REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
05/25/2021	Technical Report	

4. TITLE AND SUBTITLE
Automated Trust Analysis for Layered Attestations

5a. CONTRACT NUMBER
W56KGU-18-D-0004

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Paul D. Rowe; John D. Ramsdell; Ian D. Kretz

5d. PROJECT NUMBER
0721N6BZMA

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
The MITRE Corporation
202 Burlington Road
Bedford, MA 01730

8. PERFORMING ORGANIZATION REPORT NUMBER
PRS-21-1039

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
United States Army (ACC-APG)

10. SPONSOR/MONITOR'S ACRONYM(S)
ACC-APG

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT
In distributed systems, trust decisions are often based on layered attestations in which evidence is gathered about the integrity of subcomponents, leveraging hierarchical dependencies among the subcomponents to bolster the trustworthiness of evidence. Copland is a domain-specific language for specifying complex layered attestations. How phrases are composed bears directly on the trustworthiness of the evidence they produced, and complex phrases become quite difficult to analyze by hand. We introduce an automated method for analyzing executions of attestations specified by Copland phrases in an adversarial setting. We develop a general theory of executions with adversarial corruption and repair events.

15. SUBJECT TERMS
Computer-Communication Networks (General); Chase geometric logic; layered attestation; Copland language; saturated queries

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE		33	Susan Carpenito
					19b. TELEPHONE NUMBER *(Include area code)* 781-271-7646

Approved for Public Release; Distribution Unlimited.
Public Release Case Number 21-1039

Automated Trust Analysis for Layered Attestations

Paul D. Rowe John D. Ramsdell Ian D. Kretz

The MITRE Corporation

Abstract

In distributed systems, trust decisions are often based on layered attestations in which evidence is gathered about the integrity of subcomponents, leveraging hierarchical dependencies among the subcomponents to bolster the trustworthiness of evidence. Copland is a domain-specific language for specifying complex layered attestations. How phrases are composed bears directly on the trustworthiness of the evidence they produce, and complex phrases become quite difficult to analyze by hand. We introduce an automated method for analyzing executions of attestations specified by Copland phrases in an adversarial setting. We develop a general theory of executions with adversarial corruption and repair events. Our approach is to enrich the Copland semantics according to this theory. Using the model finder Chase, we characterize all executions consistent with a set of initial assumptions. From this set of models, an analyst can discover all ways an active adversary can corrupt subcomponents without being detected by the attestation. These efforts afford trust policymakers the ability to compare attestations expressed as Copland phrases against trust policy in a way that encompasses both static and runtime concerns.

1 Introduction

Network-based communication among computing devices increasingly relies on a notion of trust to inform the nature of their interactions. Remote attestation is a technique for allowing one entity (the target of attestation) to provide evidence of its trustworthiness to a peer (the appraiser of an attestation). It consists, in part, of processes on the target system that gather evidence by performing integrity measurements of various components of the target system. The evidence generated about these components is bundled together and transmitted to a remote peer who can appraise the evidence. The result of the appraisal can form an input to a trust decision that will govern how the network interaction will proceed. For example, a remote attestation may serve to provide a corporate VPN gateway with sufficient evidence that a machine wishing to join the network is free of malware and that it conforms to the corporate configuration before access is granted.

This technical data deliverable was developed using contract funds under Basic Contract No. W56KGU-18-D-0004.

Remote attestation faces the following interesting dilemma. If the appraiser is distrustful of the target to begin with, why should the appraiser be any more trustful of the target to faithfully gather and report accurate evidence? The answer typically lies in the layered dependencies among components of the target system. A kernel process is likely to be more trustworthy than a userspace process, if only because the interface to an operating system's kernel provides stronger protections to resist malicious corruption. It is therefore inherently more trustworthy to measure the integrity of a userspace process from the kernel than it is to measure it from another userspace process which may be just as vulnerable to corruption as the target process itself.

To illustrate this point more concretely, consider a simple attestation scenario. Alice is logged into her bank's website and attempts to initiate a high-value transfer of money out of her account. The nature of this transaction prompts the bank to ask Alice to confirm her credentials, perhaps via two-factor authentication. But it could also prompt the bank to initiate an attestation of Alice's system so that it can trust there is no malware present (perhaps in the form of a malicious browser extension) that might have hijacked the session to initiate transactions on Alice's behalf. Of course Alice herself would be willing to reveal some extra information about the current state of her system due to the sensitive nature of the transaction. One approach the bank might take is to develop its own browser extension that can list the other extensions present and compare against a whitelist of approved extensions, refusing the transaction if it sees an extension it does not know or trust. However, an adversary able to install a malicious extension may be just as capable of disabling or corrupting the bank's extension. It would be more trustworthy to inspect the browser from outside the browser itself with some special-purpose application.

This may prompt some readers to become skeptical of the trustworthiness of this special-purpose application. For similar reasons, most of the research on remote attestation has focused on how to build trust from the ground up starting in hardware. Hardware is considered significantly harder to compromise than software, so it can provide a strong root of trust for performing and storing integrity measurements. The development of the Trusted Platform Module [9] and Intel's Software Guard Extensions (SGX) [10] represent prominent products in this area. However, hardware is also less flexible than software. This has led to hardware-based attestation solutions that focus on primarily on boot-time integrity measurements at the expense of runtime integrity.

Virtualization technologies such as Xen [1] or Microsoft's Hyper-V [21] are good examples of approaches that can provide strong support for software-based runtime integrity measurement. They offer a place for measurers to stand that is better protected than the environment of the target they are measuring. They contribute to a layered architecture in which more privileged functions benefit from the protection offered by lower layers of the system. Such architectures can combine the benefits of hardware roots of trust with the flexibility afforded by software mechanisms. This suggests an approach to the bank scenario described above in which trust is built from the ground up, starting with hardware and moving up through firmware, hypervisor, and kernel to support the top-level

inspection of the browser extensions.

Such an attestation would offer strong evidence of the integrity of the browser extensions, but it may appear to some as a lot of work to generate integrity evidence for a simple property. On the other hand, the assurance gained by pushing trust all the way down to hardware could be well worth it for a more high-stakes property. Consider that a keylogger on the client's machine could learn the client's password, after which an adversary could easily masquerade as the client in the future using a "clean" machine that could pass any attestation. Thus, a more complete attestation performed before every login attempt would be able to increase the chances of discovering keyloggers before they have an opportunity to snoop the password. So it is desirable to support the ability to adjust the level of trust required and tailor it to the situation.

