurred in the history of the configuration. This rule is used for generating nonces. The rules (read) allows reading of locations. The rules (write) and (extend) allow modification of RAM and disk, and PCRs respectively. In each case, the location being modified must either be locked by the thread executing the action, or be unlocked (enforced by the side condition $\iota(l) \in \{I, _ \}$). Locks are obtained and released using the rules (lock) and (unlock). Rule (comm) allows communication between any two threads.

Rule (reset) spontaneously resets any machine m . This rule can “fire” at any time. As a result of the reset, all locations on m are set to the special symbolic value *sinit* (except disk locations, which don’t change), all locks on m are freed, and a new thread I is created to boot the machine using the program $\text{SRTM}(m)$ from Figure 3. Rule (jump) dynamically branches to program P in the thread I . Since no other parts of the configuration are affected, I retains all its locks. Rule (llaunch) ends thread I by starting a late launch session. As a result, a new thread J executing program $\text{LL}(m)$ from Figure 4 is created. All locks on dynamic PCRs on machine m are given to J , and all dynamic PCRs on

Action Predicates	R	::=	Read(I, l, e)	read l getting value e
			Write(I, l, e)	write l, e
			Send(I, e)	send e
			Extend(I, l, e)	extend l, e
			Lock(I, l)	lock l
			Unlock(I, l)	unlock l
			Receive(I, e)	receive receiving e
			Sign(I, e, K)	sign e, K
			Verify(I, e, K)	verify $SIG_{K^{-1}}\{e\}, K$
			Encrypt(I, e, K)	enc e, K
			Decrypt(I, e, K)	dec $ENC_{K^{-1}}\{e\}, K$
			SymEncrypt(I, e, K)	symenc e, K
			SymDecrypt(I, e, K)	symdec $SYMENC_K\{e\}, K$
			Hash(I, e)	hash e
			Eval(I, f, e, e')	eval f, e producing e'
			Match(I, e, e')	match e, e'
			New(I, n)	new generating nonce n
General Predicates	M	::=	Mem(l, e)	Location l contained e
			IsLocked(l, I)	Thread I had a lock on l
			Reset(m, I)	Machine m was reset, booting in thread I
			Jump(I, e)	Thread I ended with jump e
			LateLaunch(m, J)	late_launch on machine m producing thread J
			Contains(e, e')	
			$e = e' \mid t \geq t'$	
			Honest(\hat{X}, \vec{P})	
Formulas	A, B	::=	$R \mid M \mid \top \mid \perp \mid A \wedge B \mid A \vee B \mid$	
			$A \supset B \mid \neg A \mid \forall x.A \mid \exists x.A \mid A @ t$	
Defined Formulas	A on i	=	$\forall t. ((t \in i) \supset (A @ t))$	
	Reset(m)	=	$\exists I. \text{Reset}(m, I)$	
	LateLaunch(m)	=	$\exists I. \text{LateLaunch}(m, I)$	
	Jump(I)	=	$\exists e. \text{Jump}(I, e)$	
	Eval(I, f)	=	$\exists e, e'. \text{Eval}(I, f, e, e')$	
Modal Formulas	J	::=	$[P]_I^{t_b, t_e} A \mid [a]_{I,x}^{t_b, t_e} A$	

Figure 7: Syntax of LS^2

m are set to the symbolic value $dinit$, which is distinct from $sinit$.

Traces and Timed Traces. A trace is a sequence of configurations $C_0 \longrightarrow \dots \longrightarrow C_n$ such that C_{i+1} may be obtained from C_i by one of the reduction rules. A *timed trace* (denoted \mathcal{T}) is a trace in which a *time point* has been associated with each reduction. Time points are drawn from any totally ordered set with a least element $-\infty$ and a greatest element ∞ . We write a timed trace as $C_0 \xrightarrow{t_1} C_1 \dots \xrightarrow{t_n} C_n$. t_1, \dots, t_n represent points of time at which the reductions happened. We require that $t_1 < \dots < t_n$, i.e., the time points be monotonically increasing. It is assumed that the effects of a reduction, such as changes to the store or the lock map, come into effect immediately, i.e., at the time that the reduction happens.

A.3 Syntax of the Logic

The logic LS^2 is used to reason about properties of traces obtained from configurations. Its syntax is summarized in Figure 7. Predicates of the logic are divided into *action predicates* and *general predicates*. Action predicates reason about actions executed by specific threads. There is one action predicate for each action in the programming language, and all these predicates take the id of the executing thread as an argument. Their intuitive meanings are listed in Figure 7. General predicates capture properties not tied to specific threads. Mem(l, e) and IsLocked(l, I) hold whenever location l contains expression e and whenever location l is locked by thread I , respectively. Reset(m, I)

$$\begin{array}{c}
\frac{}{\text{Mayderive}(e, e, \mathcal{K})} \qquad \frac{\text{Mayderive}(e_1, e, \mathcal{K})}{\text{Mayderive}((e_1, e_2), e, \mathcal{K})} \qquad \frac{\text{Mayderive}(e_2, e, \mathcal{K})}{\text{Mayderive}((e_1, e_2), e, \mathcal{K})} \qquad \frac{\text{Mayderive}(e, e', \mathcal{K})}{\text{Mayderive}(\text{SIG}_K\{e\}, e', \mathcal{K})} \\
\\
\frac{}{\text{Mayderive}(\text{SIG}_K\{e\}, K^{-1}, \mathcal{K})} \qquad \frac{\text{Mayderive}(e, e', \mathcal{K})}{\text{Mayderive}(\text{ENC}_K\{e\}, e', \mathcal{K} \cup \{K^{-1}\})} \qquad \frac{}{\text{Mayderive}(\text{ENC}_K\{e\}, K, \mathcal{K})} \\
\\
\frac{\text{Mayderive}(e, e', \mathcal{K})}{\text{Mayderive}(\text{SYMENC}_K\{e\}, e', \mathcal{K} \cup \{K\})}
\end{array}$$

(There are no additional rules for $\text{Mayderive}(e, e', \mathcal{K})$ when $e = H(e'')$)

Figure 8: Semantic definition of predicate Mayderive

holds on a trace at any time when machine m was reset, producing thread I to reboot the machine. $\text{Jump}(I, e)$ holds whenever thread I dynamically branches to code represented by the expression e . $\text{LateLaunch}(m, J)$ holds whenever any thread on machine m initiates a late launch session creating thread J to execute the code $LL(m)$. $\text{Contains}(e, e')$ holds if e' can be derived using cryptographic and projection operations from e . This is formalized using a semantic model of containment, described in the next section. $\text{Honest}(\hat{X}, \hat{P})$ means two things: (a) Agent \hat{X} does not leak its private key, (so its signatures are always authentic), and (b) all threads owned by \hat{X} run one of the programs in the set \hat{P} . We call an agent \hat{X} *honest*, if there is a set of programs \hat{P} (possibly non-finite) such that $\text{Honest}(\hat{X}, \hat{P})$.

Predicates can be combined using the usual logical connectives of classical logic, and the special connective $A @ t$, which means that A holds at time t . There are also a number of *defined formulas* that we use often. These are also listed in the figure.

In addition to the usual formulas A , LS^2 includes two types of *modal formulas* for reasoning about programs. The formula $[P]_I^{t_b, t_e} A$ means that if the thread with id I executes all actions in P in the time interval $(t_b, t_e]$ (and no others), then formula A holds. The related formula $[a]_{I, x}^{t_b, t_e} A$ means that if thread I executes only the action a in the interval $(t_b, t_e]$, returning the result x , then A holds. As a general rule, t_b , t_e , and x are always parameters when we reason in the proof system. The formula A cannot mention variables bound in P . It may however, mention t_b , t_e , and any variables free in P or in a . In the modal formula $[a]_{I, x}^{t_b, t_e} A$, A may mention x also. This allows us to incorporate the result of executing an action into logical reasoning.

A.4 Semantics of the Logic

The formulas of LS^2 are interpreted over timed traces. We assume a priori the principals that are honest, and the programs they execute. Before defining the formal semantics, we define an auxiliary semantic predicate $\text{Mayderive}(e, e', \mathcal{K})$, which formally captures the intuition that e may be used to derive e' in the Dolev-Yao model if all the keys in the set \mathcal{K} are known. This predicate is defined inductively by the rules in Figure 8. Prominent among the rules is the fact that $\text{Mayderive}(H(e), e', \mathcal{K})$ cannot be established unless $e' = H(e)$. This is based on our assumption that the hash function H is confidentiality preserving.

The semantic judgment for formulas is written $\mathcal{T} \models^t A$. It means that A holds in the trace \mathcal{T} at time t . It is defined by induction on the structure of A .

Action Predicates. If A is an action predicate, $\mathcal{T} \models^t A$ holds if a reduction corresponding to A occurred at time t in \mathcal{T} . For example,

$\mathcal{T} \models^t \text{Read}(I, l, e)$ if thread I executed action `read` l at time t , reading e from location l .

$\mathcal{T} \models^t \text{Extend}(I, l, e)$ if thread I executed action `extend` l, e at time t .

$\mathcal{T} \models^t \text{Lock}(I, l)$ if thread I executed action `lock` l at time t .

General Predicates.

$\mathcal{T} \models^t \text{Mem}(l, e)$ if the location l contained e at time t , i.e., at time t , $\sigma(l) = e$.

$\mathcal{T} \models^t \text{IsLocked}(l, I)$ if at time t , $\iota(l) = I$.

$\mathcal{T} \models^t \text{Contains}(e, e')$ if $\text{Maydrive}(e, e', \text{pubs})$ by the rules in Figure 8, where pubs is the set of public keys of all agents.

$\mathcal{T} \models^t \text{Reset}(m, I)$ if at time t , a (reset) reduction on machine m occurs, and the new thread created has id I .

$\mathcal{T} \models^t \text{Jump}(I, P)$ if at time t , thread I reduces jump P .

$\mathcal{T} \models^t \text{LateLaunch}(m, J)$ if at time t , some thread on machine m reduces the action `late_launch` producing the new thread J .

$\mathcal{T} \models^t e = e'$ if e and e' are syntactically equal.

$\mathcal{T} \models^t t_1 \geq t_2$ if $t_1 \geq t_2$ in the total order on time points.

$\mathcal{T} \models^t \text{Honest}(\hat{X}, \vec{P})$ if it is assumed that \hat{X} is honest, and that its threads execute programs in \vec{P} only.

Formulas. Formulas are interpreted in a standard way. The only new case is that for $A @ t$.

$\mathcal{T} \models^t \top$.

$\mathcal{T} \not\models^t \perp$.

$\mathcal{T} \models^t A \wedge B$ if $\mathcal{T} \models^t A$ and $\mathcal{T} \models^t B$.

$\mathcal{T} \models^t A \vee B$ if $\mathcal{T} \models^t A$ or $\mathcal{T} \models^t B$.

$\mathcal{T} \models^t A \supset B$ if $\mathcal{T} \not\models^t A$ or $\mathcal{T} \models^t B$.

$\mathcal{T} \models^t \neg A$ if $\mathcal{T} \not\models^t A$.

$\mathcal{T} \models^t \forall x.A$ if for each ground instance v of x , $\mathcal{T} \models^t A(v/x)$.

$\mathcal{T} \models^t \exists x.A$ if there exists a ground instance v of x such that $\mathcal{T} \models^t A(v/x)$.

$\mathcal{T} \models^t A @ t'$ if $\mathcal{T} \models^{t'} A$.

It should be observed that by definition the relation $\mathcal{T} \models^t A @ t'$ is independent of t . This is consistent with semantics of other hybrid logics (e.g., [5]).

Modal Formulas. For modal formulas, the semantic judgments are written $\mathcal{T} \models [P]_I^{t_b, t_e} A$ and $\mathcal{T} \models [a]_{I,x}^{t_b, t_e} A$. As opposed to the judgment $\mathcal{T} \models^t A$, these judgments are not relativized to time because modal formulas express properties of programs and actions, and are independent of time. Intuitively, $\mathcal{T} \models [P]_I^{t_b, t_e} A$ holds if in the trace \mathcal{T} , either the thread with id I does not execute the sequence of actions P in the interval $(t_b, t_e]$ or A holds. To state this formally, we define a notion of matching between a timed trace \mathcal{T} and a modal prefix.

Matching. We say that a timed trace \mathcal{T} matches the modal prefix $[a]_{I,x}^{t_b, t_e}$ with substitution θ producing value e , written $\mathcal{T} \gg [a]_{I,e}^{t_b, t_e} \mid \theta$ if both the following hold:

- At time t_b , trace \mathcal{T} contains the thread $[y := a\theta; P\theta]_I$. Note that because actions such as jump change the program of a thread, it is possible that at time t_b itself there was an action that changed the program of I to $(y := a\theta; P\theta)$.
- At time t_e , I executes the action $a\theta$, producing value e for y .

Similarly, we say that a timed trace \mathcal{T} matches the modal prefix $[P]_I^{t_b, t_e}$ with substitution θ , written $\mathcal{T} \gg [P]_I^{t_b, t_e} \mid \theta$ if one the following holds:

- $P = \cdot$, and either I does not exist in the trace, or in the interval $(t_b, t_e]$ there is no reduction on thread I (there may, however, be a (reset) on the machine containing thread I). θ is arbitrary in this case.
- $P = \text{jump } e$, at time t_b the trace contains the thread $[\text{jump } e\theta]_I$, and during the interval $(t_b, t_e]$, I performs at least one reduction (namely, by the (jump) rule, loading program $e\theta$).
- $P = \text{late_launch}$, at time t_b the trace contains the thread $[\text{late_launch}]_I$, and during the interval $(t_b, t_e]$, I performs exactly one reduction (namely, by the (llaunch) rule, starting a new protected session).
- $P = (x := a; P')$ and there is a time $t_m \in (t_b, t_e]$ and a ground value e such that $\mathcal{T} \gg [a]_{I,e}^{t_b, t_m} \mid \theta$ and $\mathcal{T} \gg [P']_I^{t_m, t_e} \mid \theta, e/x$.

Informally, $\mathcal{T} \gg [P]_I^{t_b, t_e} \mid \theta$, if at time t_b , \mathcal{T} contains $[P_0\theta]_I$, P is a prefix of the action sequence of P_0 and during the interval $(t_b, t_e]$, exactly this prefix $P\theta$ reduces in thread I . Given this definition of matching, we define semantic satisfaction for modal formulas as follows.

$$\mathcal{T} \models [P]_I^{t_b, t_e} A \text{ if for each } \theta, \text{ and all ground time points } t, t'_b \text{ and } t'_e, \mathcal{T} \gg [P]_I^{t'_b, t'_e} \mid \theta \text{ implies } \mathcal{T} \models^t A\theta(t'_b/t_b)(t'_e/t_e).$$

$$\mathcal{T} \models [a]_{I,x}^{t_b, t_e} A \text{ if for each } \theta, \text{ all ground time points } t, t'_b \text{ and } t'_e, \text{ and each ground } e, \mathcal{T} \gg [a]_{I,e}^{t'_b, t'_e} \mid \theta \text{ implies } \mathcal{T} \models^t A\theta(t'_b/t_b)(t'_e/t_e)(e/x).$$

A.5 Proof System for LS^2

The proof system of LS^2 consists of several rules and axioms. If φ is a formula (modal or otherwise), we write $\vdash \varphi$ to mean that φ is provable using the proof system. For convenience, we have divided the proof system into several parts. Rules of inference are shown in Figure 9. Axioms are shown in Figures 10, 11 and 12, classified (roughly) by their use in reasoning. In addition to these rules and axioms, we assume a full axiomatization of first-order logic, and axioms that make the set of time points a total order. We also assume that equality of expressions is an equivalence relation. Further, we assume some axioms for the predicate $\text{Contains}(e, e')$, often relying on types of terms (for e.g., if e is a number, then $\text{Contains}(e, e') \supset e = e'$). These straightforward axioms are elided here.

Rule (NecAt) states that if A is provable, then so is $A @ t$. This rule is akin to the so called “necessitation” rule from standard modal logics. The basic rule used for reasoning about modal formulas is (Seq). If we know that the program $x := a; P$ was executed in the interval $(t_b, t_e]$, then there must be a time point t_m at which action a reduced. The (Seq) rule lets us reason about a in the interval $(t_b, t_m]$, about the remaining program P in the interval $(t_m, t_e]$, and combine this information. The side condition (t_m fresh) means that t_m should not appear free in A_1 and A_2 and that it should be distinct from both t_b and t_e . Rules (Conj1)–(Nec2) allow us to incorporate reasoning about ordinary formulas into modal formulas.

$IS(\vec{P})$ in the rule (Honesty) denotes programs that are prefixes of programs in the set \vec{P} . Formally, $IS(\cdot) = \{\cdot\}$, $IS(\text{jump } e) = \{\cdot, \text{jump } e\}$, $IS(\text{late_launch}) = \{\cdot, \text{late_launch}\}$, $IS(x := a; P) = \{\cdot\} \cup \{x := a; Q \mid Q \in IS(P)\}$, and $IS(P_1, \dots, P_n) = \{\cdot\} \cup IS(P_1) \dots IS(P_n)$. $IS(\vec{P})$ always includes the empty program \cdot . The rule (Honesty) may be interpreted as follows: if we know that thread I is executing one of the programs in the set \vec{P} (assumption $\text{Honest}(I, \vec{P})$), and on all prefixes of programs in this set some property A holds (premise), then property A must hold (conclusion).

The (Jump) rule is used for reasoning about programs that use dynamic branching. It means that if a property A holds whenever any initial prefix of program P executes (premise), and P is called at time t (assumption $\text{Jump}(I, P) @ t$), then A must hold at all points of time t' greater than t . Rules (Reset) and (LateLaunch) are similar, except that in these cases the programs being executed are fixed ($SRTM(m)$ and $LL(m)$ respectively).