Indeed, an attestation can be strengthened in layered architectures even without going all the way down to hardware. Previous work has established a model and a set of reasoning principles for the analysis of layered attestations focused on runtime measurements [18]. These principles give support to the notion that the trustworthiness of the evidence produced by a layered attestation depends on the order in which evidence is gathered on the target system. This is due to the layered dependencies among components on the target system. In particular, building up trust in components in a "bottom-up" manner is indeed generally more trustworthy than other orders. If an adversary is to avoid being detected, it must either corrupt deeper (and presumably better-protected) components of the target system or else corrupt components in opportune time windows during the attestation itself. This result is summarized by saying that bottom-up measurements force the adversary to perform corruptions that are either "recent or deep".

But it may not be necessary to force an adversary all the way down to the hardware level. In other words, the components we consider sufficiently "deep" will depend (among other things) on the context of the transaction. The primary targets of measurement for an attestation will also vary depending on the nature of the trust decision being supported by the attestation. What a bank wants to know about a system before approving a large transaction is likely different from what a corporate gateway wants to know about the same system when admitting it onto the corporate network. This suggests that the appraiser and the target will need a way to negotiate details of an attestation such as which components get measured and in what order, how deep the measurements go, and how the evidence is bundled for appraisal.

Copland [17] is a specification language designed to tailor specifications to different target device architectures and different contexts for trust decisions. Its formal semantics enables semantically clear negotiations between a target and a relying party about the details of the attestation to be performed. The flexibility allowed by Copland specifications is crucial for its ability to accommodate the full range of situations in which layered attestations might be performed. However, it underscores the importance of understanding the trust properties provided by alternative specifications. The "bottom-up" rule is insufficient in part because it does not say how deep is deep enough. An appraiser should be

able to determine what deeper attestations buy them so they can consider the trade-offs between trust and other factors. For instance, all else being equal, a speedier option may be preferred to one that takes more time to complete.

In this paper, we explicitly apply the reasoning principles developed in [18] to the trust analysis of Copland phrases as introduced in [17]. Since even moderately large Copland phrases can be prohibitively difficult to analyze by hand, we introduce a method for automating the reasoning principles. The basic goal of our analysis is to determine all the essentially different ways an adversary can corrupt a given subcomponent while avoiding detection by a layered attestation designed to attest to the integrity of the subcomponent. Our approach is to use Chase [16], a general-purpose model finder for geometric logic, to enumerate all the possible executions consistent with the Copland specification in which the adversary corrupts the target and avoids detection. This set of executions contains all the information needed to understand the conditions under which each subcomponent must be trusted.

More concretely, our novel contributions in this paper are as follows:

1. We develop a first-order theory of *saturated queries* that adapts the reasoning principles of [18] to the analysis problem of finding all models that violate the attestation goals. We identify a correctness criterion for our theory and prove that our theory meets that criterion (Thm. 14).

2. Since Chase is a model finder that works on the geometric fragment of first-order logic, we find a suitable axiomatization of our theory in special geometric form and prove the equivalence of the geometric theory to the more general first-order theory developed above (Thm. 18). This two-stage axiomatization ensures we accurately capture the reasoning principles laid out in [18] while allowing us to leverage Chase for automation.

3. We demonstrate the use of an end-to-end tool chain that compiles a Copland phrase into its event semantics [17] and then performs the trust analysis yielding a characterization of all possible ways for an adversary to defeat the attestation. From this characterization the analyst can determine the least amount of work an adversary must perform to defeat the attestation. The analyst can similarly compare two Copland phrases according to the work required to defeat each one. We also show how an analyst can alter assumptions about dependency relationships on the target and about adversary capabilities. This allows the analyst to tailor the analysis according their needs.

The remainder of the paper is structured as follows. In Section 2, we introduce a typical attestation scenario to introduce the context and some notation, and to motivate our goals. Section 3 provides a high-level overview of the syntax and semantics of Copland. Section 4 reviews the reasoning principles of [18] and adapts them into the first-order theory of saturated queries. In Section 5, we briefly introduce our Chase model finder and how it works. Section 6 provides the translation of the theory from Section 4 into special geometric form.

Section 7 walks the reader through several examples of using Chase to analyze the trust properties of Copland phrases under a variety of assumptions. After presenting some related work, we finally conclude.

2 Motivation

In this section, we use the client-bank scenario as a vehicle to illustrate the utility of Copland, a specification language for layered attestation. Via this discussion, we pose the central problem of this work: determining how an adversary can act to their own advantage among a set of partially ordered measurement events to avoid detection. In the process, we will examine Copland's syntax and semantics with respect to these sets of labeled events.

In the bank example, the client host must provide evidence to the bank that its browser is not equipped with any malicious extensions before a sensitive transaction is allowed to proceed. A very simple instance of this attestation may take place among the bank, a browser monitor bmon running at the client in userspace us, and a host-based antivirus suite av running at the client in kernelspace ks, with the ultimate target of the attestation being exts, the set of extensions in the client's browser. The bank would like the client to leverage either or both of av and bmon to generate trustworthy evidence about the benign or malign nature of exts. As we will see, the bank's choice greatly influences the trustworthiness of the evidence it receives.

Critical to this discussion is a characterization of the adversary. We envision an active adversary equipped with broad abilities to interfere with the proper function of logical components on the system, for instance bmon, av and the host's kernel, which provides clean runtime contexts for all processes on the system. This adversary may *corrupt* a component via an intervention that causes it to deviate from its intended design or function. By the same token, the adversary may decide to *repair* a corrupt component, that is, return it to its *regular* (uncorrupted) state.

The regular-corrupt concept has two important implications for measurement: a regular measurer will always accurately report the corruption state of its target, and a corrupt measurer will always report a regular state for its target.

However powerful it may be, this adversary is not omnipotent. We constrain it in the following ways:

- Tampering with evidence artifacts after they have been generated is presumed to always result in detection, so the adversary has no incentive to do this
- The adversary may not stop measurements from taking place (otherwise phrase semantics would be contradicted), but they may delay measurements to their own advantage
- All other things being equal, it is more difficult for the adversary to corrupt a component in a narrow timeframe than in a broad one

- All other things being equal, it is more difficult for the adversary to corrupt a component that is deeper in the system (subject to greater ambient protections) than a shallower one

The easiest corruption for the adversary to perform is of a shallow component without time constraints.

In the context of an attestation targeting a particular component, when the target is corrupt but evidence gathered about the target indicates the target is regular, we say the adversary has *avoided detection* at the component. A critical question is the following: assuming the adversary has avoided detection at a component of interest, what other components must also have been corrupted to make this possible? This is a central concern we must address to understand which topologies most and least constrain adversarial behavior, and therefore which we should choose.