Figure 10 lists those axioms of LS^2 that are used for reasoning about straightline code and for reasoning about network primitives and cryptography. Axioms (K) and (Disj) state basic properties of the connective $@$. These axioms, together with the rule (NecAt) imply that $(A \wedge B) @ t \equiv (A @ t) \wedge (B @ t)$ and that $(A \vee B) @ t \equiv (A @ t) \vee (B @ t)$, where $A \equiv B$ denotes logical equivalence defined as $(A \supset B) \wedge (B \supset A)$. We use the notation $R(I, x, a)$ as an abbreviation for the action predicate corresponding to the action $x := a$, performed by thread with id I . For example, if $a = \text{receive}$, then $R(I, x, a) = \text{Receive}(I, x)$, and if $a = \text{send } e$, then $R(I, x, a) = \text{Send}(I, e)$. As a syntactic convention,

$$\begin{array}{c}
\frac{\vdash A}{\vdash A @ t} \text{NecAt} \qquad \frac{\vdash [a]_{I,x}^{t_b,t_m} A_1 \quad \vdash [P]_I^{t_m,t_e} A_2 \quad (t_m \text{ fresh})}{\vdash [x := a; P]_I^{t_b,t_e} \exists t_m. \exists x. ((t_b < t_m \leq t_e) \wedge A_1 \wedge A_2)} \text{Seq} \\
\\
\frac{\vdash [P]_I^{t_b,t_e} A_1 \quad \vdash [P]_I^{t_b,t_e} A_2}{\vdash [P]_I^{t_b,t_e} A_1 \wedge A_2} \text{Conj1} \qquad \frac{\vdash [a]_{I,x}^{t_b,t_e} A_1 \quad \vdash [a]_{I,x}^{t_b,t_e} A_2}{\vdash [a]_{I,x}^{t_b,t_e} A_1 \wedge A_2} \text{Conj2} \qquad \frac{\vdash [P]_I^{t_b,t_e} A_1 \supset A_2 \quad \vdash [P]_I^{t_b,t_e} A_1}{\vdash [P]_I^{t_b,t_e} A_2} \text{Imp1} \\
\\
\frac{\vdash [a]_{I,x}^{t_b,t_e} A_1 \supset A_2 \quad \vdash [a]_{I,x}^{t_b,t_e} A_1}{\vdash [a]_{I,x}^{t_b,t_e} A_2} \text{Imp2} \qquad \frac{\vdash A}{\vdash [P]_I^{t_b,t_e} A} \text{Nec1} \qquad \frac{\vdash A}{\vdash [a]_{I,x}^{t_b,t_e} A} \text{Nec2} \qquad \frac{\forall Q \in IS(\vec{P}). \vdash [Q]_I^{t_b,t_e} A(t_b, t_e)}{\vdash \text{Honest}(\vec{I}, \vec{P}) \supset \forall t_e. A(-\infty, t_e)} \text{Honesty} \\
\\
\frac{\forall Q \in IS(\text{SRTM}(m)). \vdash [Q]_I^{t_b,t_e} A(t_b, t_e)}{\vdash \text{Reset}(m, I) @ t \supset \forall t'. (t' > t) \supset A(t, t')} \text{Reset} \\
\\
\frac{\text{For every } Q \text{ in } IS(P) : \vdash [Q]_I^{t_b,t_e} A(t_b, t_e) \quad (t_b, t_e \text{ fresh constants})}{\vdash \text{Jump}(I, P) @ t \supset \forall t'. (t' > t) \supset A(t, t')} \text{Jump} \\
\\
\frac{\forall Q \in IS(\text{LL}(m)). \vdash [Q]_I^{t_b,t_e} A(t_b, t_e)}{\vdash \text{LateLaunch}(m, I) @ t \supset \forall t'. (t' > t) \supset A(t, t')} \text{LateLaunch}
\end{array}$$

Figure 9: Proof system rules

we assume that \neg binds tighter than \wedge . Axioms (Act1)–(LL1) are used to reason about specific actions of threads. For example, the axiom (Act) states that if thread I executes the action a at time t_e , and no other action in the interval $(t_b, t_e]$ returning value x , then $R(I, x, a)$ holds at time t_e , and that $R(I, x, a)$ does not hold at any other point in the remaining interval $(t_b, t_e]$.

Axioms (Verify)–(Proj2) describe properties of values returned by actions specific actions. For example, if the action $\text{sig } e, K^{-1}$ returns value x , then x must equal $\text{SIG}_{K^{-1}}\{e\}$ (axiom (Sign)). Axiom (VER) captures the unforgeability of signatures. If a thread I verifies a signature with public key K , and the owner \hat{K} of the corresponding private key is honest, then some thread I' of \hat{K} must have either sent out the signature in a message in the past, or written the signature to a memory location in the past. Axiom (NEW) captures the freshness of nonces: if thread I generates nonce n , and n appears in an action, then the latter must have happened *after* the former.

Figure 11 shows axioms that are used to reason about values contained in locations. Axiom (READ) says that if any thread reads a value e from location l at time t , then l must have contained value e at time t . (Mem=) means that the same location cannot contain two different expressions at the same time. (MemW)–(MemE) capture the effect of specific actions on locations. (MemIR)–(MemIdP) are used to reason that a location l contains the same value throughout an interval i . In each axiom we assume enough conditions to ensure that there is no possibility of modifying l during the interval i . For example, in axiom (MemIR), that applies to locations in RAM, we assume that some thread I which does not write to l during i has a lock on the location, and that the machine on which l is situated is not reset during i . For locations of disk (axiom (MemID)), the condition that the machine is not reset is unnecessary. In case of PCRs (axioms (MemIP) and (MemIdP)), writing is replaced by extension, since PCRs expose a different interface.

Axioms (LockLL) and (LockIdP) capture in the logic how locks may be acquired. If a thread I successfully executes action $\text{lock } l$ at time t , location l is locked for I at time t (axiom (LockL)). When a late launch happens, all dynamic PCRs are automatically locked for the secure thread that is created (axiom (LockLL)). Axioms (LockI) and (LockIdP) are used for reasoning about invariance of locks on locations. These are similar to the axioms for invariance of memory.

Figure 12 lists axioms that are peculiar to PCR locations. Most of these axioms were discussed in Section 3. The axioms (PCRC) and (dPCRC) mean that it is impossible to recover a signature from the contents of any PCR. This follows from the fact that PCRs either contain hashes or constants.

(K)	$\vdash ((A \supset B) @ t) \supset ((A @ t) \supset (B @ t))$	
(Disj)	$\vdash ((A \vee B) @ t) \supset ((A @ t) \vee (B @ t))$	
(Eq)	$\vdash ((e = e') \wedge A(e/x)) \supset A(e'/x)$	
(Act1)	$\vdash [a]_{I,x}^{t_b,t_e} (R(I,x,a) @ t_e) \wedge (\neg R(I,x,a) \text{ on } (t_b,t_e))$	
(Act2)	$\vdash [a]_{I,x}^{t_b,t_e} \neg R(I,x',a') \text{ on } (t_b,t_e) \quad \text{if } x \neq x' \text{ or } a \neq a'$	
(Act3)	$\vdash [a]_{I,x}^{t_b,t_e} (\neg \text{Jump}(I,e) \text{ on } (t_b,t_e]) \wedge (\neg \text{Reset}(\text{machine}(I)) \text{ on } (t_b,t_e])$	
(ActN1)	$\vdash [\cdot]_I^{t_b,t_e} \neg R(I,e,a) \text{ on } (t_b,t_e]$	
(ActN2)	$\vdash [\cdot]_I^{t_b,t_e} \neg \text{Jump}(I,e) \text{ on } (t_b,t_e]$	
(Jump1)	$\vdash [\text{jump } e]_I^{t_b,t_e} \exists t. t \in (t_b,t_e] \wedge (\text{Jump}(I,e) @ t) \wedge (\neg \text{Jump}(I) \text{ on } (t_b,t)) \wedge (\neg \text{Reset}(m) \text{ on } (t_b,t]) \wedge$ $(\neg R(I,e_1,a_1) \text{ on } (t_b,t]) \wedge \dots \wedge (\neg R(I,e_n,a_n) \text{ on } (t_b,t])$	
(LL1)	$\vdash [\text{late_launch}]_I^{t_b,t_e} \exists t. t \in (t_b,t_e] \wedge (\exists J. \text{LateLaunch}(\text{machine}(I),J) @ t) \wedge$ $(\neg \text{Jump}(I) \text{ on } (t_b,t]) \wedge (\neg \text{Reset}(m) \text{ on } (t_b,t]) \wedge$ $(\neg R(I,e_1,a_1) \text{ on } (t_b,t]) \wedge \dots \wedge (\neg R(I,e_n,a_n) \text{ on } (t_b,t])$	
(Verify)	$\vdash [\text{verify } e, K]_{I,x}^{t_b,t_e} e = \text{SIG}_{K^{-1}}\{x\}$	
(Sign)	$\vdash [\text{sign } e, K^{-1}]_{I,x}^{t_b,t_e} x = \text{SIG}_{K^{-1}}\{e\}$	
(Enc)	$\vdash [\text{enc } e, K]_{I,x}^{t_b,t_e} x = \text{ENC}_K\{e\}$	
(Dec)	$\vdash [\text{dec } e, K^{-1}]_{I,x}^{t_b,t_e} e = \text{ENC}_K\{x\}$	
(SymEnc)	$\vdash [\text{symenc } e, K]_{I,x}^{t_b,t_e} x = \text{SYMENC}_K\{e\}$	
(SymDec)	$\vdash [\text{symdec } e, K]_{I,x}^{t_b,t_e} e = \text{SYMENC}_K\{x\}$	
(Hash)	$\vdash [\text{hash } e]_{I,x}^{t_b,t_e} x = H(e)$	
(Eval)	$\vdash [\text{eval } f, e]_{I,x}^{t_b,t_e} x = f(e)$	
(Proj1)	$\vdash [\text{proj}_1 e]_{I,x}^{t_b,t_e} \exists e'. e = (x, e')$	
(Proj2)	$\vdash [\text{proj}_2 e]_{I,x}^{t_b,t_e} \exists e'. e = (e', x)$	
(Match)	$\vdash (\text{Match}(I, e, e') @ t) \supset e = e'$	
(VER)	$\vdash ((\text{Verify}(I, e, K) @ t) \wedge (\hat{I} \neq \hat{K}) \wedge \text{Honest}(\hat{K}, \vec{P}))$ $\supset (\exists I'. \exists t'. \exists e'. (t' < t) \wedge (\hat{I}' = \hat{K}) \wedge \text{Contains}(e', \text{SIG}_{K^{-1}}\{e\})$ $\wedge ((\text{Send}(I', e') @ t') \vee \exists l. (\text{Write}(I', l, e') @ t')))$	
(NEW)	$\vdash ((\text{New}(I, n) @ t) \wedge (R(I', e, a) @ t')) \supset (t' > t)$	$(n \in a)$

Figure 10: Proof system axioms for reasoning about straightline code and network primitives

A.6 Soundness

Lemma 1 (Prefix Matching). *Let \mathcal{T} be a trace that contains the thread $[P\theta]_I$ at time t_b (where P is a program possibly containing free variables, and θ is a substitution that grounds all its free variables). Then for any $t_e \geq t_b$, there is a $Q \in \text{IS}(P)$ such that $\mathcal{T} \gg [Q]_I^{t_b,t_e} \mid \theta$.*

Proof. We prove this theorem by inducting on n , the number of reductions (in \mathcal{T}) of the thread I in the interval $(t_b, t_e]$.

Case $n = 0$. Then the thread I performs no reduction in the interval $(t_b, t_e]$. By definition of matching, $\mathcal{T} \gg [\cdot]_I^{t_b,t_e} \mid \theta$. Hence we may choose $Q = \cdot$ in this case. Note that $\cdot \in \text{IS}(P)$ for any P .

Case $n = m + 1$, $m \geq 0$. Since thread I performs at least one reduction in the interval $(t_b, t_e]$, let the time at which I performs its first reduction after t_b be t_m . We now analyze cases on the rule used for the first reduction in I after t_b :

Subcase (jump). If the first reduction in I after t_b is due to rule (jump), it must be the case that $P = \text{jump } e$, and $P\theta = \text{jump } e\theta$. By definition of matching, it follows that $\mathcal{T} \gg [\text{jump } e]_I^{t_b,t_e} \mid \theta$. Thus we may choose $Q = \text{jump } e$,

(READ)	$\vdash (\text{Read}(I, l, e) @ t) \supset (\text{Mem}(l, e) @ t)$	
(Mem=)	$\vdash ((\text{Mem}(l, e) @ t) \wedge (\text{Mem}(l, e') @ t)) \supset e = e'$	
(MemW)	$\vdash (\text{Write}(I, l, e) @ t) \supset (\text{Mem}(l, e) @ t)$	
(MemR)	$\vdash (\text{Reset}(m, I) @ t) \supset (\text{Mem}(l, \text{init}) @ t)$	$(l = m.*.*, l \neq m.\text{disk}.)$
(MemLL)	$\vdash (\text{LateLaunch}(m, I) @ t) \supset (\text{Mem}(m.\text{dpcr.k}, \text{dinit}) @ t)$	
(MemE)	$\vdash ((\text{Extend}(I, l, e) @ t) \wedge (\text{Mem}(l, \text{seq}(e_1, \dots, e_n)) \text{ on } [t', t]) \wedge (t' < t))$ $\supset (\text{Mem}(l, \text{seq}(e_1, \dots, e_n), e) @ t)$	
(MemIR)	$\vdash ((\text{Mem}(m.\text{RAM.k}, e) @ t_b) \wedge (\text{IsLocked}(m.\text{RAM.k}, I) \text{ on } i)$ $\wedge (\forall e'. \neg \text{Write}(I, m.\text{RAM.k}, e') \text{ on } i) \wedge (\neg \text{Reset}(m) \text{ on } i))$ $\supset (\text{Mem}(m.\text{RAM.k}, e) \text{ on } i)$	$(i = (t_b, t_e) \text{ or } i = (t_b, t_e])$
(MemID)	$\vdash ((\text{Mem}(m.\text{disk.k}, e) @ t_b) \wedge (\text{IsLocked}(m.\text{disk.k}, I) \text{ on } i)$ $\wedge (\forall e'. \neg \text{Write}(I, m.\text{disk.k}, e') \text{ on } i))$ $\supset (\text{Mem}(m.\text{disk.k}, e) \text{ on } i)$	$(i = (t_b, t_e) \text{ or } i = (t_b, t_e])$
(MemIP)	$\vdash ((\text{Mem}(m.\text{pcr.k}, e) @ t_b) \wedge (\text{IsLocked}(m.\text{pcr.k}, I) \text{ on } i)$ $\wedge (\forall e'. \neg \text{Extend}(I, m.l, e') \text{ on } i) \wedge (\neg \text{Reset}(m) \text{ on } i))$ $\supset (\text{Mem}(m.\text{pcr.k}, e) \text{ on } i)$	$(i = (t_b, t_e) \text{ or } i = (t_b, t_e])$
(MemIdP)	$\vdash ((\text{Mem}(m.\text{dpcr.k}, e) @ t_b) \wedge (\text{IsLocked}(m.\text{dpcr.k}, I) \text{ on } i)$ $\wedge (\forall e'. \neg \text{Extend}(I, m.\text{dpcr.k}, e') \text{ on } i) \wedge (\neg \text{Reset}(m) \text{ on } i)$ $\wedge (\neg \text{LateLaunch}(m) \text{ on } i))$ $\supset (\text{Mem}(m.\text{dpcr.k}, e) \text{ on } i)$	$(i = (t_b, t_e) \text{ or } i = (t_b, t_e])$
(LockL)	$\vdash (\text{Lock}(I, l) @ t) \supset (\text{IsLocked}(l, I) @ t)$	
(LockLL)	$\vdash (\text{LateLaunch}(m, I) @ t) \supset (\text{IsLocked}(m.\text{dpcr.k}, I) @ t)$	
(LockI)	$\vdash ((\text{IsLocked}(l, I) @ t_b) \wedge (\neg \text{Unlock}(I, l) \text{ on } i)$ $\wedge (\neg \text{Reset}(m) \text{ on } i))$ $\supset (\text{IsLocked}(l, I) \text{ on } i)$	$(i = (t_b, t_e) \text{ or } i = (t_b, t_e]),$ $l = m.*.*, l \neq m.\text{dpcr}.)$
(LockIdP)	$\vdash ((\text{IsLocked}(m.\text{dpcr.k}, I) @ t_b) \wedge (\neg \text{Unlock}(I, m.\text{dpcr.k}) \text{ on } i)$ $\wedge (\neg \text{Reset}(m) \text{ on } i) \wedge (\neg \text{LateLaunch}(m) \text{ on } i))$ $\supset (\text{IsLocked}(m.\text{dpcr.k}, I) \text{ on } i)$	$(i = (t_b, t_e) \text{ or } i = (t_b, t_e])$

Figure 11: Proof system axioms for reasoning about memory and its protection

which is indeed in the set $IS(P)$ since $P = \text{jump } e$.

Subcase (llaunch). If the first reduction in I after t_b is due to rule (llaunch), it must be the case that $P = \text{late_launch}$, and $P\theta = \text{late_launch } \theta$. By definition of matching, it follows that $\mathcal{T} \gg [\text{late_launch}]_I^{t_b, t_e} \mid \theta$. Thus we may choose $Q = \text{late_launch}$, which is indeed in the set $IS(P)$ since $P = \text{late_launch}$.