Returning to the simple bank example, in light of av's more privileged place in the client's kernel, av's measurement may be all that is required to convince the bank of exts' integrity, depending on its own requirements.

By the same token, because bmon is collocated in us with exts, evidence gathered by bmon about exts is generally less trustworthy compared to that gathered by av. Moreover, as a userspace process, bmon is subject to the same protections afforded by the client's operating system as exts and fewer than av, which runs in the kernel. It might therefore be relatively easy for the adversary to corrupt bmon to their advantage. The bank might therefore also request that av perform a measurement of bmon and submit this with the latter's measurement of exts. Each measurement topology we might select has its own benefits and shortcomings.

The Copland phrase in Example 1 describes the single-measurement topology in which av measures exts directly.

Example 1.

$$*\mathsf{bank} : @_{\mathsf{ks}}\,[\mathsf{av\ us\ exts}]$$

A full explanation of Copland syntax and semantics follows in the next section. At a high level, this phrase says that the place bank initiates the attestation and receives the final evidence product. The requested attestation consists of a request that av in ks take a measurement of exts in us. So long as the bank trusts that av is beyond the reach of the adversary to corrupt, this measurement will accurately describe the corruption state of client browser extensions at the time they are measured. The downside of this phrase is that the browser monitor is underutilized.

The attestation specified in Example 2 encompasses two measurements taken in parallel, indicated by the $\sim$ operator joining them. Here, bank issues separate requests for measurement to av and bmon and forms the final evidence by joining the results according to a bundling strategy for composing parallel measurements. We envision less simultaneous measurements (although that is certainly possible) and more that any concrete precedence ordering between

the measurements is acceptable. This has the effect of leaving the ordering unconstrained.

Example 2.

$$*\mathsf{bank} : @_{\mathsf{ks}}\,[\mathsf{av\ us\ bmon}] -\sim- @_{\mathsf{us}}\,[\mathsf{bmon\ us\ exts}]$$

In one of these measurements, av from its position in ks measures bmon in us. In the other, bmon from its position in us measures exts, also in us. The idea is that av will collect evidence on bmon from its privileged position in the kernel, and bmon in turn will collect evidence on exts.

One way of looking at the Copland phrase in Example 2 is as a precise shorthand for a partially ordered set of labeled events. A subset of possible events includes requests for attestation and replies to those requests with collected evidence, measurement events, and branching and joining events. Every legal Copland phrase induces a well-defined partially ordered set of non-adversarial events. This relationship between phrases and events is covered extensively in the next section. Apart from branch/join and request/reply pairs, the phrase in Example 2 encompasses the two measurement events in which av measures bmon and bmon measures exts. The labels associated with these events are msp(ks.av, us.bmon) and msp(us.bmon, us.exts), respectively.

Unfortunately, applying no constraints to the precedence relationship between the two measurements opens this scheme up to relatively easy subversion by the adversary. Suppose a malicious extension has been installed sometime before the attestation begins. The adversary can delay av's measurement of bmon. During this interval, it may corrupt bmon and subsequently allow it to perform its measurement of exts. The adversary then repairs bmon before finally allowing av's measurement to take place. Via this procedure, the adversary has as much time as they need to corrupt bmon, a shallow component subject to relatively few protections. This is the ideal outcome for an adversary seeking to avoid detection at exts. That this attack pattern is viable is a consequence of the way this measurement topology fails to constrain event precedence sufficiently.

The Copland phrase in Example 3 is identical to that in Example 2 but for one critical distinction: av's measurement of bmon is now required to occur before the latter measures exts. This is indicated by the $<$ operator joining the two measurements. The bank will not request bmon measure exts before it has requested that av measure bmon and received acceptable evidence in response. Once again, bank will form the final evidence by joining the two results of measurement according to a bundling strategy for sequential measurements.

Example 3.

$$*\mathsf{bank} : @_{\mathsf{ks}}\,[\mathsf{av\ us\ bmon}] -<- @_{\mathsf{us}}\,[\mathsf{bmon\ us\ exts}]$$

Delaying av's measurement now provides no advantage to the adversary, as this will also delay bmon's measurement. In order to avoid detection at exts under this topology, the adversary has two choices: corrupt av and bmon

before the attestation begins or else corrupt bmon after it is measured by av but before it measures exts. The former requires the adversary to perform a *deep* corruption of av, while the latter requires the adversary to perform a *recent* corruption of bmon in a narrow timeframe during the attestation. Thus, the adversary is forced to pursue one of two difficult corruption strategies by this choice of topology. In fact, prior theoretical work has shown this to be a general feature of such *bottom-up* measurement topologies, in which better protected or more trusted components take measurements of less protected, less trusted ones before these perform their own measurements.[18]

The adversary's choices may be even further constrained by adding additional measurements. The Copland phrase in Example 4 ("trust but verify"), for instance, elaborates on that in Example 3 by adding a direct av measurement of exts sequentially with the bottom-up measurement chain from av through bmon to exts.

Example 4.

$$*\mathsf{bank} : (@_{\mathsf{ks}}\,[\mathsf{av\ us\ bmon} -<- @_{\mathsf{us}}\,[\mathsf{bmon\ us\ exts}]]) \\ -<- @_{\mathsf{ks}}\,[\mathsf{av\ us\ exts}]$$

This phrase removes the possibility of the adversary getting away with only a recent corruption of bmon. Rather, they will be forced to perform the deep corruption of av as well as the shallow corruption of bmon.

This discussion illustrates how deciding which components measure, which are measured, and the order in which measurements are taken can have striking implications for the trustworthiness of evidence. Relative to this formalism about measurement and corruption, it is easy enough to work out these implications by hand when there are only a handful of measurements to contend with. However, dependency structures and trust relationships in real-world systems can be extremely complex, and the sheer number of components one must be prepared to deal with in such systems would easily overwhelm any human analyst. The goal of this paper is to develop a method for automating this analysis and thereby scaling it up to meet the challenges of real-world attestation problems.

3 Copland Layered Attestation Language

Copland is a language for specifying layered attestations. It is used to specify the location and means by which a component is measured, the methods for combining the evidence collected from each measurement, and the constraints on the sequencing of measurement actions. It has a precise semantics based on partially ordered sets of events [17, Sec. 5]. Possible events include, among others, measurement actions to generate evidence, and hashing and signing to protect previously generated evidence.

Copland assumes that each action is performed at a location called a *place*. The origin of an attestation request p for phrase t is specified with the syntax

C	$\leftarrow$	$S\ P\ S$	Measurement (Probe Place Target)
	$\mid$	$@_P\ [\ C\ ]$	At place
	$\mid$	$\{\}$	Nullify
	$\mid$	$-$	Copy
	$\mid$	$!$	Sign
	$\mid$	$\#$	Hash
	$\mid$	$C\ \rightarrow\ C$	Linear sequence
	$\mid$	$C\ \ D{<}D\ \ C$	Sequential branching
	$\mid$	$C\ \ D{\sim}D\ \ C$	Parallel branching
	$\mid$	$(\ C\ )$	Grouping
D	$\leftarrow$	$-\mid+$	Splitting specification

Figure 1: Copland Syntax

$*p : t$. All phrases we consider in this paper are of the form $*\mathsf{bank} : t$. The syntax of Copland phrases is presented next.

A symbol S is a sequence of lowercase letters. A place P is a sequence of digits or a symbol. A legal Copland phrase t is in the syntactic category C as defined in Fig. 1. The branching operators are non-associative and have the same precedence, and the linear sequencing operator is right associative and has a higher precedence.

A complete description of the combining operators in Copland requires a description of the flow of evidence collected by a phrase, however, for this paper, the particulars of evidence collection are irrelevant. Instead, we focus on the events used to collect and combine evidence, and their orderings. Features of the language that are relevant only for reasoning about evidence will be identified in the presentation of the language.

Semantically, a Copland phrase specifies the mapping of input evidence to some output evidence via one or more attestation events. Each event occurs at a well defined place. Events involve taking measurements, signing or hashing evidence, or routing evidence so as to produce combinations of evidence.

To describe the semantics of executing Copland phrase C at place P, we describe how it transforms evidence making use of diagrams of the form:

The box represents an event, and the arrows show the flow of evidence. The ordering of events respects the flow of evidence. The syntax of an event label is given in Figure 2.

The most basic Copland phrase is a measurement $m\ q\ t$, for symbols m and t, and place q, where m names the probe, t names the target of the measurement, and q is the place at which the target resides. When at place p, $m\ q\ t$ means that p should receive some evidence, perform m targeting t at q, and then emit the resulting evidence.

$$
\begin{array}{rcl}
L & \leftarrow & \mathsf{msp}(P.S, P.S) \mid \mathsf{nul}(P) \mid \mathsf{cpy}(P) \mid \mathsf{sig}(P) \mid \mathsf{hsh}(P) \\
 & \mid & \mathsf{req}(P, P) \mid \mathsf{rpy}(P, P) \mid \mathsf{split}(P, D, O, D) \mid \mathsf{join}(P, P) \\
O & \leftarrow & < \mid \sim
\end{array}
$$

Figure 2: Event Label Syntax

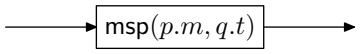

The semantics of the phrases nullify {}, copy _, sign !, and hash # have the same form as a measurement. When executing at place p, their corresponding event labels are $\mathsf{nul}(p)$, $\mathsf{cpy}(p)$, $\mathsf{sig}(p)$, and $\mathsf{hsh}(p)$.

Let $p : \{c\}$ be the events and their orderings associated with executing phrase c at p. Measurements can be combined in a pipeline fashion using the $\rightarrow$ operator. Thus when at p, $c_1 \rightarrow c_2$ means

A measurement can be taken at a remote location using the @ operator. When at p, $@_q\,[c]$ means

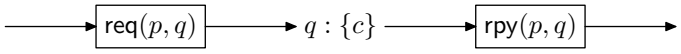

Phrases c_1 and c_2 can be combined using branching. There are two ways of combining phrases using branching, sequential ($o = <$) and parallel ($o = \sim$) combination. They both follow the same split-join pattern. When at p, $c_1\ d_1 o d_2\ c_2$ means

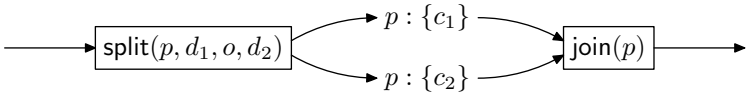

The split specifications d_1 and d_2 only effect the flow of evidence. When $d = +$, evidence is passed, and when $d = -$, evidence is dropped. There are additional orderings associated with sequential branching. When $o = <$, all of the events associated with c_1 precede the events associated with c_2.

Formally, the semantics of a Copland phrase is given by an event system.

Definition 5 (Event System)**.** An *event system* $(E, \prec, \ell)$ consists of

1. set E of events,
2. relation $\prec \subseteq E \times E$, a strict partial order, and
3. function $\ell : E \rightarrow L$, a map from events to labels.

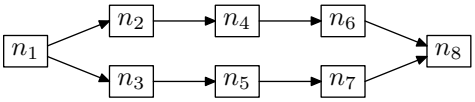

$\ell(n_1) = \mathsf{split}(\mathsf{rp}, -, \sim, -)$ $\quad$ $\ell(n_5) = \mathsf{msp}(\mathsf{us.bmon}, \mathsf{us.bser})$
$\ell(n_2) = \mathsf{req}(\mathsf{rp}, \mathsf{ks})$ $\quad$ $\ell(n_6) = \mathsf{rpy}(\mathsf{rp}, \mathsf{ks})$
$\ell(n_3) = \mathsf{req}(\mathsf{rp}, \mathsf{us})$ $\quad$ $\ell(n_7) = \mathsf{rpy}(\mathsf{rp}, \mathsf{us})$
$\ell(n_4) = \mathsf{msp}(\mathsf{ks.av}, \mathsf{us.bmon})$ $\quad$ $\ell(n_8) = \mathsf{join}(\mathsf{rp})$

Figure 3: Event System for Example 2

The event system for Example 2 is presented in Figure 3. The branching operation $\sim$ contributes nodes n_1 and n_8. Measurements occur at nodes n_4 and n_5. Events n_2 and n_6 causes the measurement at n_4 to occur at ks. Events n_3 and n_7 causes the measurement at n_5 to occur at us. The strict partial order $\prec$ is the transitive closure of the $\rightarrow$ relation shown in the diagram in Figure 3.

The events and labels for Example 3 are the same as they are for Example 2 except that $\ell(n_1) = \mathsf{split}(\mathsf{rp}, -, <, -)$. The other difference is the nodes are linearly ordered.

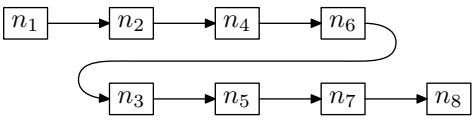

In this work, we will focus solely on measurement events, and abstract away all other kinds of events. As a result, sequential branching and pipelines have the same semantics.