Subcase (any action rule, i.e., any of the rules from (sign)–(comm)). If the first reduction in I after t_b is due to an action, it must be the case that $P = (x := a; P')$, and $P\theta = (x := a\theta; P'\theta)$. Further, this first reduction must produce a value, say e , and by our assumption it occurs at time t_m . By definition of matching,

$$\mathcal{T} \gg [a]_{I, e}^{t_b, t_m} \mid \theta \quad (\text{A.1})$$

Further as a result of the reduction, \mathcal{T} contains the thread $[P'\theta(e/x)]_I$ at time t_m . In the interval $(t_m, t_e]$, this thread performs exactly m reductions, which is less than n . Hence by the induction hypothesis, there is a program $Q' \in IS(P')$

$$\begin{aligned}
(\text{PCRC}) \quad & \vdash (\text{Mem}(m.pcr.k, e) @ t) \supset \neg \text{Contains}(e, \text{SIG}_K\{e'\}) \\
(\text{PCR1}) \quad & \vdash (\text{Mem}(m.pcr.k, \text{seq}(sinit, e_1, \dots, e_n)) @ t) \\
& \supset (\exists t'. (t' < t) \wedge (\text{Mem}(m.pcr.k, \text{seq}(sinit, e_1, \dots, e_{n-1})) @ t') \wedge (\neg \text{Reset}(m) \text{ on } [t', t])) \quad (n \geq 1) \\
(\text{PCR2}) \quad & \vdash (\text{Mem}(m.pcr.k, sinit) @ t) \supset (\exists t'. (t' \leq t) \wedge \exists J. (\text{Reset}(m, J) @ t') \wedge (\neg \text{Reset}(m) \text{ on } [t', t])) \\
(\text{PCR}=\) \quad & \vdash ((\text{Mem}(m.pcr.k, \text{seq}(sinit, e_1, \dots, e_n)) @ t) \wedge (\text{Mem}(m.pcr.k, \text{seq}(sinit, e'_1, \dots, e'_n, \dots)) @ t') \\
& \wedge (t' > t) \wedge (\neg \text{Reset}(m) \text{ on } [t, t'])) \\
& \supset ((e_1 = e'_1) \wedge \dots \wedge (e_n = e'_n)) \\
(\text{dPCRC}) \quad & \vdash (\text{Mem}(m.dpcr.k, e) @ t) \supset \neg \text{Contains}(e, \text{SIG}_K\{e'\}) \\
(\text{dPCR1}) \quad & \vdash (\text{Mem}(m.dpcr.k, \text{seq}(dinit, e_1, \dots, e_n)) @ t) \\
& \supset (\exists t'. (t' < t) \wedge (\text{Mem}(m.dpcr.k, \text{seq}(dinit, e_1, \dots, e_{n-1})) @ t') \\
& \wedge (\neg \text{Reset}(m) \text{ on } [t', t]) \wedge (\neg \text{LateLaunch}(m) \text{ on } [t', t])) \quad (n \geq 1) \\
(\text{dPCR2}) \quad & \vdash (\text{Mem}(m.dpcr.k, dinit) @ t) \\
& \supset (\exists t'. (t' \leq t) \wedge \exists J. (\text{LateLaunch}(m, J) @ t') \\
& \wedge (\neg \text{Reset}(m) \text{ on } [t', t]) \wedge (\neg \text{LateLaunch}(m) \text{ on } [t', t])) \\
(\text{dPCR}=\) \quad & \vdash ((\text{Mem}(m.dpcr.k, \text{seq}(dinit, e_1, \dots, e_n)) @ t) \wedge (\text{Mem}(m.dpcr.k, \text{seq}(dinit, e'_1, \dots, e'_n, \dots)) @ t') \\
& \wedge (t' > t) \wedge (\neg \text{Reset}(m) \text{ on } [t, t']) \wedge (\neg \text{LateLaunch}(m) \text{ on } [t, t'])) \\
& \supset ((e_1 = e'_1) \wedge \dots \wedge (e_n = e'_n))
\end{aligned}$$

Figure 12: Proof system axioms for reasoning about PCRs

such that

$$\mathcal{T} \gg [Q]_I^{tm, te} \mid \theta, e/x \quad (\text{A.2})$$

Now observe two facts:

1. $(x := a; Q') \in IS(x := a; P')$ (since $Q' \in IS(P')$)
2. $\mathcal{T} \gg [x := a; Q]_I^{tb, te} \mid \theta$ (from A.1 and A.2)

From (1) and (2) it follows that $Q = (x := a; Q')$ satisfies the requirements of the statement of the theorem. \square

Statement of Soundness. Let Γ denote a set of non-modal formulas. We write $\Gamma \vdash \varphi$ to mean that there is a proof of φ in the proof system of LS^2 , assuming that each assumption in Γ is provable. We define $(\mathcal{T} \models A)$ to mean $(\forall t. \mathcal{T} \models^t A)$, $(\mathcal{T} \models \Gamma)$ to mean $(\forall B \in \Gamma. \mathcal{T} \models B)$, $(\Gamma \models A)$ to mean $(\forall \mathcal{T}. (\mathcal{T} \models \Gamma) \text{ implies } (\mathcal{T} \models A))$, and $(\Gamma \models J)$ to mean $(\forall \mathcal{T}. (\mathcal{T} \models \Gamma) \text{ implies } (\mathcal{T} \models J))$.

The soundness theorem for LS^2 may be stated as follows:

Theorem 5 (Soundness). For $\varphi = A$ or J , if $\Gamma \vdash \varphi$ then $\Gamma \models \varphi$.

Proof of soundness. In proving the soundness theorem, we make the following assumptions about traces.

1. In the starting configuration of any trace, expressions signed by honest agents do not exist in threads of other agents, nor in memory locations. This assumption is needed to prove the soundness of axiom (VER).
2. In the starting configuration of any trace, all PCRs contain values they cannot otherwise contain, i.e., values different from *sinit*, *dinit*, and hashes. This assumption is needed to prove soundness of axioms (PCR1), (PCR2), (dPCR1), and (dPCR2).

To prove soundness, we assume $\Gamma \vdash \varphi$, pick an arbitrary timed trace \mathcal{T} , assume that $\mathcal{T} \models \Gamma$, and show that $\mathcal{T} \models \varphi$. The proof proceeds by induction on the assumed derivation of $\Gamma \vdash \varphi$ in LS^2 's proof system. We case analyze the last rule in the derivation, showing below some of the representative cases.

Case (HYP). This is the case where $\Gamma \vdash A$ because $A \in \Gamma$. Since we have assumed $\mathcal{T} \models \Gamma$, and $A \in \Gamma$, we have $\mathcal{T} \models A$ by definition.

Case (NecAt).

$$\frac{\vdash A}{\vdash A @ t} \text{NecAt}$$

We need to show $\mathcal{T} \models A @ t$, i.e., for any time t' , $\mathcal{T} \models^{t'} A @ t$. By definition of \models , it suffices to show that $\mathcal{T} \models^{t'} A$. By the i.h., $\mathcal{T} \models^{t''} A$ for each t'' . In particular, $\mathcal{T} \models^t A$, as required.

Case (Seq).

$$\frac{\vdash [a]_{I,x}^{t_b,t_m} A_1 \quad \vdash [P]_I^{t_m,t_e} A_2 \quad (t_m \text{ fresh})}{\vdash [x := a; P]_I^{t_b,t_e} \exists t_m. \exists x. ((t_b < t_m \leq t_e) \wedge A_1 \wedge A_2)} \text{Seq}$$

Suppose for some ground time points t'_b and t'_e , $\mathcal{T} \gg [x := a; P]_I^{t'_b,t'_e} \mid \theta$. By definition of matching, there is an expression e and a time $t'_m \in (t'_b, t'_e]$ such that the following hold.

1. $\mathcal{T} \gg [a]_{I,e}^{t'_b,t'_m} \mid \theta$
2. $\mathcal{T} \gg [P]_I^{t'_m,t'_e} \mid \theta, e/x$

From i.h. on the first premise and (1) we obtain that for any t' , $\mathcal{T} \models^{t'} A_1 \theta(t'_b/t_b)(t'_m/t_m)(e/x)$. From i.h. on the second premise and from (2) we obtain that for any t' , $\mathcal{T} \models^{t'} A_2 \theta(e/x)(t'_m/t_m)(t'_e/t_e)$. Further since $t_b \notin A_2$ and $t_e \notin A_1$ (due to syntactic restrictions on modal formulas), we obtain by definition of \models that $\mathcal{T} \models^{t'} (A_1 \wedge A_2) \theta(e/x)(t'_b/t_b)(t'_m/t_m)(t'_e/t_e)$. This immediately implies that $\mathcal{T} \models^{t'} (\exists t_m. \exists x. ((t_b < t_m \leq t_e) \wedge A_1 \wedge A_2)) \theta(t'_b/t_b)(t'_e/t_e)$. This is what we wanted to show.

Case (Conj1).

$$\frac{\vdash [P]_I^{t_b,t_e} A_1 \quad \vdash [P]_I^{t_b,t_e} A_2}{\vdash [P]_I^{t_b,t_e} A_1 \wedge A_2} \text{Conj1}$$

Suppose for some ground time points t'_b and t'_e , $\mathcal{T} \gg [P]_I^{t'_b,t'_e} \mid \theta$. By i.h., for any ground t , $\mathcal{T} \models^t A_1 \theta(t'_b/t_b)(t'_e/t_e)$ and $\mathcal{T} \models^t A_2 \theta(t'_b/t_b)(t'_e/t_e)$. By definition of satisfaction, $\mathcal{T} \models^t (A_1 \wedge A_2) \theta(t'_b/t_b)(t'_e/t_e)$, as required.

Cases (Conj2), (Imp1), (Imp2). These are similar to the previous case.

Case (Nec1).

$$\frac{\vdash A}{\vdash [P]_I^{t_b,t_e} A} \text{Nec1}$$

Suppose that for some t'_b, t'_e , $\mathcal{T} \gg [P]_I^{t'_b,t'_e} \mid \theta$. Given any t , we need to show that $\mathcal{T} \models^t (A \theta)(t'_b/t_b)(t'_e/t_e)$. Since $\vdash A$, $\vdash A \phi$ with an equal derivation for every substitution ϕ (this is a fundamental property of first-order logic). In particular, we may choose $\phi = \theta, t'_b/t_b, t'_e/t_e$ to get $\vdash (A \theta)(t'_b/t_b)(t'_e/t_e)$. By i.h. we obtain $\mathcal{T} \models^t (A \theta)(t'_b/t_b)(t'_e/t_e)$, as required.

Case (Nec2). Similar to above case.

Case (Honesty).

$$\frac{\forall Q \in IS(\vec{P}). \vdash [Q]_I^{t_b, t_e} A(t_b, t_e)}{\vdash \text{Honest}(\hat{I}, \vec{P}) \supset \forall t_e. A(-\infty, t_e)} \text{Honesty}$$

We have to show that for any ground t it is the case that $\mathcal{T} \models^t \text{Honest}(\hat{X}, \vec{P}) \supset \forall t_e. A(-\infty, t_e)$. So, suppose that $\mathcal{T} \models^t \text{Honest}(\hat{X}, \vec{P})$, and pick any ground time point t'_e . It suffices to show that $\mathcal{T} \models^t A(-\infty, t'_e)$.

There are two possibilities. Either there is no thread belonging to the agent \hat{I} in \mathcal{T} , or there is at least one. In the former case, let I be a hypothetical thread belonging to \hat{I} . By definition of matching we have $\mathcal{T} \gg [\cdot]_I^{-\infty, t'_e} | \cdot$. Using this and i.h. on the premise with $Q = \cdot$, we get for any t that $\mathcal{T} \models^t (A(t_b, t_e))(-\infty/t_b)(t'_e/t_e)$, i.e., $\mathcal{T} \models^t A(-\infty, t'_e)$, as required.

If there is at least one thread belonging to agent \hat{I} in \mathcal{T} , choose any one such thread I . Due to the assumption $\mathcal{T} \models^t \text{Honest}(\hat{X}, \vec{P})$, \mathcal{T} must contain $[P]_I$ at time $-\infty$, where $P \in \vec{P}$ is a program. By Lemma 1, there is a $Q \in IS(P)$ such that $\mathcal{T} \gg [Q]_I^{-\infty, t'_e} | \cdot$. Using this fact, and the premise corresponding to exactly this Q , we obtain that $\mathcal{T} \models^t (A(t_b, t_e))(-\infty/t_b)(t'_e/t_e)$, i.e., $\mathcal{T} \models^t A(-\infty, t'_e)$, as required.

Case (Reset).

$$\frac{\forall Q \in IS(\text{SRTM}(m)). \vdash [Q]_I^{t_b, t_e} A(t_b, t_e)}{\vdash \text{Reset}(m, I) @ t \supset \forall t'. (t' > t) \supset A(t, t')} \text{Reset}$$

We have to show that for any ground t_0 it is the case that $\mathcal{T} \models^{t_0} \text{Reset}(m, I) @ t \supset \forall t'. (t' > t) \supset A(t, t')$. So assume that $\mathcal{T} \models^{t_0} \text{Reset}(m, I) @ t$, and pick an arbitrary t' such that $t' > t$. It suffices to show that $\mathcal{T} \models^{t_0} A(t, t')$.

From the assumption $\mathcal{T} \models^{t_0} \text{Reset}(m, I) @ t$, it follows that $\mathcal{T} \models^t \text{Reset}(m, I)$, and hence by definition of \models that the reduction (reset) happened on machine m at time t . Hence at time t , \mathcal{T} contains $[\text{SRTM}(m)]_I$. By Lemma 1, there is a program $Q \in IS(\text{SRTM}(m))$ such that $\mathcal{T} \gg [Q]_I^{t, t'} | \cdot$. Using this, and i.h. on the premise corresponding exactly to this Q , we get, $\mathcal{T} \models^{t_0} (A(t_b, t_e))(t/t_b)(t'/t_e)$, i.e., $\mathcal{T} \models^{t_0} A(t, t')$, as required.

Case (Jump).

$$\frac{\text{For every } Q \in IS(P) : \vdash [Q]_I^{t_b, t_e} A(t_b, t_e) \quad (t_b, t_e \text{ fresh constants})}{\vdash \text{Jump}(I, P) @ t \supset \forall t'. (t' > t) \supset A(t, t')} \text{Jump}$$

We have to show that for any ground t_0 it is the case that $\mathcal{T} \models^{t_0} \text{Jump}(I, P) @ t \supset \forall t'. (t' > t) \supset A(t, t')$. So assume that $\mathcal{T} \models^{t_0} \text{Jump}(I, P) @ t$, and pick an arbitrary t' such that $t' > t$. It suffices to show that $\mathcal{T} \models^{t_0} A(t, t')$.

From the assumption $\mathcal{T} \models^{t_0} \text{Jump}(I, P) @ t$, it follows that $\mathcal{T} \models^t \text{Jump}(I, P)$, and hence by definition of \models that the action jump P reduced in thread I at time t . Hence at time t , \mathcal{T} contains $[P]_I$. By Lemma 1, there is a program $Q \in IS(P)$ such that $\mathcal{T} \gg [Q]_I^{t, t'} | \cdot$. Using this, and i.h. on the premise corresponding exactly to this Q , we get, $\mathcal{T} \models^{t_0} (A(t_b, t_e))(t/t_b)(t'/t_e)$, i.e., $\mathcal{T} \models^{t_0} A(t, t')$, as required.

Case (LateLaunch).

$$\frac{\forall Q \in IS(\text{LL}(m)). \vdash [Q]_I^{t_b, t_e} A(t_b, t_e)}{\vdash \text{LateLaunch}(m, I) @ t \supset \forall t'. (t' > t) \supset A(t, t')} \text{LateLaunch}$$

We have to show that for any ground t_0 it is the case that $\mathcal{T} \models^{t_0} \text{LateLaunch}(m, I) @ t \supset \forall t'. (t' > t) \supset A(t, t')$. So assume that $\mathcal{T} \models^{t_0} \text{LateLaunch}(m, I) @ t$, and pick an arbitrary t' such that $t' > t$. It suffices to show that $\mathcal{T} \models^{t_0} A(t, t')$.

From the assumption $\mathcal{T} \models^{t_0} \text{LateLaunch}(m, I) @ t$, it follows that $\mathcal{T} \models^t \text{LateLaunch}(m, I)$, and hence by definition of \models that the reduction (llaunch) happened on machine m at time t . Hence at time t , \mathcal{T} contains $[\text{LL}(m)]_I$. By Lemma 1, there is a program $Q \in IS(\text{LL}(m))$ such that $\mathcal{T} \gg [Q]_I^{t, t'} | \cdot$. Using this, and i.h. on the premise corresponding exactly to this Q , we get, $\mathcal{T} \models^{t_0} (A(t_b, t_e))(t/t_b)(t'/t_e)$, i.e., $\mathcal{T} \models^{t_0} A(t, t')$, as required.

Case (K).

$$\vdash ((A \supset B) @ t) \supset ((A @ t) \supset (B @ t))$$

We have to show that for each time t' , $\mathcal{T} \models^{t'} ((A \supset B) @ t) \supset ((A @ t) \supset (B @ t))$. Assume that $\mathcal{T} \models^{t'} (A \supset B) @ t$ and $\mathcal{T} \models^{t'} A @ t$. We must show that $\mathcal{T} \models^{t'} B @ t$. From our assumptions and definition of \models it follows that $\mathcal{T} \models^t A \supset B$ and that $\mathcal{T} \models^t A$. Hence, $\mathcal{T} \models^t B$, from which we obtain $\mathcal{T} \models^{t'} B @ t$ as required.

Case (Disj).

$$\vdash ((A \vee B) @ t) \supset ((A @ t) \vee (B @ t))$$

We have to show that for each time t' , $\mathcal{T} \models^{t'} ((A \vee B) @ t) \supset ((A @ t) \vee (B @ t))$. Assume that $\mathcal{T} \models^{t'} (A \vee B) @ t$. It suffices to show that $\mathcal{T} \models^{t'} (A @ t) \vee (B @ t)$. By assumption and definition of \models , we obtain $\mathcal{T} \models^t A \vee B$. Thus either $\mathcal{T} \models^t A$ or $\mathcal{T} \models^t B$. In the former case, we obtain $\mathcal{T} \models^{t'} A @ t$. In the latter case, $\mathcal{T} \models^{t'} B @ t$. In each case, $\mathcal{T} \models^{t'} (A @ t) \vee (B @ t)$, as required.

Case (Eq).

$$\vdash ((e = e') \wedge A(e/x)) \supset A(e'/x)$$

We have to show that for each time t' , $\mathcal{T} \models^{t'} ((e = e') \wedge A(e/x)) \supset A(e'/x)$. Assume that $\mathcal{T} \models^{t'} e = e'$ and that $\mathcal{T} \models^{t'} A(e/x)$. Then, by definition, e and e' are syntactically equal. Thus $A(e/x) = A(e'/x)$. Hence $\mathcal{T} \models^{t'} A(e'/x)$.

Case (Act1).

$$\vdash [a]_{I,x}^{t_b, t_e} \exists t. t \in (t_b, t_e] \wedge (R(I, x, a) @ t) \wedge (\neg R(I, x, a) \text{ on } (t_b, t)) \wedge (\neg R(I, x, a) \text{ on } (t, t_e])$$

Suppose that for some ground t'_b , t'_e and e , $\mathcal{T} \gg [a]_{I,e}^{t'_b, t'_e} | \theta$. We have to show that $\mathcal{T} \models^{t_0} ((R(I, x, a) @ t'_e) \wedge (\neg R(I, x, a) \text{ on } (t'_b, t'_e))) \theta(e/x)$. By definition of matching, $a\theta$ happened in thread I at time t'_e (hence $(R(I, x, a) @ t'_e) \theta(e/x)$ holds), and that no other action happened in thread I at any other time in the interval $(t'_b, t'_e]$. From the latter it follows that $(\neg R(I, x, a) \text{ on } (t'_b, t'_e)) \theta(e/x)$ holds. This is what we had to show.

Cases (Act2) – (ActN2). These are similar to the previous case.

Case (Jump1).

$$\vdash [\text{jump } e]_I^{t_b, t_e} \exists t. t \in (t_b, t] \wedge (\text{Jump}(I, e) @ t) \wedge (\neg \text{Jump}(I) \text{ on } (t_b, t)) \wedge (\neg \text{Reset}(\text{machine}(I)) \text{ on } (t_b, t]) \wedge (\neg R(I, x_1, a_1) \text{ on } (t_b, t]) \wedge \dots \wedge (\neg R(I, x_n, a_n) \text{ on } (t_b, t])$$

Suppose that for some ground t'_b , t'_e , and θ , it is the case that $\mathcal{T} \gg [\text{jump } e]_I^{t'_b, t'_e} | \theta$. Then it suffices to show that $\mathcal{T} \models^{t_0} (\exists t. t \in (t_b, t_e] \wedge (\text{Jump}(I, e) @ t) \wedge (\neg \text{Jump}(I) \text{ on } (t_b, t)) \wedge (\neg \text{Reset}(\text{machine}(I)) \text{ on } (t_b, t]) \wedge (\neg R(I, x_1, a_1) \text{ on } (t_b, t]) \wedge \dots \wedge (\neg R(I, x_n, a_n) \text{ on } (t_b, t])) \theta(t'_b/t_b)(t'_e/t_e)$, i.e., $\mathcal{T} \models^{t_0} \exists t. t \in (t'_b, t'_e] \wedge (\text{Jump}(I, e\theta) @ t) \wedge (\neg \text{Jump}(I) \text{ on } (t'_b, t)) \wedge (\neg \text{Reset}(\text{machine}(I)) \text{ on } (t'_b, t]) \wedge (\neg R(I, x_1\theta, a_1\theta) \text{ on } (t'_b, t]) \wedge \dots \wedge (\neg R(I, x_n\theta, a_n\theta) \text{ on } (t'_b, t])$. By definition of matching, $[\text{jump } e\theta]_I$ exists in \mathcal{T} at time t'_b . Further, this reduces at some time $t \in (t_b, t_e]$ (hence $(\text{Jump}(I, e\theta) @ t)$ holds). Clearly, there cannot be any reduction in I in the interval (t'_b, t) , so $(\neg \text{Jump}(I) \text{ on } (t'_b, t)) \wedge (\neg R(I, x_1\theta, a_1\theta) \text{ on } (t'_b, t]) \wedge \dots \wedge (\neg R(I, x_n\theta, a_n\theta) \text{ on } (t'_b, t])$ also holds. Further there cannot be a reset on $\text{machine}(I)$ in the interval $(t'_b, t]$ because that would have killed thread I , implying that the jump could not have happened. Hence $(\neg \text{Reset}(\text{machine}(I)) \text{ on } (t'_b, t])$ must also hold.

Case (LL1). This is similar to the previous case, except that in place of jump, we have a late launch.

Case (Verify).

$$\vdash [\text{verify } e, K]_{I,x}^{t_b, t_e} e = \text{SIG}_{K-1}\{x\}$$

Suppose that for some ground points t'_b , t'_e , and some ground expression e' , $\mathcal{T} \gg [\text{verify } e, K]_{I,e'}^{t'_b, t'_e} | \theta$. By definition of the match, at time t'_e , the action $\text{verify } e\theta, K\theta$ reduces, returning e' . This forces $e\theta = \text{SIG}_{K-1}\{e'\}$. We now need

to show that $\mathcal{T} \models^t (e = \text{SIG}_{K^{-1}}\{\{x\}\})\theta(e'/x)$, i.e., that $e\theta(e'/x) = \text{SIG}_{K^{-1}\theta(e'/x)}\{e'\}$. But due to syntactic restrictions on the modal formula ($[\text{verify } e, K]_{l,x}^{b,e}$), e and K cannot mention x . Hence it suffices to show $e\theta = \text{SIG}_{K^{-1}\theta}\{e'\}$. We have already established this.

Cases (Sign)–(Proj2). These are similar to the previous case.

Case (Match).

$$\vdash (\text{Match}(I, e, e') @ t) \supset e = e'$$

Assume that $\mathcal{T} \models^{t_0} \text{Match}(I, e, e') @ t$. It suffices to show that $e = e'$. By assumption, we obtain that $\mathcal{T} \models^t \text{Match}(I, e, e')$, or equivalently, that thread I executed action $\text{match } e, e'$ at time t . From the operational semantics, this can happen only if $e = e'$.

Case (VER).

$$\begin{aligned} \vdash & ((\text{Verify}(I, e, K) @ t) \wedge (\hat{I} \neq \hat{K}) \wedge \text{Honest}(\hat{K})) \\ & \supset (\exists I'. \exists t'. \exists e'. (t' < t) \wedge (\hat{I}' = \hat{K}) \wedge \text{Contains}(e', \text{SIG}_{K^{-1}}\{e\}) \\ & \wedge ((\text{Send}(I', e') @ t') \vee \exists l. (\text{Write}(I', l, e') @ t'))) \end{aligned}$$

Suppose $\mathcal{T} \models^{t_0} ((\text{Verify}(I, e, K) @ t) \wedge (\hat{I} \neq \hat{K}) \wedge \text{Honest}(\hat{K}))$. It follows that at time t , I executed the action $x := \text{verify } \text{SIG}_{K^{-1}}\{e\}, K$ in \mathcal{T} . Since in the initial configuration, I cannot contain $\text{SIG}_{\hat{K}}\{e\}$ (because \hat{K} is honest, and $\hat{I} \neq \hat{K}$), at some earlier time point $\text{SIG}_{\hat{K}}\{e\}$ must have appeared in I 's thread for the first time. This could only have happened in two ways: either some other thread sent it to I , or I read it from a location in RAM or on disk (it is impossible to extract a signature from anything written to a PCR). In the latter case, some other thread Y wrote it to the location. In either case, some other thread either sent the signature to I , or wrote it to memory at an earlier time. If this thread belongs to \hat{K} , we are done, else we can repeat the argument on thread Y (the argument terminates because we are moving backwards on the trace, which is finite).

Case (NEW).

$$\vdash ((\text{New}(I, n) @ t) \wedge (R(I', e, a) @ t')) \supset (t' > t) \quad (n \in a)$$

Suppose $\mathcal{T} \models^{t_0} (\text{New}(I, n) @ t) \wedge (\text{Receive}(I', e) @ t')$. By definition, $\mathcal{T} \models^t \text{New}(I, n)$ and $\mathcal{T} \models^{t'} R(I', e, a)$. Thus at time t' , thread I' executed an action a that contained n in it. Suppose for the sake of contradiction that $t' \leq t$. Then, since we assume that actions happen at distinct time points, $t' < t$. It follows that at time t , I executed new resulting in expression n , which existed earlier in the configuration (at time t'). This contradicts the freshness of n in the reduction rule (new). Hence we must have $t' > t$.

Case (READ).

$$\vdash (\text{Read}(I, l, e) @ t) \supset (\text{Mem}(l, e) @ t)$$

Suppose $\mathcal{T} \models^{t_0} \text{Read}(I, l, e) @ t$. By definition, $\mathcal{T} \models^t \text{Read}(I, l, e)$. Again by definition, thread I executed action $\text{read } l$ at time t , obtaining value e . Hence $\sigma(l) = e$ at time t by the side condition on the rule (read). By definition $\mathcal{T} \models^t \text{Mem}(l, e)$. Thus $\mathcal{T} \models^{t_0} \text{Mem}(l, e) @ t$, as required.

Case (Mem=).

$$\vdash ((\text{Mem}(l, e) @ t) \wedge (\text{Mem}(l, e') @ t)) \supset e = e'$$

Suppose that $\mathcal{T} \models^{t_0} (\text{Mem}(l, e) @ t) \wedge (\text{Mem}(l, e') @ t)$. By definition, $\mathcal{T} \models^t \text{Mem}(l, e)$ and $\mathcal{T} \models^t \text{Mem}(l, e')$. Thus in \mathcal{T} , at time t , $\sigma(l) = e$ and $\sigma(l) = e'$. But σ is a function, so e and e' must be syntactically equal. Hence $\mathcal{T} \models^{t_0} e = e'$ as required.

Case (MemW).

$$\vdash (\text{Write}(I, l, e) @ t) \supset (\text{Mem}(l, e) @ t)$$

Let us assume that $\mathcal{T} \models^{t_0} \text{Write}(I, l, e) @ t$, i.e., $\mathcal{T} \models^t \text{Write}(I, l, e)$. It suffices to show that $\mathcal{T} \models^{t_0} \text{Mem}(l, e) @ t$. By assumption, thread I executed `write` l, e at time t . Since the results of an action take effect at the time at which the action occurs, l must contain e at time t . Hence, $\mathcal{T} \models^t \text{Mem}(l, e)$, or equivalently, $\mathcal{T} \models^{t_0} \text{Mem}(l, e) @ t$, which is what we had to show.

Cases (MemR), (MemLL). These are similar to the previous case, except that in these cases, the location of memory is changed due to a (reset) reduction and due to a (llaunch) reduction respectively (instead of (write)).

Case (MemE).

$$\begin{aligned} & \vdash (\text{Extend}(I, l, e) @ t) \wedge (\text{Mem}(l, \text{seq}(e_1, \dots, e_n)) \text{ on } [t', t]) \wedge (t' < t) \\ & \supset (\text{Mem}(l, \text{seq}(e_1, \dots, e_n, e)) @ t) \end{aligned}$$

Let us assume that $\mathcal{T} \models^{t_0} \text{Extend}(I, l, e) @ t$, and that $\mathcal{T} \models^{t_0} \text{Mem}(l, \text{seq}(e_1, \dots, e_n)) \text{ on } [t', t]$ for $t' < t$. It suffices to show that $\mathcal{T} \models^t \text{Mem}(l, \text{seq}(e_1, \dots, e_n, e))$.

From our assumptions we know that at time t , l (which must be a PCR) was extended with value e . Further, just before t , l contained $\text{seq}(e_1, \dots, e_n)$. (The condition $t' < t$ ensures that $[t', t]$ is not empty.) From these and the reduction (extend), it follows that the extension wrote value $\text{seq}(e_1, \dots, e_n, e)$ to l . Thus at time t , l contains $\text{seq}(e_1, \dots, e_n, e)$. Hence by definition of \models we get $\mathcal{T} \models^t \text{Mem}(l, \text{seq}(e_1, \dots, e_n, e))$, as required.

Case (MemIR).

$$\begin{aligned} & \vdash ((\text{Mem}(m.RAM.k, e) @ t_b) \wedge (\text{IsLocked}(m.RAM.k, I) \text{ on } i) \\ & \wedge (\forall e'. \neg \text{Write}(I, m.RAM.k, e') \text{ on } i) \wedge (\neg \text{Reset}(m) \text{ on } i)) \\ & \supset (\text{Mem}(m.RAM.k, e) \text{ on } i) \quad (i = (t_b, t_e) \text{ or } i = (t_b, t_e]) \end{aligned}$$

We consider here the case $i = (t_b, t_e]$. The other case $i = (t_b, t_e)$ is very similar (we just replace the interval $(t_b, t_e]$ by (t_b, t_e)). Assume that $\mathcal{T} \models^{t_0} \text{Mem}(m.RAM.k, e) @ t_b$, that $\mathcal{T} \models^{t_0} \text{IsLocked}(m.RAM.k, I) \text{ on } (t_b, t_e]$, that $\mathcal{T} \models^{t_0} \forall e'. \neg \text{Write}(I, m.RAM.k, e') \text{ on } (t_b, t_e]$, and that $\mathcal{T} \models^{t_0} \neg \text{Reset}(m) \text{ on } (t_b, t_e]$. It suffices to show that $\mathcal{T} \models^{t_0} \text{Mem}(m.RAM.k, e) \text{ on } (t_b, t_e]$.

By assumption, we know that at time t_b , $m.RAM.k$ contains e . We observe that only two rules in the reduction semantics can change the value in $m.RAM.k$. These are (write) and (reset). It follows that the value in $m.RAM.k$ could have changed in the interval $(t_b, t_e]$ only if one of the following happened:

1. In the interval $(t_b, t_e]$, some thread other than I executed `write` $m.RAM.k, e'$ for some e' .
2. In the interval $(t_b, t_e]$, thread I executed `write` $m.RAM.k, e'$ for some e' .
3. In the interval $(t_b, t_e]$, (reset) was applied to machine m .

However each of these possibilities is ruled out by the given assumptions. In particular, (1) is ruled out by the assumption $\mathcal{T} \models^{t_0} \text{IsLocked}(m.RAM.k, I) \text{ on } (t_b, t_e]$, (2) is ruled out by the assumption $\mathcal{T} \models^{t_0} \forall e'. \neg \text{Write}(I, m.RAM.k, e') \text{ on } (t_b, t_e]$, and (3) is ruled out by the assumption $\mathcal{T} \models^{t_0} \neg \text{Reset}(m) \text{ on } (t_b, t_e]$. It follows that the value in $m.RAM.k$ cannot change in the interval $(t_b, t_e]$. Thus we must have $\mathcal{T} \models^{t_0} \text{Mem}(m.RAM.k, e) \text{ on } (t_b, t_e]$ as required.

Cases (MemID), (MemIP), (MemIDP). These are similar to the previous case.

Case (LockL).

$$\vdash (\text{Lock}(I, l) @ t) \supset (\text{IsLocked}(l, I) @ t)$$

Let us assume that $\mathcal{T} \models^{t_0} \text{Lock}(I, l) @ t$, i.e., $\mathcal{T} \models^t \text{Lock}(I, l)$. It suffices to show that $\mathcal{T} \models^{t_0} \text{IsLocked}(l, I) @ t$. By assumption, thread I executed `lock` l at time t . Since the results of an action take effect at the time at which the action occurs, $\iota(l) = I$ at time t . Hence, $\mathcal{T} \models^t \text{IsLocked}(l, I)$, or equivalently, $\mathcal{T} \models^{t_0} \text{IsLocked}(l, I) @ t$, which is what we had to show.

Case (LockLL). This case is similar to the previous case.

Case (LockI).

$$\begin{aligned} & \vdash ((\text{IsLocked}(l, I) @ t_b) \wedge (\neg \text{Unlock}(I, l) \text{ on } i) \\ & \quad \wedge (\neg \text{Reset}(m) \text{ on } i)) \quad (i = (t_b, t_e) \text{ or } i = (t_b, t_e], \\ & \quad \supset (\text{IsLocked}(l, I) \text{ on } i) \quad l = m.*.*, l \neq m.dpcr.*) \end{aligned}$$

We consider here the case $i = (t_b, t_e]$. The other case of $i = (t_b, t_e)$ is similar. Let us assume that $\mathcal{S} \models^{t_0} \text{IsLocked}(l, I) @ t_b$, that $\mathcal{S} \models^{t_0} \neg \text{Unlock}(I, l) \text{ on } (t_b, t_e]$, and that $\mathcal{S} \models^{t_0} \neg \text{Reset}(m) \text{ on } (t_b, t_e]$. It suffices to show that $\mathcal{S} \models^{t_0} \text{IsLocked}(l, I) \text{ on } (t_b, t_e]$. By assumption we know that I has a lock on l at time t_b . There are only three reductions in the operational semantics of the language that can change this: (lock), (unlock) and (reset). (The only other rule (llaunch) that changes the lock map ι is ruled out here because l is not of the form $m.dpcr.*$). This implies that for the lock on l to have changed, one of the following must have happened:

1. In the interval $(t_b, t_e]$, I executed `unlock l`.
2. In the interval $(t_b, t_e]$, some thread other than I executed `lock l`.
3. In the interval $(t_b, t_e]$, rule (reset) happened on machine m .

Suppose, for the sake of contradiction, that one of (1)–(3) does happen. So let t be minimum time in the interval $(t_b, t_e]$ at which one of (1)–(3) happens. Then clearly, since t is the minimum such time, $\iota(l) = I$ at t . If (1) happens at t , it violates the assumption $\mathcal{S} \models^{t_0} \neg \text{Unlock}(I, l) \text{ on } (t_b, t_e]$. If (2) happens at t , then $\iota(l) = _$ at t ; this violates our earlier observation that $\iota(l) = I$ at t . If (3) happens at t , this violates the assumption $\mathcal{S} \models^{t_0} \neg \text{Reset}(m) \text{ on } (t_b, t_e]$. In all cases we get a contradiction. Thus neither of (1)–(3) can hold, and therefore, I must have the lock on location l throughout the interval $(t_b, t_e]$. Accordingly we obtain $\mathcal{S} \models^{t_0} \text{IsLocked}(l, I) \text{ on } (t_b, t_e]$ as required.

Case (LockIdP). The analysis here is similar to the previous case, except that we must also consider a fourth possible reduction rule (llaunch). This possibility is ruled out by the assumption $(\neg \text{LateLaunch}(m) \text{ on } i)$ in the rule.