4 Theory of Layered Attestation

In the previous section, we defined the semantics of Copland phrases in the absence of adversary intervention. This section concerns understanding how an adversary can interfere with an attestation. It is primarily a summary of the main definitions and results of [18].

The Copland semantics of the previous section does not identify any adversary events. It represents the ideal execution in the absence of an adversary. However, the purpose of an attestation is to reliably gather enough evidence of the system state to determine that some components of interest have not been corrupted by an adversary (and therefore can reasonably be expected to behave predictably). To analyze the security of remote attestations, we must account for adversary events and their consequences.

We start by enriching the partially ordered event semantics for Copland with labels for two new types of adversary events: corruption of a component, and repair of a component.

$\ell(e) = \mathsf{cor}(p.c)$ asserts that event e corrupts component $p.c$, and

$\ell(e) = \mathsf{rep}(p.c)$ asserts that event e repairs component $p.c$.

Adversary events have no associated evidence.

These adversary events serve to toggle the corruption state of system components between being *regular* or *corrupted*. In turn, the corruption state of components will affect the behavior of other events. In the current work, we consider an adversary model in which corrupted components only effect the outcome of measurement events. For all other events (req, rpy, etc.) corrupted components have no effect. The reason for this focus is that it corresponds to the adversary model discussed in [18] which only contained measurement events. Since the principles developed in that work are the basis of the current study, we restrict ourselves to such an adversary. We plan to study adversaries with more advanced capabilities (such as tampering with evidence it receives from other components) in future work.

Measurement events produce new pieces of evidence about the target being measured. While the actual values of the evidence data will be complex and varied, our primary concern is whether the evidence passes or fails appraisal at the relying party. We therefore adopt the following idealization of the outcome of measurement. A measurement event e can generate evidence that will either pass or fail appraisal. We also assume that a corrupted measurer has the incentive and ability to always generate evidence that will pass appraisal, regardless of whether or not the corruption state of the target of measurement. Conversely, (with one exception described below) a regular measurer will always generate evidence whose appraisal accurately describes the corruption state of the target. So if the target is regular at the time of measurement the measurer will produce evidence that passes appraisal, and if the target is corrupt, the measurer will generate evidence that will not pass appraisal. In the latter case we say that the measurement event detects the corruption.

The one exception to the above rule is when the measurer relies on some other component in order to work effectively. A good example of this is anti-virus software. Its access to the file system is mediated through the operating system kernel. So, if the kernel is corrupted, it could hide a corrupted target file by presenting the anti-virus software with a pristine target instead. Thus, whenever a measurer relies on another component in this way, and that component is corrupted, we assume it always has the incentive and ability to cause the measurer to generate evidence that will pass appraisal. These assumptions are summarized in Table 1, where we use * to indicate that the outcome of measurement is the same regardless of the corruption state.

Because of the effect the dependency relation (i.e. the relation between anti-virus and kernel) has on the outcome of measurement, we must keep track of these potential context dependencies that may arise in a given attestation. A component may have zero, one, or more such dependencies. When $p.m$ depends on $q.c$ we write $\mathsf{Depends}(p.m, q.c)$. Typically the places p and q will be the same so we make that simplifying assumption throughout this paper. Since, according to Table 1, the appraisal only fails when the target is actually corrupt, this

Table 1: The effect of corruption on measurement.

Measurer	Context	Target	Evidence Appraisal
corrupt	*	*	passes
*	corrupt	*	passes
regular	regular	regular	passes
regular	regular	corrupt	fails

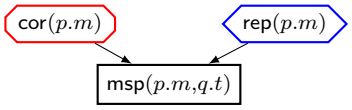

Figure 4: An example in which the notion of corruption state is ill-defined.

means these cases accurately detect corruption. We therefore write $\mathsf{Detects}(e)$ to denote the fact that the measurement event e produces evidence that will fail appraisal allowing the relying party to detect a corruption. As we will see below, since we are most interested in cases where the relying party is fooled into trusting a corrupted system, we will focus on instances in which the Detects predicate is empty, i.e. all measurement events generate evidence that will pass appraisal.

Adversary ordered. The assumptions codified in Table 1 only make sense if the corruption state of a component at an event is well defined. To see why there might be an issue, consider the partially ordered set of events in Fig. 4. Since the corruption of $p.m$ is neither before nor after its repair, it is unclear whether we should consider $p.m$ to be corrupt or regular at the measurement event. The problem is that there are two adversary events affecting $p.m$ that are maximal in the precedence ordering before the measurement event. In the adversary-enriched semantics, we need to ensure there is only ever at most one maximal adversary event affecting a given component prior to any event.

Notice also that the only corruption states that matter at the measurement event are those of $p.m$, of $q.t$, and any components $p.m$ depends on to perform its measurement. Other corrupted components cannot interfere with the outcome of this event. Similarly, if the repair event affected $q.t$ instead, there would be no ambiguity. So, we only need to require at most one maximal adversary event affecting each component that is involved in the event.

Definition 6. Component $p.c$ is *relevant* to event e iff

1. $\ell(e) = \mathsf{msp}(p.c, q.t)$ or $\ell(e) = \mathsf{msp}(q.m, p.c)$, or
2. $\ell(e) = \mathsf{msp}(p.m, q.t)$ and $\mathsf{Depends}(p.m, p.c)$, or
3. $\ell(e) = \mathsf{cor}(p.c)$ or $\ell(e) = \mathsf{rep}(p.c)$.

There is a simple condition that will guarantee the corruption state of components is well defined.

Definition 7 (Adversary-Ordered)**.** Event system $(E, \prec, \ell)$ is *adversary-ordered* iff for each $e_1, e_2 \in E$, if $p.c$ is is relevant to both e_1 and e_2, and e_2 is an adversary event, then e_1 and e_2 are comparable events, that is $e_1 \prec e_2$, $e_2 \prec e_1$, or $e_1 = e_2$.