Case (PCRC).

$$\vdash (\text{Mem}(m.pcr.k, e) @ t) \supset \neg \text{Contains}(e, \text{SIG}_K\{e'\})$$

Suppose $\mathcal{S} \models^{t_0} \text{Mem}(m.pcr.k, e) @ t$, i.e., $\mathcal{S} \models^t \text{Mem}(m.pcr.k, e)$. Thus at time t , $\sigma(m.pcr.k) = e$. Since the value in a static PCR may change only by reset (which writes a constant `sinit` to it), or by an extend (which writes a hash to it), e must have the form `sinit` or $H(e'')$. Assume for the sake of contradiction that it is *not the case that* $\mathcal{S} \models^{t_0} \neg \text{Contains}(e, \text{SIG}_K\{e'\})$. It follows from definition of \models that $\mathcal{S} \models^{t_0} \text{Contains}(e, \text{SIG}_K\{e'\})$. Hence $\text{Mayderive}(e, \text{SIG}_K\{e'\}, \text{pubs})$. But since e is either a hash or a constant, from definition of Mayderive it follows that $e = \text{SIG}_K\{e'\}$. The latter is impossible in our symbolic model. Thus $\mathcal{S} \models^{t_0} \neg \text{Contains}(e, \text{SIG}_K\{e'\})$ must hold.

Case (PCR1).

$$\begin{aligned} & \vdash (\text{Mem}(m.pcr.k, \text{seq}(\text{sinit}, e_1, \dots, e_n)) @ t) \\ & \quad \supset (\exists t'. (t' < t) \wedge (\text{Mem}(m.pcr.k, \text{seq}(\text{sinit}, e_1, \dots, e_{n-1})) @ t') \wedge (\neg \text{Reset}(m) \text{ on } (t', t])) \quad (n \geq 1) \end{aligned}$$

Assume that $\mathcal{S} \models^{t_0} \text{Mem}(m.pcr.k, \text{seq}(\text{sinit}, e_1, \dots, e_n)) @ t$, i.e., $\mathcal{S} \models^t \text{Mem}(m.pcr.k, \text{seq}(\text{sinit}, e_1, \dots, e_n))$. It suffices to show that there is a $t' < t$ such that $\mathcal{S} \models^{t'} \text{Mem}(m.pcr.k, \text{seq}(\text{sinit}, e_1, \dots, e_{n-1}))$ and $\mathcal{S} \models^{t_0} (\neg \text{Reset}(m) \text{ on } (t', t])$. By assumption, $m.pcr.k$ contains $\text{seq}(\text{sinit}, e_1, \dots, e_n)$ at time t . Since we assume that in the starting configuration all PCRs contain a special value that does not exist otherwise, the only way to obtain this value in $m.pcr.k$ is either by a reset or by an extend. Since reset puts `sinit` in a PCR, and $\text{sinit} \neq \text{seq}(\text{sinit}, e_1, \dots, e_n)$ (in our symbolic model under the assumption $n \geq 1$), the value $\text{seq}(\text{sinit}, e_1, \dots, e_n)$ must have appeared in $m.pcr.k$ by an extend reduction at some time $t'' \leq t$. Thus at time t'' , $\sigma(m.pcr.k) = \text{seq}(\text{sinit}, e_1, \dots, e_{n-1})$. Choose t' to be *last time point* before t'' at which either $m.pcr.k$ was extended, or machine m was reset, whichever is later. Clearly, such a t' must exist since the value $\text{seq}(\text{sinit}, e_1, \dots, e_{n-1})$ exists in $m.pcr.k$ at time t'' . Also, the value in $m.pcr.k$ at t' must be

exactly $seq(sinit, e_1, \dots, e_{n-1})$. Further, there cannot be a reset on machine m in the interval (t', t'') (by assumption), nor can there be a reset on machine m at time t'' (since we assumed that there was an extend operation at that time), nor can there be a reset on machine m in the interval $(t'', t]$ (since that would put the value $sinit$ in $m.pcr.k$, but we know that $m.pcr.k$ contains $seq(sinit, e_1, \dots, e_n)$ at time t). Thus there is no reset on m in the interval $(t', t]$.

Case (PCR2).

$$\begin{aligned} & \vdash (\text{Mem}(m.pcr.k, sinit) @ t) \\ & \quad \supset (\exists t'. (t' \leq t) \wedge \exists J. (\text{Reset}(m, J) @ t') \wedge (\neg \text{Reset}(m) \text{ on } (t', t])) \end{aligned}$$

Let us assume that $\mathcal{T} \models^{t_0} \text{Mem}(m.pcr.k, sinit) @ t$, i.e., $\mathcal{T} \models^t \text{Mem}(m.pcr.k, sinit)$. It suffices to show that there is a time $t' \leq t$ and a thread J such that $\mathcal{T} \models^{t'} \text{Reset}(m, J)$, and $\mathcal{T} \models^{t_0} \neg \text{Reset}(m) \text{ on } (t', t]$. By assumption, $m.pcr.k$ contains $sinit$ at time t . Since we assume that in the initial configuration PCRs contain a special value that does not appear otherwise, and because extends cannot write $sinit$ to a PCR, the value $sinit$ must have appeared in $m.pcr.k$ due to a reset. Let $t' \leq t$ be the last time at which the machine m was reset. Clearly then, there must be a J such that $\mathcal{T} \models^{t'} \text{Reset}(m, J)$, and further, there is no reset in the interval $(t', t]$ so $\mathcal{T} \models^{t_0} \neg \text{Reset}(m) \text{ on } (t', t]$ also holds.

Case (PCR=).

$$\begin{aligned} & \vdash ((\text{Mem}(m.pcr.k, seq(sinit, e_1, \dots, e_n)) @ t) \wedge (\text{Mem}(m.pcr.k, seq(sinit, e'_1, \dots, e'_n, \dots)) @ t') \\ & \quad \wedge (t' > t) \wedge (\neg \text{Reset}(m) \text{ on } [t, t'])) \\ & \quad \supset ((e_1 = e'_1) \wedge \dots \wedge (e_n = e'_n)) \end{aligned}$$

Assume that $t' > t$, that $\mathcal{T} \models^{t_0} \text{Mem}(m.pcr.k, seq(sinit, e_1, \dots, e_n)) @ t$, that $\mathcal{T} \models^{t_0} \text{Mem}(m.pcr.k, seq(sinit, e'_1, \dots, e'_n, e'_{n+1}, \dots, e'_m)) @ t'$, and that $\mathcal{T} \models^{t_0} \neg \text{Reset}(m) \text{ on } [t, t']$. It suffices to show that $\mathcal{T} \models^{t_0} (e_1 = e'_1) \wedge \dots \wedge (e_n = e'_n)$. By our assumptions it follows that $m.pcr.k$ contains $seq(sinit, e_1, \dots, e_n)$ at time t , that it contains $seq(sinit, e'_1, \dots, e'_n, e'_{n+1}, \dots, e'_m)$ at a later time t' , and that there is no reset on m in the interim. Thus the only way that $m.pcr.k$ changed between t and t' is by extends. By definition of extension, it follows that $e_i = e'_i$ for $1 \leq i \leq n$, which is what we wanted to show.

Cases (dPCRC), (dPCR1), (dPCR2), (dPCR=). These are similar to cases (PCRC), (PCR1), (PCR2), and (PCR=) respectively.

B Attestation using a Static Root of Trust Measurement (SRTM)

The SRTM protocol is shown in Figure 3. We first prove its measurement property. We remind the reader that we defined $\text{ProtectedSRTM}(m)$ and $\text{MeasuredBoot}_{\text{SRTM}}(m, t)$ as follows.

$$\begin{aligned} \text{ProtectedSRTM}(m) = \\ \forall t, I. (\text{Reset}(m, I) @ t) \supset (\text{IsLocked}(m.pcr.s, I) @ t) \end{aligned}$$

$$\begin{aligned} \text{MeasuredBoot}_{\text{SRTM}}(m, t) = \\ \exists t_T. \exists t_B. \exists t_O. \exists J. (t_T < t_B < t_O < t) \wedge \\ (\text{Reset}(m, J) @ t_T) \wedge (\text{Jump}(J, BL(m)) @ t_B) \wedge \\ (\text{Jump}(J, OS(m)) @ t_O) \wedge (\neg \text{Reset}(m) \text{ on } (t_T, t]) \\ (\neg \text{Jump}(J) \text{ on } (t_T, t_B)) \wedge (\neg \text{Jump}(J) \text{ on } (t_B, t_O)) \end{aligned}$$

Theorem 6 (Security of integrity measurement). *The following is provable in LS^2 :*

$$\begin{aligned} \text{ProtectedSRTM}(m) \vdash \\ \text{Mem}(m.pcr.s, \text{seq}(sinit, BL(m), OS(m), APP(m))) @ t \\ \supset \text{MeasuredBoot}_{\text{SRTM}}(m, t) \end{aligned}$$

Proof. Let us assume $\text{ProtectedSRTM}(m)$ and $\text{Mem}(m.pcr.s, \text{seq}(sinit, BL(m), OS(m), APP(m))) @ t$. We start by using axioms (PCR1) and (PCR2) repeatedly to obtain the following:

$$\begin{aligned} \exists t_T, t_1, t_2, t_3, J. (t_T \leq t_1 < t_2 < t_3 < t) \\ \wedge (\text{Mem}(m.pcr.s, \text{seq}(sinit, BL(m), OS(m), APP(m))) @ t) \\ \wedge (\text{Mem}(m.pcr.s, \text{seq}(sinit, BL(m), OS(m))) @ t_3) \\ \wedge (\text{Mem}(m.pcr.s, \text{seq}(sinit, BL(m))) @ t_2) \\ \wedge (\text{Mem}(m.pcr.s, sinit) @ t_1) \\ \wedge (\text{Reset}(m, J) @ t_T) \\ \wedge (\neg \text{Reset}(m) \text{ on } (t_T, t]) \end{aligned} \tag{B.1}$$

We know from the above that there is a thread J such that $\text{Reset}(m, J) @ t_T$. Next we would like to apply the (Reset) rule. In order to do that, we must prove an invariant of the program $\text{SRTM}(m)$. The specific invariant we prove is that for each $Q \in IS(\text{SRTM}(m))$, it is the case that,

$$\begin{aligned} [Q]_J^{t_b, t_e} \quad \forall t, b, o. \\ ((\text{Reset}(m, J) @ t_b) \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e]) \\ \wedge (t_b < t \leq t_e) \wedge (\text{Mem}(m.pcr.s, \text{seq}(sinit, b, o)) @ t)) \\ \supset \exists t'. ((t_b < t' < t) \wedge (\text{Jump}(J, b) @ t') \\ \wedge (\neg \text{Jump}(J) \text{ on } (t_b, t')) \\ \wedge (\text{Mem}(m.pcr.s, \text{seq}(sinit, b)) @ t') \\ \wedge (\text{IsLocked}(m.pcr.s, J) @ t')) \end{aligned} \tag{B.2}$$

This invariant is not difficult to prove, but the proof is tedious. As an illustration, we verify this invariant for some cases. We start with $Q = \cdot$. In that case, we prove the following stronger property. (The reader may readily verify that this property implies the above for $Q = \cdot$.)

$$\begin{aligned} [\cdot]_J^{t_b, t_e} \quad \forall t, b, o. \\ ((\text{Reset}(m, J) @ t_b) \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e]) \\ \wedge (t_b < t \leq t_e)) \\ \supset (\neg \text{Mem}(m.pcr.s, \text{seq}(sinit, b, o)) @ t) \end{aligned} \tag{B.3}$$

To prove this, we first apply rule (ActN1) to deduce that

$$\begin{aligned} [\cdot]_J^{t_b, t_e} \quad (\neg \text{Unlock}(J, m.pcr.s) \text{ on } (t_b, t_e]) \\ \wedge (\forall e. \neg \text{Extend}(J, m.pcr.s, e) \text{ on } (t_b, t_e]) \end{aligned} \tag{B.4}$$

Further, using the assumption $\text{ProtectedSRTM}(m)$, we get

$$[\cdot]_J^{t_b, t_e} \quad (\text{Reset}(m, J) @ t_b) \supset (\text{IsLocked}(m.pcr.s, J) @ t_b) \quad (\text{B.5})$$

Similarly, using axiom (MemR), we get

$$[\cdot]_J^{t_b, t_e} \quad (\text{Reset}(m, J) @ t_b) \supset (\text{Mem}(m.pcr.s, \text{sinit}) @ t_b) \quad (\text{B.6})$$

Combining B.4, B.5, and B.6, and weakening slightly with an extra assumption gives,

$$\begin{aligned} [\cdot]_J^{t_b, t_e} \quad & ((\text{Reset}(m, J) @ t_b) \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e])) \\ & \supset ((\neg \text{Unlock}(J, m.pcr.s) \text{ on } (t_b, t_e]) \\ & \quad \wedge (\forall e. \neg \text{Extend}(J, m.pcr.s, e) \text{ on } (t_b, t_e]) \\ & \quad \wedge (\text{Mem}(m.pcr.s, \text{sinit}) @ t_b) \\ & \quad \wedge (\text{IsLocked}(m.pcr.s, J) @ t_b)) \end{aligned} \quad (\text{B.7})$$

Using axiom (LockI), we obtain

$$\begin{aligned} [\cdot]_J^{t_b, t_e} \quad & ((\text{Reset}(m, J) @ t_b) \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e])) \\ & \supset ((\neg \text{Unlock}(J, m.pcr.s) \text{ on } (t_b, t_e]) \\ & \quad \wedge (\forall e. \neg \text{Extend}(J, m.pcr.s, e) \text{ on } (t_b, t_e]) \\ & \quad \wedge (\text{Mem}(m.pcr.s, \text{sinit}) @ t_b) \\ & \quad \wedge (\text{IsLocked}(m.pcr.s, J) @ t_b) \\ & \quad \wedge (\text{IsLocked}(m.pcr.s, J) @ (t_b, t_e])) \end{aligned} \quad (\text{B.8})$$

Next, we use axiom (MemIP) and the above formula to deduce that

$$[\cdot]_J^{t_b, t_e} \quad ((\text{Reset}(m, J) @ t_b) \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e])) \supset (\text{Mem}(m.pcr.s, \text{sinit}) \text{ on } (t_b, t_e]) \quad (\text{B.9})$$

Or equivalently, by expanding the definition of A on I ,

$$[\cdot]_J^{t_b, t_e} \quad ((\text{Reset}(m, J) @ t_b) \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e])) \supset (\forall t. (t_b < t \leq t_e) \supset (\text{Mem}(m.pcr.s, \text{sinit}) @ t)) \quad (\text{B.10})$$

Using axiom (Mem=) we get,

$$[\cdot]_J^{t_b, t_e} \quad ((\text{Reset}(m, J) @ t_b) \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e])) \supset (\forall t. (t_b < t \leq t_e) \supset (\forall b, o. \neg \text{Mem}(m.pcr.s, \text{seq}(\text{sinit}, b, o)) @ t)) \quad (\text{B.11})$$

Reorganizing slightly gives us the required property from B.3:

$$\begin{aligned} [\cdot]_J^{t_b, t_e} \quad & \forall t, b, o. \\ & ((\text{Reset}(m, J) @ t_b) \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e])) \\ & \quad \wedge (t_b < t \leq t_e) \\ & \quad \supset (\neg \text{Mem}(m.pcr.s, \text{seq}(\text{sinit}, b, o)) @ t) \end{aligned} \quad (\text{B.12})$$

As another illustrative case, let us take $Q = \text{SRTM}(m)$ (i.e. the whole program). In this case, we will establish the invariant in B.2 directly. First using the rule (Seq), and axioms (Act1) – (Jump1), we deduce that,

$$\begin{aligned} [\text{SRTM}(m)]_J^{t_b, t_e} \quad & (\exists t_E, t_C, b. (t_b < t_E < t_C \leq t_e) \\ & \quad \wedge (\text{Extend}(J, m.pcr.s, b) @ t_E) \\ & \quad \wedge (\forall e. (\neg \text{Extend}(J, m.pcr.s, e) \text{ on } (t_b, t_E]) \\ & \quad \quad \wedge (\neg \text{Extend}(J, m.pcr.s, e) \text{ on } (t_E, t_C])) \\ & \quad \wedge (\text{Jump}(J, b) @ t_C) \\ & \quad \wedge (\neg \text{Jump}(J) \text{ on } (t_b, t_C)) \\ & \quad \wedge (\neg \text{Unlock}(J, m.pcr.s) \text{ on } (t_b, t_C)) \end{aligned} \quad (\text{B.13})$$

We can now weaken this by adding more assumptions,

$$\begin{aligned}
[SRTM(m)]_J^{t_b, t_e} & ((\text{Reset}(m, J) @ t_b) \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e])) \\
& \supset (\exists t_E, t_C, b. (t_b < t_E < t_C \leq t_e) \\
& \quad \wedge (\text{Extend}(J, m.pcr.s, b) @ t_E) \\
& \quad \wedge (\forall e. (\neg \text{Extend}(J, m.pcr.s, e) \text{ on } (t_b, t_E)) \\
& \quad \quad \wedge (\neg \text{Extend}(J, m.pcr.s, e) \text{ on } (t_E, t_C])) \\
& \quad \wedge (\text{Jump}(J, b) @ t_C) \\
& \quad \wedge (\neg \text{Jump}(J) \text{ on } (t_b, t_C)) \\
& \quad \wedge (\neg \text{Unlock}(J, m.pcr.s) \text{ on } (t_b, t_C]))
\end{aligned} \tag{B.14}$$

As in the case of $Q = \cdot$, we use axiom (MemR) and assumption ProtectedSRTM(m) to deduce,

$$\begin{aligned}
[SRTM(m)]_J^{t_b, t_e} & ((\text{Reset}(m, J) @ t_b) \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e])) \\
& \supset (\exists t_E, t_C, b. (t_b < t_E < t_C \leq t_e) \\
& \quad \wedge (\text{Extend}(J, m.pcr.s, b) @ t_E) \\
& \quad \wedge (\forall e. (\neg \text{Extend}(J, m.pcr.s, e) \text{ on } (t_b, t_E)) \\
& \quad \quad \wedge (\neg \text{Extend}(J, m.pcr.s, e) \text{ on } (t_E, t_C])) \\
& \quad \wedge (\text{Jump}(J, b) @ t_C) \\
& \quad \wedge (\neg \text{Jump}(J) \text{ on } (t_b, t_C)) \\
& \quad \wedge (\neg \text{Unlock}(J, m.pcr.s) \text{ on } (t_b, t_C]) \\
& \quad \wedge (\text{Mem}(m.pcr.s, sinit) @ t_b) \\
& \quad \wedge (\text{IsLocked}(m.pcr.s, J) @ t_b))
\end{aligned} \tag{B.15}$$

Using (LockI), we get,

$$\begin{aligned}
[SRTM(m)]_J^{t_b, t_e} & ((\text{Reset}(m, J) @ t_b) \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e])) \\
& \supset (\exists t_E, t_C, b. (t_b < t_E < t_C \leq t_e) \\
& \quad \wedge (\text{Extend}(J, m.pcr.s, b) @ t_E) \\
& \quad \wedge (\forall e. (\neg \text{Extend}(J, m.pcr.s, e) \text{ on } (t_b, t_E)) \\
& \quad \quad \wedge (\neg \text{Extend}(J, m.pcr.s, e) \text{ on } (t_E, t_C])) \\
& \quad \wedge (\text{Jump}(J, b) @ t_C) \\
& \quad \wedge (\neg \text{Jump}(J) \text{ on } (t_b, t_C)) \\
& \quad \wedge (\neg \text{Unlock}(J, m.pcr.s) \text{ on } (t_b, t_C]) \\
& \quad \wedge (\text{Mem}(m.pcr.s, sinit) @ t_b) \\
& \quad \wedge (\text{IsLocked}(m.pcr.s, J) \text{ on } (t_b, t_C]))
\end{aligned} \tag{B.16}$$

Using axiom (MemIP), we obtain,

$$\begin{aligned}
[SRTM(m)]_J^{t_b, t_e} & ((\text{Reset}(m, J) @ t_b) \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e])) \\
& \supset (\exists t_E, t_C, b. (t_b < t_E < t_C \leq t_e) \\
& \quad \wedge (\text{Extend}(J, m.pcr.s, b) @ t_E) \\
& \quad \wedge (\forall e. (\neg \text{Extend}(J, m.pcr.s, e) \text{ on } (t_E, t_C])) \\
& \quad \wedge (\text{Jump}(J, b) @ t_C) \\
& \quad \wedge (\neg \text{Jump}(J) \text{ on } (t_b, t_C)) \\
& \quad \wedge (\neg \text{Unlock}(J, m.pcr.s) \text{ on } (t_b, t_C]) \\
& \quad \wedge (\text{IsLocked}(m.pcr.s, J) \text{ on } (t_b, t_C]) \\
& \quad \wedge (\text{Mem}(m.pcr.s, sinit) \text{ on } [t_b, t_E]))
\end{aligned} \tag{B.17}$$

Next, using axiom (MemE) we obtain,

$$\begin{aligned}
[SRTM(m)]_J^{t_b, t_e} & ((\text{Reset}(m, J) @ t_b) \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e])) \\
& \supset (\exists t_E, t_C, b. (t_b < t_E < t_C \leq t_e) \\
& \quad \wedge (\forall e. (\neg \text{Extend}(J, m.pcr.s, e) \text{ on } (t_E, t_C])) \\
& \quad \wedge (\text{Jump}(J, b) @ t_C) \\
& \quad \wedge (\neg \text{Jump}(J) \text{ on } (t_b, t_C)) \\
& \quad \wedge (\neg \text{Unlock}(J, m.pcr.s) \text{ on } (t_b, t_C]) \\
& \quad \wedge (\text{IsLocked}(m.pcr.s, J) \text{ on } (t_b, t_C]) \\
& \quad \wedge (\text{Mem}(m.pcr.s, sinit) \text{ on } [t_b, t_E)) \\
& \quad \wedge (\text{Mem}(m.pcr.s, seq(sinit, b)) @ t_E))
\end{aligned} \tag{B.18}$$

Using axiom (MemIP) we obtain

$$\begin{aligned}
[SRTM(m)]_J^{t_b, t_e} & ((\text{Reset}(m, J) @ t_b) \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e])) \\
& \supset (\exists t_E, t_C, b. (t_b < t_E < t_C \leq t_e) \\
& \quad \wedge (\text{Jump}(J, b) @ t_C) \\
& \quad \wedge (\neg \text{Jump}(J) \text{ on } (t_b, t_C)) \\
& \quad \wedge (\text{IsLocked}(m.pcr.s, J) \text{ on } (t_b, t_C]) \\
& \quad \wedge (\text{Mem}(m.pcr.s, sinit) \text{ on } [t_b, t_E)) \\
& \quad \wedge (\text{Mem}(m.pcr.s, seq(sinit, b)) @ t_E) \\
& \quad \wedge (\text{Mem}(m.pcr.s, seq(sinit, b)) \text{ on } (t_E, t_C]))
\end{aligned} \tag{B.19}$$

Simplifying slightly, we obtain:

$$\begin{aligned}
[SRTM(m)]_J^{t_b, t_e} & ((\text{Reset}(m, J) @ t_b) \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e])) \\
& \supset (\exists t_C, b. (t_b < t_C \leq t_e) \\
& \quad \wedge (\text{Jump}(J, b) @ t_C) \\
& \quad \wedge (\neg \text{Jump}(J) \text{ on } (t_b, t_C)) \\
& \quad \wedge (\text{IsLocked}(m.pcr.s, J) @ t_C) \\
& \quad \wedge (\text{Mem}(m.pcr.s, seq(sinit, b)) @ t_C) \\
& \quad \wedge ((\text{Mem}(m.pcr.s, sinit) \vee \text{Mem}(m.pcr.s, seq(sinit, b))) \text{ on } (t_b, t_C]))
\end{aligned} \tag{B.20}$$

Next, we use the axiom (Mem=) to deduce that

$$\begin{aligned}
[SRTM(m)]_J^{t_b, t_e} & ((\text{Reset}(m, J) @ t_b) \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e])) \\
& \supset (\exists t_C, b. (t_b < t_C \leq t_e) \\
& \quad \wedge (\text{Jump}(J, b) @ t_C) \\
& \quad \wedge (\neg \text{Jump}(J) \text{ on } (t_b, t_C)) \\
& \quad \wedge (\text{IsLocked}(m.pcr.s, J) @ t_C) \\
& \quad \wedge (\text{Mem}(m.pcr.s, seq(sinit, b)) @ t_C) \\
& \quad \wedge (\forall b', o, t. (t_b < t \leq t_e) \wedge (\text{Mem}(m.pcr.s, seq(sinit, b'), o)) @ t) \\
& \quad \supset (t_C < t \leq t_e))
\end{aligned} \tag{B.21}$$

Using axiom (PCR=), we get

$$\begin{aligned}
[SRTM(m)]_J^{t_b, t_e} & ((\text{Reset}(m, J) @ t_b) \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e])) \\
& \supset (\exists t_C, b. (t_b < t_C \leq t_e) \\
& \quad \wedge (\text{Jump}(J, b) @ t_C) \\
& \quad \wedge (\neg \text{Jump}(J) \text{ on } (t_b, t_C)) \\
& \quad \wedge (\text{IsLocked}(m.pcr.s, J) @ t_C) \\
& \quad \wedge (\text{Mem}(m.pcr.s, seq(sinit, b)) @ t_C) \\
& \quad \wedge (\forall b', o, t. (t_b < t \leq t_e) \wedge (\text{Mem}(m.pcr.s, seq(sinit, b'), o)) @ t) \\
& \quad \supset (t_C < t \leq t_e) \wedge (b' = b))
\end{aligned} \tag{B.22}$$

Now we rotate the existential and universal quantifiers in this formula to obtain,

$$\begin{aligned}
[SRTM(m)]_J^{t_b, t_e} & ((\text{Reset}(m, J) @ t_b) \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e])) \\
& \supset \forall b', o, t. (\exists t_C, b. (t_b < t_C \leq t_e) \\
& \quad \wedge (\text{Jump}(J, b) @ t_C) \\
& \quad \wedge (\neg \text{Jump}(J) \text{ on } (t_b, t_C)) \\
& \quad \wedge (\text{IsLocked}(m.pcr.s, J) @ t_C) \\
& \quad \wedge (\text{Mem}(m.pcr.s, seq(sinit, b)) @ t_C) \\
& \quad \wedge ((t_b < t \leq t_e) \wedge (\text{Mem}(m.pcr.s, seq(sinit, b'), o)) @ t) \\
& \quad \supset (t_C < t \leq t_e) \wedge (b' = b))
\end{aligned} \tag{B.23}$$

Reorganizing,

$$\begin{aligned}
[SRTM(m)]_J^{t_b, t_e} & \forall t, b', o. \\
& ((\text{Reset}(m, J) @ t_b) \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e])) \\
& \wedge (t_b < t \leq t_e) \wedge (\text{Mem}(m.pcr.s, seq(sinit, b'), o)) @ t) \\
& \supset (\exists t_C, b. (t_b < t_C \leq t_e) \\
& \quad \wedge (\text{Jump}(J, b) @ t_C) \\
& \quad \wedge (\neg \text{Jump}(J) \text{ on } (t_b, t_C)) \\
& \quad \wedge (\text{IsLocked}(m.pcr.s, J) @ t_C) \\
& \quad \wedge (\text{Mem}(m.pcr.s, seq(sinit, b)) @ t_C) \\
& \quad \wedge (t_C < t \leq t_e) \wedge (b' = b))
\end{aligned} \tag{B.24}$$

Simplifying using axiom (Eq), we get

$$\begin{aligned}
[SRTM(m)]_J^{t_b, t_e} & \forall t, b', o. \\
& ((\text{Reset}(m, J) @ t_b) \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e])) \\
& \wedge (t_b < t \leq t_e) \wedge (\text{Mem}(m.pcr.s, seq(sinit, b'), o)) @ t) \\
& \supset (\exists t_C, b. (t_b < t_C < t \leq t_e) \\
& \quad \wedge (\text{Jump}(J, b') @ t_C) \\
& \quad \wedge (\neg \text{Jump}(J) \text{ on } (t_b, t_C)) \\
& \quad \wedge (\text{IsLocked}(m.pcr.s, J) @ t_C) \\
& \quad \wedge (\text{Mem}(m.pcr.s, seq(sinit, b')) @ t_C))
\end{aligned} \tag{B.25}$$

Dropping some of the unnecessary parts, we get,

$$\begin{aligned}
[SRTM(m)]_J^{t_b, t_e} & \forall t, b', o. \\
& ((\text{Reset}(m, J) @ t_b) \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e])) \\
& \wedge (t_b < t \leq t_e) \wedge (\text{Mem}(m.pcr.s, seq(sinit, b'), o)) @ t) \\
& \supset (\exists t_C. (t_b < t_C < t) \\
& \quad \wedge (\text{Jump}(J, b') @ t_C) \\
& \quad \wedge (\neg \text{Jump}(J) \text{ on } (t_b, t_C)) \\
& \quad \wedge (\text{IsLocked}(m.pcr.s, J) @ t_C) \\
& \quad \wedge (\text{Mem}(m.pcr.s, seq(sinit, b')) @ t_C))
\end{aligned} \tag{B.26}$$

This is exactly what we set out to prove in B.2, except that b and t' have been α renamed to b' and t_C respectively. In this manner, we can establish the invariant of B.2. Thus using the (Reset) rule we deduce,

$$\begin{aligned}
\text{Reset}(m, J) @ t_b & \supset \forall t_e. (t_e > t_b) \supset \forall t, b, o. \\
& ((\text{Reset}(m, J) @ t_b) \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e])) \\
& \wedge (t_b < t \leq t_e) \wedge (\text{Mem}(m.pcr.s, seq(sinit, b, o)) @ t) \\
& \supset \exists t'. ((t_b < t' < t) \wedge (\text{Jump}(J, b) @ t') \\
& \quad \wedge (\neg \text{Jump}(J) \text{ on } (t_b, t')) \\
& \quad \wedge (\text{Mem}(m.pcr.s, seq(sinit, b)) @ t') \\
& \quad \wedge (\text{IsLocked}(m.pcr.s, J) @ t'))
\end{aligned} \tag{B.27}$$

Now we instantiate the above formula with values from B.1. We choose $t_b = t_T$, $t_e = t$, $t = t_3$, $b = BL(m)$, $o = OS(m)$, and α -rename t' to t_B to obtain,

$$\begin{aligned}
\text{Reset}(m, J) @ t_T \supset (t > t_T) \supset & \\
& \wedge ((\text{Reset}(m, J) @ t_T) \wedge (\neg \text{Reset}(m) \text{ on } (t_T, t])) \\
& \wedge (t_T < t_3 \leq t) \wedge (\text{Mem}(m.pcr.s, seq(sinit, BL(m), OS(m))) @ t_3)) \\
& \supset \exists t_B. ((t_T < t_B < t_3) \wedge (\text{Jump}(J, BL(m)) @ t_B) \\
& \quad \wedge (\neg \text{Jump}(J) \text{ on } (t_T, t_B)) \\
& \quad \wedge (\text{Mem}(m.pcr.s, seq(sinit, BL(m))) @ t_B) \\
& \quad \wedge (\text{IsLocked}(m.pcr.s, J) @ t_B))
\end{aligned} \tag{B.28}$$

Combining with B.1 and simplifying, we obtain,

$$\begin{aligned}
& \exists t_T, t_B, J. (t_T < t_B < t) \\
& \wedge (\text{Mem}(m.pcr.s, seq(sinit, BL(m), OS(m), APP(m))) @ t) \\
& \wedge (\neg \text{Reset}(m) \text{ on } (t_T, t]) \\
& \wedge (\text{Mem}(m.pcr.s, seq(sinit, BL(m))) @ t_B) \\
& \wedge (\text{Reset}(m, J) @ t_T) \\
& \wedge (\text{Jump}(J, BL(m)) @ t_B) \\
& \wedge (\neg \text{Jump}(J) \text{ on } (t_T, t_B)) \\
& \wedge (\text{IsLocked}(m.pcr.s, J) @ t_B)
\end{aligned} \tag{B.29}$$

Next we want to use the (Jump) rule to reason from the assumption $\text{Jump}(J, BL(m)) @ t_B$. In this case, we want to prove an invariant about the program $BL(m)$. Specifically, we want to show that for each $Q \in IS(BL(m))$, it is the case that,

$$\begin{aligned}
[Q]_J^{t_b, t_e} \quad \forall t, o, a. & \\
& ((\neg \text{Reset}(m) \text{ on } (t_b, t_e]) \\
& \wedge (\text{Mem}(m.pcr.s, seq(sinit, BL(m))) @ t_b) \\
& \wedge (\text{IsLocked}(m.pcr.s, J) @ t_b) \\
& \wedge (t_b < t \leq t_e) \wedge (\text{Mem}(m.pcr.s, seq(sinit, BL(m), o, a)) @ t)) \\
& \supset \exists t'. ((t_b < t' < t) \wedge (\text{Jump}(J, o) @ t') \wedge (\neg \text{Jump}(J) \text{ on } (t_b, t')))
\end{aligned} \tag{B.30}$$

This invariant can be proved almost as we proved B.2 for $SRTM(m)$. This is because the two programs $SRTM(m)$ and $BL(m)$ are very similar. The main difference is that we do not need to use the assumption ProtectedSRTM(m) here. In fact the proof is simpler. Having established this invariant, we use the rule (Jump) to deduce that,

$$\begin{aligned}
\text{Jump}(J, BL(m)) @ t_b \supset \quad \forall t_e. (t_e > t_b) \supset \forall t, o, a. & \\
& ((\neg \text{Reset}(m) \text{ on } (t_b, t_e]) \\
& \wedge (\text{Mem}(m.pcr.s, seq(sinit, BL(m))) @ t_b) \\
& \wedge (\text{IsLocked}(m.pcr.s, J) @ t_b) \\
& \wedge (t_b < t \leq t_e) \wedge (\text{Mem}(m.pcr.s, seq(sinit, BL(m), o, a)) @ t)) \\
& \supset \exists t'. ((t_b < t' < t) \wedge (\text{Jump}(J, o) @ t') \wedge (\neg \text{Jump}(J) \text{ on } (t_b, t')))
\end{aligned} \tag{B.31}$$

We now instantiate this formula with values from B.29. We choose $t_b = t_B$, $t_e = t$, $t = t$, $o = OS(m)$, $a = APP(m)$, and α -rename t' to t_O to get,

$$\begin{aligned}
\text{Jump}(J, BL(m)) @ t_B \supset (t > t_B) \supset & \\
& ((\neg \text{Reset}(m) \text{ on } (t_B, t]) \\
& \wedge (\text{Mem}(m.pcr.s, seq(sinit, BL(m))) @ t_B) \\
& \wedge (\text{IsLocked}(m.pcr.s, J) @ t_B) \\
& \wedge (t_B < t \leq t) \\
& \wedge (\text{Mem}(m.pcr.s, seq(sinit, BL(m), OS(m), APP(m))) @ t)) \\
& \supset \exists t_O. ((t_B < t_O < t) \wedge (\text{Jump}(J, OS(m)) @ t_O) \wedge (\neg \text{Jump}(J) \text{ on } (t_B, t_O)))
\end{aligned} \tag{B.32}$$