It turns out (c.f. [18] Lemma 1) that adversary-ordered event systems guarantee that the corruption state of components is well-defined at events they are relevant to. That is, any adversary-ordered event system E induces a predicate Cor_E that identifies all and only those components that are corrupt at an event to which they are relevant. To determine whether $\mathsf{Cor}_E(p.c, e)$ holds, we simply look backwards in the ordering to find the (unique) most recent adversary event for the given component. If it is a corruption, then $\mathsf{Cor}_E(p.c, e)$ holds. If it is a repair, or if there are no prior adversary events for the given component, then $\neg\, \mathsf{Cor}_E(p.c, e)$. Following the rules in Table 1, adversary-ordered event systems also uniquely determine the $\mathsf{Detects}_E$ predicate that identifies which events accurately detect corruptions. To be explicit, $\mathsf{Detects}_E(e)$ iff $\mathsf{Cor}_E(p.t, e)$ holds for the target of measurement $p.t$, and $\neg\, \mathsf{Cor}_E(q.c, e)$ for all other components $q.c$ relevant to e. Otherwise $\neg\, \mathsf{Detects}_E(e)$. For the remainder of the paper, we restrict our attention to adversary-ordered event systems.

The output of the Copland semantics for a phrase t is trivially adversary-ordered. Since there are no corruption events, Cor_E is the empty predicate. However, the purpose of performing an attestation is to detect corruptions of components of interest. We therefore are interested in all the ways to enrich the (adversary-free) Copland event semantics with adversary events.

Definition 8 (Adversary-Enriched)**.** Let $(E, \prec, \ell)$ be an (adversary-ordered) event system. Then (adversary-ordered) event system $(E', \prec', \ell')$ *adversary-enriches* $(E, \prec, \ell)$ iff

1. $E \subseteq E'$,
2. $\prec \subseteq \prec'$,
3. for $e \in E$, $\ell(e) = \ell'(e)$, and
4. every $e \in E' \setminus E$, is an adversary event.

We often abuse notation and write E to denote the triple $(E, \prec, \ell)$ leaving $\prec$ and ℓ implicit. Definition 8 defines a natural partial order on executions: $E \leq E'$ iff E' adversary-enriches E.

Definition 9 (Execution)**.** An *execution* of Copland phrase t is an event system that adversary-enriches the Copland semantics for t. We write $\mathcal{E}(t)$ to denote the set of all executions of phrase t.

The set $\mathcal{E}(t)$ will contain many executions,[1] some of which detect corruptions (i.e. $\mathsf{Detects}_E(e)$ for some measurement event e) and some of which do not

[1] In fact, $\mathcal{E}(t)$ is infinite because of useless chains of cor and rep events. We may want to define *minimal* executions. There are still exponentially many of those.

(i.e. $\mathsf{Detects}_E$ is the empty predicate). Our aim is to provide a mechanism for querying $\mathcal{E}(t)$ to discover the set of executions that satisfy some constraints or assumptions of interest. The most important queries are those that yield executions in which the adversary successfully avoids detection. Thus we may assume that $\mathsf{Detects}_E$ is empty, and that $\mathsf{Cor}_E(p.c, e)$ for some component $p.c$ and some measurement event e. We then want to find all executions in $\mathcal{E}(t)$ consistent with these assumptions.

More generally, we consider queries of the form (E, φ) where φ is a predicate identifying which components we are assuming to be corrupt at which events. That is $\varphi(p.c, e)$ indicates an assumption that $p.c$ is corrupted at event e.

Just as we defined a partial order on event systems above, we can define a partial order on corruption predicates φ by saying $\varphi \leq \varphi'$ iff for all $p.c$ and for all e, $\varphi(p.c, e)$ implies $\varphi'(p.c, e)$. These two orderings combined create a natural partial order on queries (E, φ): $(E, \varphi) \leq (E', \varphi')$ iff $E \leq E'$ and $\varphi \leq \varphi'$. The ordering is strict iff either of the constituent orderings is strict. We use this ordering to formalize our search goal from two paragraphs ago.

Definition 10. The *denotation* of query (E, φ), written $[\![(E, \varphi)]\!]$, is the set of executions defined by the following rules. $E' \in [\![(E, \varphi)]\!]$ iff all of the following hold:

(a) $E \leq E'$

(b) $\varphi \leq \mathsf{Cor}_{E'}$

(c) $\mathsf{Detects}_{E'}$ is the empty predicate.

Condition (a) says we are interested only in adversary enrichments of the given execution E. Condition (b) says that E' satisfies any assumption we made in φ about components being corrupt. Condition (c) says that we are only interested in executions that do not detect corruptions.

Definition 10 is a clear definition of the set we would like to enumerate, but it does not immediately suggest any procedure for doing so. In fact, our approach is not to enumerate all executions in (E, φ), but rather to enumerate only those executions that are minimal in the $\leq$ ordering on event systems. All other executions in the denotation require the adversary to perform at least as much work as one of the minimal ones. In that sense, we don't aim to enumerate all executions in the denotation, but rather, we aim to *characterize* all executions in the denotation by enumerating the minimal ones. The core idea is to perform a search by climbing in the $\leq$ ordering on queries until we reach a stopping condition for a query (E', φ') that allows us to include E' in our enumeration. We next define the stopping condition which relies on the following definition.

Definition 11. The query (E, φ) is *saturated* if and only if

(a) $\varphi = \mathsf{Cor}_E$, and

(b) $\mathsf{Detects}_E$ is the empty predicate.

Condition (a) here is similar to condition (b) in Def. 10, except it further restricts how Cor_E is allowed to extend φ. In particular, it says that E cannot contain any corrupted components at events not already identified by φ. Saturated queries signal that a search in the query ordering can stop, as indicated by the following two lemmas.

Lemma 12. If (E, φ) is saturated, then $E \in [\![(E, \varphi)]\!]$.

Proof. We check the three conditions of Def. 10. Clearly, $E \leq E$. By Def. 11, $\mathsf{Detects}_E$ is empty. Finally, $\mathsf{Cor}_E = \varphi$ implies $\varphi \leq \mathsf{Cor}_E$. □

Lemma 13. If (E, φ) is saturated and $(E, \varphi) < (E', \varphi')$ then $E \notin [\![(E', \varphi')]\!]$.

Proof. Either $E < E'$ or $\varphi < \varphi'$. In the first case, either E' has strictly more events than E or E' has strictly more orderings than E. Either way, this extra structure is preserved by all $\hat{E} \geq E'$, so all elements of $[\![(E', \varphi')]\!]$ must also contain this extra structure. It follows immediately that $E \notin [\![(E', \varphi')]\!]$.

In the other case where $\varphi < \varphi'$, there is some $(p.c, e)$ such that $\neg\varphi(p.c, e)$ and $\varphi'(p.c, e)$. By Def. 10 condition (b), for every element $\hat{E} \in [\![(E', \varphi')]\!]$, $\mathsf{Cor}_{\hat{E}}(p.c, e)$. But since $\mathsf{Cor}_E = \varphi$ and $\neg\varphi(p.c, e)$, $E \notin [\![(E', \varphi')]\!]$. □

Lemmas 12 and 13 show why the definition of saturated is a useful and natural stopping condition. By Lemma 12, saturated queries identify elements of the denotation, and by Lemma 13 if we didn't stop the search, we would miss the execution encoded by the saturated query. The following theorem tells us that by enumerating the minimal saturated queries above (E, φ), we capture precisely the denotation of (E, φ).