Combining with B.29 and simplifying we get,

$$\begin{aligned}
& \exists t_T, t_B, t_O, J. (t_T < t_B < t_O < t) \\
& \wedge (\text{Mem}(m.pcr.s, seq(sinit, BL(m), OS(m), APP(m))) @ t) \\
& \wedge (\neg \text{Reset}(m) \text{ on } (t_T, t]) \\
& \wedge (\text{Mem}(m.pcr.s, seq(sinit, BL(m))) @ t_B) \\
& \wedge (\text{Reset}(m, J) @ t_T) \\
& \wedge (\text{Jump}(J, BL(m)) @ t_B) \\
& \wedge (\neg \text{Jump}(J) \text{ on } (t_T, t_B)) \\
& \wedge (\text{IsLocked}(m.pcr.s, J) @ t_B) \\
& \wedge (\text{Jump}(J, OS(m)) @ t_O) \\
& \wedge (\neg \text{Jump}(J) \text{ on } (t_B, t_O))
\end{aligned} \tag{B.33}$$

Simplifying,

$$\begin{aligned}
& \exists t_T, t_B, t_O, J. (t_T < t_B < t_O < t) \\
& \wedge (\neg \text{Reset}(m) \text{ on } (t_T, t]) \\
& \wedge (\text{Reset}(m, J) @ t_T) \\
& \wedge (\text{Jump}(J, BL(m)) @ t_B) \\
& \wedge (\neg \text{Jump}(J) \text{ on } (t_T, t_B)) \\
& \wedge (\text{Jump}(J, OS(m)) @ t_O) \\
& \wedge (\neg \text{Jump}(J) \text{ on } (t_B, t_O))
\end{aligned} \tag{B.34}$$

This is what we set out to prove. □

Next, we turn to the integrity reporting protocol. We define the assumptions:

$$\begin{aligned}
\Gamma_{SRTM} = & \\
& \{\hat{V} \neq AI\hat{K}(m), \\
& \text{Honest}(AI\hat{K}(m), \{TPM_{SRTM}(m), TPM_{DRTM}(m)\})\}
\end{aligned}$$

Theorem 7 (Correctness of integrity reporting). *The following are provable in LS^2 :*

$$\begin{aligned}
\Gamma_{SRTM} \vdash & \\
& [\text{Verifier}(m)]_V^{t_b, t_e} \exists t. (t < t_e) \wedge \\
& (\text{Mem}(m.pcr.s, seq(sinit, BL(m), OS(m), APP(m))) @ t)
\end{aligned}$$

$$\begin{aligned}
\Gamma_{SRTM}, \text{ProtectedSRTM}(m) \vdash & \\
& [\text{Verifier}(m)]_V^{t_b, t_e} \exists t. (t < t_e) \wedge \text{MeasuredBoot}_{SRTM}(m, t)
\end{aligned}$$

Proof. We prove the first property. The second follows immediately from the first property and correctness of measurement. We begin by analyzing the code of $\text{Verifier}(m)$ using the rule (Seq) and the axiom (Act1). In two steps we deduce,

$$\begin{aligned}
[\text{Verifier}(m)]_V^{t_b, t_e} \exists t_m. (t_b < t_m \leq t_e) \\
\wedge \text{Verify}(V, (PCR(s), seq(sinit, BL(m), OS(m), APP(m))), AIK(m)) @ t_m
\end{aligned} \tag{B.35}$$

Using axiom (VER), and the assumptions $\text{Honest}(AI\hat{K}(m), \{TPM_{SRTM}(m), TPM_{DRTM}(m)\})$ and $\hat{V} \neq AI\hat{K}(m)$ together with the above formula, we deduce:

$$\begin{aligned}
[\text{Verifier}(m)]_V^{t_b, t_e} \exists t_m. (t_b < t_m \leq t_e) \\
\wedge \text{Verify}(V, (PCR(s), seq(sinit, BL(m), OS(m), APP(m))), AIK(m)) @ t_m \\
\wedge \exists I', t', e'. (t' < t_m) \wedge (\hat{I}' = AI\hat{K}(m)) \\
\wedge \text{Contains}(e', SIG_{AIK(m)-1}\{(PCR(s), seq(sinit, BL(m), OS(m), APP(m)))\}) \\
\wedge ((\text{Send}(I', e') @ t') \vee \exists l. (\text{Write}(I', l, e') @ t'))
\end{aligned} \tag{B.36}$$

Simplifying slightly, and α -renaming t' to t_S (S for ‘send’), I' to I , and e' to e , we obtain,

$$\begin{aligned} [\text{Verifier}(m)]_V^{t_b, t_e} \quad & \exists t_S, e, I. (t_S < t_e) \wedge \hat{I} = \text{AIK}(m) \\ & \wedge \text{Contains}(e, \text{SIG}_{\text{AIK}(m)^{-1}}\{\{(PCR(s), \text{seq}(\text{sinit}, BL(m), OS(m), APP(m)))\}\}) \\ & \wedge ((\text{Send}(I, e) @ t_S) \vee \exists l. (\text{Write}(I, l, e) @ t_S)) \end{aligned} \quad (\text{B.37})$$

Next we wish to prove an invariant about the programs that are executed by the TPM, namely $TPM_{SRTM}(m)$ and $TPM_{DRTM}(m)$. We want to show that for each $Q \in IS(TPM_{SRTM}(m), TPM_{DRTM}(m))$, it is the case that

$$\begin{aligned} [Q]_I^{t_b, t_e} \quad & (\forall l, e, t. (t \in (t_b, t_e]) \supset \neg \text{Write}(I, l, e) @ t) \\ & \wedge (\forall t', e'. ((t' \in (t_b, t_e]) \wedge \text{Send}(I, e') @ t') \\ & \quad \supset (\exists e'', t_R. (t_R < t') \wedge (\text{Read}(I, m.pcr.s, e'') @ t_R) \wedge e' = \text{SIG}_{\text{AIK}(m)^{-1}}\{\{(PCR(s), e'')\}\}) \\ & \quad \vee (\neg \text{Contains}(e', PCR(s)))))) \end{aligned} \quad (\text{B.38})$$

Proving the above is rather straightforward. For $Q \in IS(TPM_{SRTM}(m))$, we prove a stronger invariant:

$$\begin{aligned} [Q]_I^{t_b, t_e} \quad & (\forall l, e, t. (t \in (t_b, t_e]) \supset \neg \text{Write}(I, l, e) @ t) \\ & \wedge (\forall t', e'. ((t' \in (t_b, t_e]) \wedge \text{Send}(I, e') @ t') \\ & \quad \supset (\exists e'', t_R. (t_R < t') \wedge (\text{Read}(I, m.pcr.s, e'') @ t_R) \wedge e' = \text{SIG}_{\text{AIK}(m)^{-1}}\{\{(PCR(s), e'')\}\})) \end{aligned} \quad (\text{B.39})$$

For $Q \neq TPM_{SRTM}(m)$, this follows just from rule (Seq) and axioms (Act2) and (ActN1). For $Q = TPM_{SRTM}(m)$, we use axioms (Sign) and (Act1) in addition. For $Q \in IS(TPM_{DRTM}(m))$, we prove that,

$$\begin{aligned} [Q]_I^{t_b, t_e} \quad & (\forall l, e, t. (t \in (t_b, t_e]) \supset \neg \text{Write}(I, l, e) @ t) \\ & \wedge (\forall t', e'. ((t' \in (t_b, t_e]) \wedge \text{Send}(I, e') @ t') \\ & \quad \supset (\neg \text{Contains}(e', PCR(s)))) \end{aligned} \quad (\text{B.40})$$

This follows immediately using basic properties of the predicate Contains. Having established the invariant in (B.38), we use the rule (Honesty) to deduce that

$$\begin{aligned} \text{Honest}(\hat{I}, \{TPM_{SRTM}(m), TPM_{DRTM}(m)\}) \supset & \forall t_e. (\forall l, e, t. (t \in (-\infty, t_e]) \supset \neg \text{Write}(I, l, e) @ t) \\ & \wedge (\forall t', e'. ((t' \in (-\infty, t_e]) \wedge \text{Send}(I, e') @ t') \\ & \quad \supset (\exists e'', t_R. (t_R < t') \\ & \quad \quad \wedge (\text{Read}(I, m.pcr.s, e'') @ t_R) \\ & \quad \quad \wedge e' = \text{SIG}_{\text{AIK}(m)^{-1}}\{\{(PCR(s), e'')\}\}) \\ & \quad \quad \vee (\neg \text{Contains}(e', PCR(s)))))) \end{aligned} \quad (\text{B.41})$$

Setting $I = \text{AIK}(m)$, and using the fact that $\text{Honest}(\hat{\text{AIK}}(m), \{TPM_{SRTM}(m), TPM_{DRTM}(m)\})$ (from Γ_{SRTM}), we get

$$\begin{aligned} \forall t_e. (\forall l, e, t. (t \in (-\infty, t_e]) \supset & \neg \text{Write}(\text{AIK}(m), l, e) @ t) \\ & \wedge (\forall t', e'. ((t' \in (-\infty, t_e]) \wedge \text{Send}(\text{AIK}(m), e') @ t') \\ & \quad \supset (\exists e'', t_R. (t_R < t') \wedge (\text{Read}(\text{AIK}(m), m.pcr.s, e'') @ t_R) \wedge e' = \text{SIG}_{\text{AIK}(m)^{-1}}\{\{(PCR(s), e'')\}\}) \\ & \quad \vee (\neg \text{Contains}(e', PCR(s)))))) \end{aligned} \quad (\text{B.42})$$

Next, we instantiate $l = m.pcr.s$, $e = e$, $t = t_S$, $t' = t_S$, $e' = e$ to get the parametric formula

$$\begin{aligned} (t_S \leq t_e) \supset & \neg \text{Write}(\text{AIK}(m), m.pcr.s, e) @ t_S \\ & \wedge ((t_S \leq t_e) \wedge \text{Send}(\text{AIK}(m), e) @ t_S) \\ & \quad \supset (\exists e'', t_R. (t_R < t_S) \wedge (\text{Read}(\text{AIK}(m), m.pcr.s, e'') @ t_R) \wedge e = \text{SIG}_{\text{AIK}(m)^{-1}}\{\{(PCR(s), e'')\}\}) \\ & \quad \vee (\neg \text{Contains}(e, PCR(s)))) \end{aligned} \quad (\text{B.43})$$

This implies the weaker formula:

$$\begin{aligned}
& (t_S < t_e) \supset \neg \text{Write}(\text{AIK}(m), m.pcr.s, e) @ t_S \\
& \wedge ((t_S < t_e) \wedge \text{Send}(\text{AIK}(m), e) @ t_S) \\
& \quad \supset (\exists e'', t_R. (t_R < t_S) \wedge (\text{Read}(\text{AIK}(m), m.pcr.s, e'') @ t_R) \wedge e = \text{SIG}_{\text{AIK}(m)^{-1}}\{(\text{PCR}(s), e'')\}) \\
& \quad \quad \vee (\neg \text{Contains}(e, \text{PCR}(s)))
\end{aligned} \tag{B.44}$$

Since t_S and e are parameters, we can choose them to be the same as those in the existential quantifier in B.37. Thus combining the two formulas and simplifying, we get,

$$\begin{aligned}
[\text{Verifier}(m)]_V^{t_b, t_e} & \exists t_S, e, I. (t_S < t_e) \wedge \hat{I} = \text{AIK}\hat{K}(m) \\
& \wedge \text{Contains}(e, \text{SIG}_{\text{AIK}(m)^{-1}}\{(\text{PCR}(s), \text{seq}(\text{sinit}, \text{BL}(m), \text{OS}(m), \text{APP}(m)))\}) \\
& \wedge ((\exists e'', t_R. (t_R < t_S) \wedge (\text{Read}(I, m.pcr.s, e'') @ t_R) \wedge e = \text{SIG}_{\text{AIK}(m)^{-1}}\{(\text{PCR}(s), e'')\}) \\
& \quad \vee (\neg \text{Contains}(e, \text{PCR}(s))))
\end{aligned} \tag{B.45}$$

From the predicate $\text{Contains}(e, \text{SIG}_{\text{AIK}(m)^{-1}}\{(\text{PCR}(s), \text{seq}(\text{sinit}, \text{BL}(m), \text{OS}(m), \text{APP}(m)))\})$, we deduce $\text{Contains}(e, \text{PCR}(s))$. Hence we may simplify the above equation to get,

$$\begin{aligned}
[\text{Verifier}(m)]_V^{t_b, t_e} & \exists t_S, e, I. (t_S < t_e) \wedge \hat{I} = \text{AIK}\hat{K}(m) \\
& \wedge \text{Contains}(e, \text{SIG}_{\text{AIK}(m)^{-1}}\{(\text{PCR}(s), \text{seq}(\text{sinit}, \text{BL}(m), \text{OS}(m), \text{APP}(m)))\}) \\
& \wedge (\exists e'', t_R. (t_R < t_S) \wedge (\text{Read}(I, m.pcr.s, e'') @ t_R) \wedge e = \text{SIG}_{\text{AIK}(m)^{-1}}\{(\text{PCR}(s), e'')\})
\end{aligned} \tag{B.46}$$

Reorganizing and simplifying, we get,

$$\begin{aligned}
[\text{Verifier}(m)]_V^{t_b, t_e} & \exists t_R, e, e'', I. (t_R < t_e) \wedge \hat{I} = \text{AIK}\hat{K}(m) \\
& \wedge \text{Contains}(e, \text{SIG}_{\text{AIK}(m)^{-1}}\{(\text{PCR}(s), \text{seq}(\text{sinit}, \text{BL}(m), \text{OS}(m), \text{APP}(m)))\}) \\
& \wedge (\text{Read}(I, m.pcr.s, e'') @ t_R) \wedge e = \text{SIG}_{\text{AIK}(m)^{-1}}\{(\text{PCR}(s), e'')\}
\end{aligned} \tag{B.47}$$

Simplifying further using axiom (Eq), we get,

$$\begin{aligned}
[\text{Verifier}(m)]_V^{t_b, t_e} & \exists t_R, e'', I. (t_R < t_e) \wedge \hat{I} = \text{AIK}\hat{K}(m) \\
& \wedge \text{Contains}(\text{SIG}_{\text{AIK}(m)^{-1}}\{(\text{PCR}(s), e'')\}, \text{SIG}_{\text{AIK}(m)^{-1}}\{(\text{PCR}(s), \text{seq}(\text{sinit}, \text{BL}(m), \text{OS}(m), \text{APP}(m)))\}) \\
& \wedge (\text{Read}(I, m.pcr.s, e'') @ t_R)
\end{aligned} \tag{B.48}$$

Using axiom (READ), we obtain,

$$\begin{aligned}
[\text{Verifier}(m)]_V^{t_b, t_e} & \exists t_R, e''. (t_R < t_e) \\
& \wedge \text{Contains}(\text{SIG}_{\text{AIK}(m)^{-1}}\{(\text{PCR}(s), e'')\}, \text{SIG}_{\text{AIK}(m)^{-1}}\{(\text{PCR}(s), \text{seq}(\text{sinit}, \text{BL}(m), \text{OS}(m), \text{APP}(m)))\}) \\
& \wedge (\text{Mem}(m.pcr.s, e'') @ t_R)
\end{aligned} \tag{B.49}$$

Now using basic properties of containment we get (in the context of the above existential quantifiers),

$$\begin{aligned}
& (e'' = \text{seq}(\text{sinit}, \text{BL}(m), \text{OS}(m), \text{APP}(m))) \\
& \vee \text{Contains}(e'', \text{SIG}_{\text{AIK}(m)^{-1}}\{(\text{PCR}(s), \text{seq}(\text{sinit}, \text{BL}(m), \text{OS}(m), \text{APP}(m)))\})
\end{aligned} \tag{B.50}$$

The latter disjunct together with axiom (PCRC) and the fact $\text{Mem}(m.pcr.s, e'') @ t_R$ (equation B.49) gives a contradiction. Hence the former disjunct must hold, i.e., $e'' = \text{seq}(\text{sinit}, \text{BL}(m), \text{OS}(m), \text{APP}(m))$. Combining this with B.49 using axiom (Eq), we get,

$$\begin{aligned}
[\text{Verifier}(m)]_V^{t_b, t_e} & \exists t_R. (t_R < t_e) \\
& \wedge (\text{Mem}(m.pcr.s, \text{seq}(\text{sinit}, \text{BL}(m), \text{OS}(m), \text{APP}(m))) @ t_R)
\end{aligned} \tag{B.51}$$

This is what we set out to prove, except that t has been α -renamed to t_R . □

C Attestation using a Dynamic Root of Trust for Measurement (DRTM)

The protocol is shown in Figure 4. We are trying to prove the following property for the protocol:

$$\begin{aligned}
J_{DRTM} = [Verifier(m)]_V^{t_b, t_e} \quad & \exists J, t_X, t_E, t_N, t_L, t_C, n. \\
& \wedge (t_L < t_C < t_E < t_X < t_e) \wedge (t_b < t_N < t_E) \\
& \wedge (\text{New}(V, n) @ t_N) \\
& \wedge (\text{LateLaunch}(m, J) @ t_L) \\
& \wedge (\neg \text{LateLaunch}(m) \text{ on } (t_L, t_X]) \\
& \wedge (\neg \text{Reset}(m) \text{ on } (t_L, t_X]) \\
& \wedge (\text{Jump}(J, P(m)) @ t_C) \\
& \wedge (\neg \text{Jump}(J) \text{ on } (t_L, t_C)) \\
& \wedge (\text{Extend}(J, m.dpcr.k, EOL) @ t_X) \\
& \wedge (\text{Eval}(J, f) @ t_E) \\
& \wedge (\neg \text{Eval}(J, f) \text{ on } (t_C, t_E)) \\
& \wedge (\neg \text{Eval}(J, f) \text{ on } (t_E, t_X)) \\
& \wedge (\text{IsLocked}(m.dpcr.k, J) \text{ on } (t_L, t_X])
\end{aligned}$$

In this proof we make the following assumptions:

$$\Gamma_{DRTM} = \text{Honest}(AI\hat{K}(m), \{TPM_{SRTM}(m), TPM_{DRTM}(m)\}), \hat{V} \neq AI\hat{K}(m)$$