Theorem 14. Let $R = \{(E', \varphi') \mid (E, \varphi) \leq (E', \varphi')$ and (E', φ') is saturated$\}$. Let R_{min} be the $\leq$-minimal members of R. Then

$$[\![(E, \varphi)]\!] = \bigcup_{(E', \varphi') \in R_{min}} [\![(E', \varphi')]\!].$$

Proof. The reverse inclusion is easy to show. We know $E \leq E'$ and $\varphi \leq \varphi'$ by the definition of R_{min}. Suppose $\hat{E} \in \bigcup_{(E', \varphi') \in R_{min}} [\![(E', \varphi')]\!]$. Then for some $(E', \varphi') \in R_{min}$, we have $E' \leq \hat{E}$, $\varphi' \leq \mathsf{Cor}_{\hat{E}}$, and $\mathsf{Detects}_{\hat{E}}$ is empty. By the transitivity of $\leq$, we easily conclude that $E \leq E$ and $\varphi \leq \mathsf{Cor}_{\hat{E}}$.

Now consider the forward inclusion. Suppose $\hat{E} \in [\![(E, \varphi)]\!]$. Then by Def. 10, $E \leq \hat{E}$, $\varphi \leq \mathsf{Cor}_{\hat{E}}$, and $\mathsf{Detects}_{\hat{E}}$ is empty. Now consider the query $(\hat{E}, \mathsf{Cor}_{\hat{E}})$. It is saturated because, trivially, $\mathsf{Cor}_{\hat{E}} = \mathsf{Cor}_{\hat{E}}$, and we already saw $\mathsf{Detects}_{\hat{E}}$ is empty. Furthermore, we already knew that $\varphi \leq \mathsf{Cor}_{\hat{E}}$. So $(\hat{E}, \mathsf{Cor}_{\hat{E}}) \in R$. Thus there must be some $(E', \varphi') \in R_{min}$ such that $(E', \varphi') \leq (\hat{E}, \mathsf{Cor}_{\hat{E}})$. So $E' \leq \hat{E}$ and $\varphi' \leq \mathsf{Cor}_{\hat{E}}$. So $\hat{E} \in [\![(E', \varphi')]\!]$ as desired. □

Up to now, we have described the theory of saturated queries and showed why that theory correctly captures the denotation we want. We have not produced any search algorithm to enumerate saturated queries. Our approach is

to leverage a general-purpose model finder for geometric logic to implement the search. We will provide the model finder with an axiomatization of the theory of saturated queries, and the job of the model finder is to enumerate minimal (and possibly non-minimal) models of the theory. Theorem 14 tells us that if we correctly axiomatize the theory and if the model finder performs correctly, then it will enumerate the set we want. Before turning to our axiomatization of the theory of saturated queries, we first provide an overview of the model finder discussing what it does and how it works.

5 Model Finding with Chase

Chase [16] is a model finder for first-order logic with equality. It finds minimal models of a theory expressed in finitary special geometric form, where functions in models may be partial. A formula is in *finitary special geometric form* if it is a finite sentence consisting of a single implication, the antecedent is a conjunction of atomic formulas, and the consequent is a disjunction. Each disjunct is a possibly existentially quantified conjunction of atomic formulas.

$$\forall \vec{x}.\, P_1(\vec{x}) \wedge \cdots \wedge P_n(\vec{x}) \Rightarrow \bigvee_i \exists \vec{y}_i.\, Q_{i,1}(\vec{x}, \vec{y}_i) \wedge \cdots \wedge Q_{i,n_i}(\vec{x}, \vec{y}_i)$$

A function is *partial* if it is defined only on a proper subset of its domain. A sentence in first-order logic is *finitary geometric* iff it is logically equivalent to a finite set of sentences in finitary special geometric form. Finitary geometric logic is also called coherent logic.

We will assume familiarity with basic ideas and results from first-order mathematical logic; notions that are not defined here are treated in any text on logic [6] with allowances for partial functions. When a structure A satisfies theory T, we write $A \models T$ and call A a model of T. The definition of a homomorphism must account for partial functions.

Definition 15 (Homomorphism)**.** Let A and B be structures. A *homomorphism* h of A into B is a function with these properties

1. For each n-place predicate P,

$$P(a_1, \ldots, a_n) \in A \text{ implies } P(h(a_1), \ldots, h(a_n)) \in B.$$

2. For each n-place function f,

$$f(a_1, \ldots, a_n) = a_0 \in A \text{ implies } f(h(a_1), \ldots, h(a_n)) = h(a_0) \in B.$$

We write $A \ll B$ when there is a homomorphism from A into B.

Definition 16 (Minimal Model)**.** Model A of T is *minimal* iff for all models B of T, whenever $B \ll A$ then $A \ll B$.

Definition 17 (Set of Support)**.** A set of models M is a *set of support* for theory T iff whenever $B \models T$, there exists a model $A \in M$ such that $A \ll B$.

When given a theory, Chase produces a set of support whenever it terminates successfully. It may produce some models that are not minimal.

5.1 Input Syntax

The input to the Chase program is a set of sequents written in a slight variant of Geolog [7] syntax. The syntax is inspired by Prolog. Quantification is implicit and constants and variables are distinguished using capitalization, where variables are capitalized.

A Chase sequent has the form

$$A_1 \texttt{ \& } A_2 \texttt{ \& } \cdots \texttt{ \& } A_m \texttt{ => } C_1 \texttt{ | } C_2 \texttt{ | } \cdots \texttt{ | } C_n \texttt{.}$$

The formula to the left of `=>` is the antecedent, and the consequent is to the right. Each conjunct A_i in the antecedent is an atomic formula. The consequent is a disjunction. Each disjunct C_j is a conjunction of the form

$$B_{j,1} \texttt{ \& } B_{j,2} \texttt{ \& } \cdots \texttt{ \& } B_{j,p_j}$$

where each conjunct $B_{j,k}$ is an atomic formula. A consequent with no disjuncts is indicated by the reserve symbol `false`. A conjunction with no conjuncts is indicated by the reserved symbol `true`. When the antecedent is true, the tokens `true =>` may be omitted.