Theorem 8. $\Gamma_{DRTM} \vdash J_{DRTM}$

Proof. We begin by analyzing the code of $Verifier(m)$ using the rule (Seq) and the axiom (Act1). In two steps we deduce,

$$\begin{aligned}
[Verifier(m)]_V^{t_b, t_e} \quad & \exists t_N, t_m, n. (t_b < t_N < t_m \leq t_e) \wedge (\text{New}(V, n) @ t_N) \\
& \wedge (\text{Verify}(V, (dPCR(k), seq(dinit, P(m), n, EOL)), AIK(m)) @ t_m)
\end{aligned} \tag{C.1}$$

Using axiom (VER), and the assumptions $\text{Honest}(AI\hat{K}(m), \{TPM_{SRTM}(m), TPM_{DRTM}(m)\})$ and $\hat{V} \neq AI\hat{K}(m)$ (from Γ_{DRTM}) together with the above formula, we deduce:

$$\begin{aligned}
[Verifier(m)]_V^{t_b, t_e} \quad & \exists t_N, t_m, n. (t_b < t_N < t_m \leq t_e) \wedge (\text{New}(V, n) @ t_N) \\
& \wedge \text{Verify}(V, seq(dinit, P(m), n, EOL), AIK(m)) @ t_m \\
& \wedge \exists I', t', e'. (t' < t_m) \wedge (\hat{I}' = AI\hat{K}(m)) \\
& \wedge \text{Contains}(e', SIG_{AIK(m)^{-1}}\{(dPCR(k), seq(dinit, P(m), n, EOL))\}) \\
& \wedge ((\text{Send}(I', e') @ t') \vee \exists l. (\text{Write}(I', l, e') @ t'))
\end{aligned} \tag{C.2}$$

Simplifying slightly, and α -renaming t' to t_S (S for 'send'), I' to I , and e' to e , we obtain,

$$\begin{aligned}
[Verifier(m)]_V^{t_b, t_e} \quad & \exists t_N, t_S, e, I, n. (t_b < t_N < t_e) \wedge (\text{New}(V, n) @ t_N) \\
& \wedge (t_S < t_e) \wedge \hat{I} = AI\hat{K}(m) \\
& \wedge \text{Contains}(e, SIG_{AIK(m)^{-1}}\{(dPCR(k), seq(dinit, P(m), n, EOL))\}) \\
& \wedge ((\text{Send}(I, e) @ t_S) \vee \exists l. (\text{Write}(I, l, e) @ t_S))
\end{aligned} \tag{C.3}$$

Next we prove an invariant about the programs executed by the TPM, namely $TPM_{SRTM}(m)$ and $TPM_{DRTM}(m)$ (as we did for SRTM). Specifically, we show that for each $Q \in IS(TPM_{SRTM}(m), TPM_{DRTM}(m))$, it is the case that

$$\begin{aligned}
[Q]_I^{t_b, t_e} \quad & (\forall l, e, t. (t \in (t_b, t_e]) \supset \neg \text{Write}(I, l, e) @ t) \\
& \wedge (\forall t', e'. ((t' \in (t_b, t_e]) \wedge \text{Send}(I, e') @ t') \\
& \quad \supset (\exists e'', t_R. (t_R < t') \wedge (\text{Read}(I, m.dpcr.k, e'') @ t_R) \wedge e' = SIG_{AIK(m)^{-1}}\{(dPCR(k), e'')\}) \\
& \quad \vee (\neg \text{Contains}(e', dPCR(k))))
\end{aligned} \tag{C.4}$$

We omit this proof since it is similar to invariance proofs described earlier. Next we use equations C.3 and C.4, and follow the proof of Theorem 7 (steps B.38 to the end), replacing $m.pcr.s$ by $m.dpcr.k$, $PCR(s)$ by $dPCR(k)$, and $seq(sinit, BL(m), OS(m), APP(m))$ by $seq(dinit, P(m), n, EOL)$ to deduce that

$$[Verif\ ier(m)]_V^{t_b, t_e} \quad \exists t_N, t_R, n. (t_b < t_N < t_e) \wedge (\text{New}(V, n) @ t_N) \wedge (t_R < t_e) \\ \wedge (\text{Mem}(m.dpcr.k, seq(dinit, P(m), n, EOL)) @ t_R) \quad (\text{C.5})$$

Next, we use axioms (dPCR1) and (dPCR2) repeatedly to obtain the following:

$$[Verif\ ier(m)]_V^{t_b, t_e} \quad \exists J, t_N, t_R, t_3, t_2, t_1, t_L, n. (t_b < t_N < t_e) \wedge (\text{New}(V, n) @ t_N) \\ \wedge (t_L \leq t_1 < t_2 < t_3 < t_R < t_e) \\ \wedge (\text{Mem}(m.dpcr.k, seq(dinit, P(m), n, EOL)) @ t_R) \\ \wedge (\text{Mem}(m.dpcr.k, seq(dinit, P(m), n)) @ t_3) \\ \wedge (\text{Mem}(m.dpcr.k, seq(dinit, P(m))) @ t_2) \\ \wedge (\text{Mem}(m.dpcr.k, dinit)) @ t_1 \\ \wedge (\text{LateLaunch}(m, J) @ t_L) \\ \wedge (\neg \text{LateLaunch}(m) \text{ on } (t_L, t_R]) \\ \wedge (\neg \text{Reset}(m) \text{ on } (t_L, t_R]) \quad (\text{C.6})$$

We know from the above that there is a thread J such that $\text{LateLaunch}(m, J) @ t_L$. We would like to apply the (LateLaunch) rule to reason further. In order to do that, we must prove an invariant of the program $LL(m)$. The specific invariant we prove is that for each $Q \in IS(LL(m))$, it is the case that,

$$[Q]_J^{t_b, t_e} \quad \forall t, P, x. ((\text{LateLaunch}(m, J) @ t_b) \wedge (\neg \text{LateLaunch}(m) \text{ on } (t_b, t_e]) \\ \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e]) \wedge (t_b < t \leq t_e) \wedge (\text{Mem}(m.dpcr.k, seq(dinit, P, x)) @ t)) \\ \supset \exists t'. ((t_b < t' < t) \wedge (\text{Jump}(J, P) @ t') \\ \wedge (\neg \text{Jump}(J) \text{ on } (t_b, t')) \\ \wedge (\text{Mem}(m.dpcr.k, seq(dinit, P)) @ t') \\ \wedge (\text{IsLocked}(m.dpcr.k, J) \text{ on } (t_b, t')))) \quad (\text{C.7})$$

The proof of the invariant in C.7 is similar to previous proofs, hence we omit the details. Briefly, the proof uses the axioms (Mem=) and (Jump1) to establish that a jump happened, and (MemLL), (LockLL), (LockIdP), and (MemIdP) to infer the properties about memory and locks. Next, using the (LateLaunch) rule we deduce,

$$\text{LateLaunch}(m, J) @ t_b \supset \forall t_e. (t_e > t_b) \supset \forall t, P, x. \\ ((\text{LateLaunch}(m, J) @ t_b) \wedge (\neg \text{LateLaunch}(m) \text{ on } (t_b, t_e]) \\ \wedge (\neg \text{Reset}(m) \text{ on } (t_b, t_e]) \wedge (t_b < t \leq t_e) \wedge (\text{Mem}(m.dpcr.k, seq(dinit, P, x)) @ t)) \\ \supset \exists t'. ((t_b < t' < t) \wedge (\text{Jump}(J, P) @ t') \\ \wedge (\neg \text{Jump}(J) \text{ on } (t_b, t')) \\ \wedge (\text{Mem}(m.dpcr.k, seq(dinit, P)) @ t') \\ \wedge (\text{IsLocked}(m.dpcr.k, J) \text{ on } (t_b, t')))) \quad (\text{C.8})$$

Now we instantiate the above formula with values from C.10. We choose $t_b = t_L$, $t_e = t_R$, $t = t_3$, $P = P(m)$, $x = n$, and α -rename t' to t_C to obtain,

$$\text{LateLaunch}(m, J) @ t_L \supset (t_R > t_L) \supset ((\text{LateLaunch}(m, J) @ t_L) \wedge (\neg \text{LateLaunch}(m) \text{ on } (t_L, t_R]) \\ \wedge (\neg \text{Reset}(m) \text{ on } (t_L, t_R]) \wedge (t_L < t_3 \leq t_R) \wedge (\text{Mem}(m.dpcr.k, seq(dinit, P, n)) @ t_3)) \\ \supset \exists t_C. ((t_L < t_C < t_3) \wedge (\text{Jump}(J, P(m)) @ t_C) \\ \wedge (\neg \text{Jump}(J) \text{ on } (t_L, t_C)) \\ \wedge (\text{Mem}(m.dpcr.k, seq(dinit, P(m))) @ t_C) \\ \wedge (\text{IsLocked}(m.dpcr.k, J) \text{ on } (t_L, t_C])) \quad (\text{C.9})$$

Combining with C.10 and simplifying, we obtain,

$$\begin{aligned}
[Verifier(m)]_V^{t_b, t_e} \quad & \exists J, t_N, t_R, t_3, t_2, t_1, t_L, t_C, n. (t_b < t_N < t_e) \wedge (\text{New}(V, n) @ t_N) \\
& \wedge (t_L \leq t_1 < t_2 < t_3 < t_R < t_e) \wedge (t_L < t_C < t_3) \\
& \wedge (\text{Mem}(m.dpcr.k, seq(dinit, P(m), n, EOL)) @ t_R) \\
& \wedge (\text{Mem}(m.dpcr.k, seq(dinit, P(m), n)) @ t_3) \\
& \wedge (\text{Mem}(m.dpcr.k, seq(dinit, P(m))) @ t_2) \\
& \wedge (\text{Mem}(m.dpcr.k, dinit)) @ t_1) \\
& \wedge (\text{LateLaunch}(m, J) @ t_L) \\
& \wedge (\neg \text{LateLaunch}(m) \text{ on } (t_L, t_R]) \\
& \wedge (\neg \text{Reset}(m) \text{ on } (t_L, t_R]) \\
& \wedge (\text{Jump}(J, P(m)) @ t_C) \\
& \wedge (\neg \text{Jump}(J) \text{ on } (t_L, t_C)) \\
& \wedge (\text{Mem}(m.dpcr.k, seq(dinit, P(m))) @ t_C) \\
& \wedge (\text{IsLocked}(m.dpcr.k, J) \text{ on } (t_L, t_C])
\end{aligned} \tag{C.10}$$

Next we want to use the (Jump) rule to reason from the assumption $\text{Jump}(J, P(m)) @ t_C$. In this case, we want to prove an invariant about the program $P(m)$. Specifically, we want to show that for each $Q \in IS(P(m))$, it is the case that,

$$\begin{aligned}
[Q]_J^{t_b, t_e} \quad & \forall t, x. \\
& ((\neg \text{Reset}(m) \text{ on } (t_b, t_e]) \\
& \wedge (\neg \text{LateLaunch}(m) \text{ on } (t_b, t_e]) \\
& \wedge (\text{Mem}(m.dpcr.k, seq(dinit, P(m))) @ t_b) \\
& \wedge (\text{IsLocked}(m.dpcr.k, J) @ t_b) \\
& \wedge (t_b < t \leq t_e) \wedge (\text{Mem}(m.dpcr.k, seq(dinit, P(m), x, EOL)) @ t)) \\
& \supset \exists t_n, t_E, t_X. ((t_b < t_n < t_E < t_X < t) \wedge (\text{Eval}(J, f) @ t_E) \\
& \wedge (\text{Extend}(J, m.dpcr.k, x) @ t_n) \\
& \wedge (\text{Extend}(J, m.dpcr.k, EOL) @ t_X) \\
& \wedge (\neg \text{Eval}(J, f) \text{ on } (t_b, t_E)) \\
& \wedge (\neg \text{Eval}(J, f) \text{ on } (t_E, t_X)) \\
& \wedge (\text{IsLocked}(m.dpcr.k, J) \text{ on } (t_b, t_X]))
\end{aligned} \tag{C.11}$$

Again, we omit the details of this invariant's proof. Having established this invariant, we use the rule (Jump) to deduce that,

$$\begin{aligned}
\text{Jump}(J, P(m)) @ t_b \supset \quad & \forall t_e. (t_e > t_b) \supset \forall t, x. \\
& ((\neg \text{Reset}(m) \text{ on } (t_b, t_e]) \\
& \wedge (\neg \text{LateLaunch}(m) \text{ on } (t_b, t_e]) \\
& \wedge (\text{Mem}(m.dpcr.k, seq(dinit, P(m))) @ t_b) \\
& \wedge (\text{IsLocked}(m.dpcr.k, J) @ t_b) \\
& \wedge (t_b < t \leq t_e) \wedge (\text{Mem}(m.dpcr.k, seq(dinit, P(m), x, EOL)) @ t)) \\
& \supset \exists t_n, t_E, t_X. ((t_b < t_n < t_E < t_X < t) \wedge (\text{Eval}(J, f) @ t_E) \\
& \wedge (\text{Extend}(J, m.dpcr.k, x) @ t_n) \\
& \wedge (\text{Extend}(J, m.dpcr.k, EOL) @ t_X) \\
& \wedge (\neg \text{Eval}(J, f) \text{ on } (t_b, t_E)) \\
& \wedge (\neg \text{Eval}(J, f) \text{ on } (t_E, t_X)) \\
& \wedge (\text{IsLocked}(m.dpcr.k, J) \text{ on } (t_b, t_X]))
\end{aligned} \tag{C.12}$$

We now instantiate this formula with values from C.10. We choose $t_b = t_C$, $t_e = t_R$, $t = t_R$, and $x = n$.

$$\begin{aligned}
\text{Jump}(J, P(m)) @ t_C \supset & (t_R > t_C) \supset \\
& ((\neg \text{Reset}(m) \text{ on } (t_C, t_R]) \\
& \wedge (\neg \text{LateLaunch}(m) \text{ on } (t_C, t_R]) \\
& \wedge (\text{Mem}(m.dpcr.k, \text{seq}(\text{dinit}, P(m))) @ t_C) \\
& \wedge (\text{IsLocked}(m.dpcr.k, J) @ t_C) \\
& \wedge (t_C < t_R \leq t_R) \wedge (\text{Mem}(m.dpcr.k, \text{seq}(\text{dinit}, P(m), n, \text{EOL})) @ t_R)) \\
& \supset \exists t_n, t_E, t_X. ((t_C < t_n < t_E < t_X < t_R) \wedge (\text{Eval}(J, f) @ t_E) \\
& \wedge (\text{Extend}(J, m.dpcr.k, n) @ t_n) \\
& \wedge (\text{Extend}(J, m.dpcr.k, \text{EOL}) @ t_X) \\
& \wedge (\neg \text{Eval}(J, f) \text{ on } (t_C, t_E)) \\
& \wedge (\neg \text{Eval}(J, f) \text{ on } (t_E, t_X)) \\
& \wedge (\text{IsLocked}(m.dpcr.k, J) \text{ on } (t_C, t_X]))
\end{aligned} \tag{C.13}$$

Combining with C.10 and simplifying we get,

$$\begin{aligned}
[\text{Verifier}(m)]_V^{t_b, t_e} \exists J, t_n, t_X, t_E, t_N, t_R, t_L, t_C, n. & (t_b < t_N < t_e) \wedge (\text{New}(V, n) @ t_N) \\
& \wedge (t_L \leq t_R < t_e) \wedge (t_L < t_C < t_n < t_E < t_X < t_R) \\
& \wedge (\text{LateLaunch}(m, J) @ t_L) \\
& \wedge (\neg \text{LateLaunch}(m) \text{ on } (t_L, t_R]) \\
& \wedge (\neg \text{Reset}(m) \text{ on } (t_L, t_R]) \\
& \wedge (\text{Jump}(J, P(m)) @ t_C) \\
& \wedge (\neg \text{Jump}(J) \text{ on } (t_L, t_C)) \\
& \wedge (\text{Mem}(m.dpcr.k, \text{seq}(\text{dinit}, P(m))) @ t_C) \\
& \wedge (\text{IsLocked}(m.dpcr.k, J) @ \text{on}(t_L, t_C]) \\
& \wedge (\text{Extend}(J, m.dpcr.k, n) @ t_n) \\
& \wedge (\text{Extend}(J, m.dpcr.k, \text{EOL}) @ t_X) \\
& \wedge (\text{Eval}(J, f) @ t_E) \\
& \wedge (\neg \text{Eval}(J, f) \text{ on } (t_C, t_E)) \\
& \wedge (\neg \text{Eval}(J, f) \text{ on } (t_E, t_X)) \\
& \wedge (\text{IsLocked}(m.dpcr.k, J) \text{ on } (t_C, t_X])
\end{aligned} \tag{C.14}$$

Next, we use the facts $\text{New}(V, n) @ t_N$ and $\text{Extend}(J, m.dpcr.k) @ t_n$ together with the axiom (NEW) to deduce that $t_N < t_n$. Simplifying using this fact we obtain,

$$\begin{aligned}
[\text{Verifier}(m)]_V^{t_b, t_e} \exists J, t_X, t_E, t_N, t_L, t_C, n. & \\
& \wedge (t_L < t_C < t_E < t_X < t_e) \wedge (t_b < t_N < t_E) \\
& \wedge (\text{New}(V, n) @ t_N) \\
& \wedge (\text{LateLaunch}(m, J) @ t_L) \\
& \wedge (\neg \text{LateLaunch}(m) \text{ on } (t_L, t_X]) \\
& \wedge (\neg \text{Reset}(m) \text{ on } (t_L, t_X]) \\
& \wedge (\text{Jump}(J, P(m)) @ t_C) \\
& \wedge (\neg \text{Jump}(J) \text{ on } (t_L, t_C)) \\
& \wedge (\text{Extend}(J, m.dpcr.k, \text{EOL}) @ t_X) \\
& \wedge (\text{Eval}(J, f) @ t_E) \\
& \wedge (\neg \text{Eval}(J, f) \text{ on } (t_C, t_E)) \\
& \wedge (\neg \text{Eval}(J, f) \text{ on } (t_E, t_X)) \\
& \wedge (\text{IsLocked}(m.dpcr.k, J) \text{ on } (t_L, t_X])
\end{aligned} \tag{C.15}$$

This is what we set out to prove. □