A symbol is a letter followed by a sequence of letters, dollar signs ($), and underscores (_). An atomic formula is a predicate symbol (a symbol not capitalized) applied to a parenthesized sequence of comma separated terms, or an equality consisting of two terms separated by the = sign. A term is a variable (a capitalized symbol), a constant (a symbol not capitalized), or a function symbol (a symbol not capitalized) applied to a parenthesized sequence of comma separated terms.

Comments start with `%` and continue until the end of the line. An example of a theory for input follows.

```
% Total Ordering
num(a) & num(b).
num(X) & num(Y) => lt(X, Y) | X = Y | lt(Y, X).
```

The variables that occur in the antecedent of a sequent are universally quantified. The variables that occur in disjunct B_j and not in the antecedent are existentially quantified over the disjunct. Thus,

```
p(X) | q(X).
```

means $(\exists X.\, p(X)) \vee (\exists X.\, q(X))$ *not* $\exists X.\, (p(X) \vee q(X))$.

To specify that X is universally quantified in the consequent, add `X=X` to the antecedent. Thus $\forall X.\, p(X)$ is written as

```
X = X => p(X).
```

Within this paper, we typeset theories using connectives from mathematical logic. Thus the Total Ordering Example from above becomes

```
% Total Ordering
num(a) ∧ num(b).
num(X) ∧ num(Y) ⇒ lt(X, Y) ∨ X = Y ∨ lt(Y, X).
```

5.2 Structures

A chase structure for theory T is a set of facts. A *fact* is a ground atomic formula that has one of two forms

1. $P(c_1, \ldots, c_n)$, or
2. $f(c_1, \ldots, c_n) = c_0$.

where P and f are in the signature of T.

The universe of structure A is the least set U of constants such that

1. $P(c_1, \ldots, c_n) \in A$ implies $c_1, \ldots, c_n \in U$, and
2. $f(c_1, \ldots, c_n) = c_0 \in A$ implies $c_0, \ldots, c_n \in U$.

Let C be the set of constants that occur in theory T. A structure A produced by Chase for theory T has the following properties.

1. Functions may be partial.
2. Each constant in C is the left hand side of a fact in A.
3. Equality in A is closed under congruence.
4. Each element in U is the canonical representative of an equivalence class induced by congruence closure.

One of the three models found by Chase for the Total Ordering Example is

$$\{\mathtt{num(a)}, \mathtt{num(b)}, \mathtt{lt(a, b)}, \mathtt{a = a}, \mathtt{b = b}\}.$$

5.3 Algorithm

Models of theory T are found using an algorithm called the chase. The procedure starts with a structure in which each constant in T is equated to itself. Queue Q is created containing the initial structure, and the main loop begins.

The chase for theory T repeats the following steps until queue Q is empty.

1. Take structure A from Q.
2. If A models T then output A.

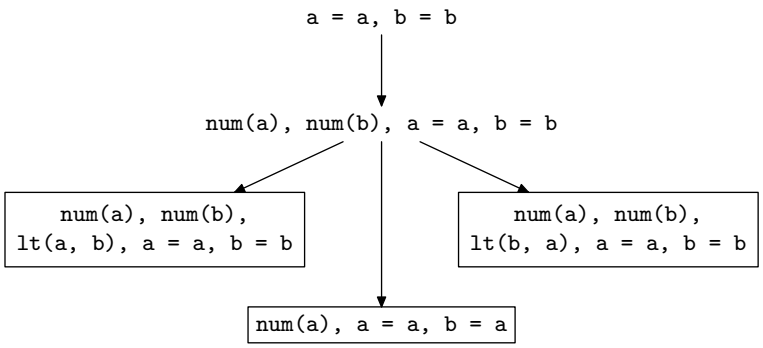

Figure 5: Structures Generated From the Total Ordering Example

3. Otherwise, choose a formula F in T not satisfied by A.
   (a) Find a variable assignment S for the universally quantified variables in F such that its antecedent is satisfied, but its consequent is not.
   (b) Apply S to each disjunct in the consequent.
   (c) For each disjunct, substitute a freshly generated constant for each existentially quantified variable, and add to the queue a structure produced by augmenting A with the disjunct. Mark A as being the parent of the new structure.

The structures generated by a run of Chase on the Total Ordering Example are shown in Fig. 5. An arrow connects a parent structure with a child. The boxed structures are models of the theory.

The initial structure just equates constants. The second structure generated adds `num` facts. At this point, the only applicable sentence is the one that imposes a total ordering. It produces models, and Chase terminates successfully.

In general, when structure A is the parent of B, there exists a homomorphism from A into B. Every chase step is structure-preserving.

6 Chase Theory for Layered Attestation

At the end of Section 4 we identified the goal of correctly axiomatizing the theory of saturated queries. As queries (E, φ) contain adversary-ordered event systems as one component of the pair, we first axiomatize the theory of adversary-ordered event systems.

Throughout this section, we will omit numerous rules in the axiomatization that do not provide any insight, such as rules that express the injectivity of event labels. Such rules are important for a correct axiomatization, but they represent facts that are often taken for granted by humans and detract from the underlying logical principles.

We note that it suffices to axiomatize strict partial orders whose events satisfy the adversary-ordered condition of Def. 7. We do not need to include rules that state all events have a label, or that all labels are drawn from some finite set as long as our rules preserve such facts (which they will). We do need to include auxiliary rules that introduce predicates used by our core axiomatization. Typically such rules summarize several facts with one predicate, thereby making the syntax for other rules much more natural to read. For example, we include a formula to introduce the abbreviation `ms_evt` for measurement events. Fortunately, the axioms for strict partial orders and the definition of adversary-ordered already present themselves naturally in finitary special geometric form as shown in Fig. 6.

```
% Strict partial order:
prec(E, E) ⇒ false.
prec(E1, E2) ∧ prec(E2, E3) ⇒ prec(E1, E3).

% Definition of relevant:
l(E) = msp(P2, M, P1, T) ∧ dep(P2, C, M)
 ⇒ relevant(P2, M, E) ∧ relevant(P1, T, E)
  ∧ relevant(P2, C, E).
l(E) = msp(P2, M, P1, T)
 ⇒ relevant(P2, M, E) ∧ relevant(P1, T, E).
l(E) = cor(P, C) ⇒ relevant(P, C, E).
l(E) = rep(P, C) ⇒ relevant(P, C, E).

l(E) = msp(P2, M, P1, T) ∧ relevant(P, C, E)
 ⇒ P = P1 ∧ C = T ∨ P = P2 ∧ C = M ∨ P = P2 ∧ dep(P2, C, M).

% Adversary ordered:
relevant(P, C, E1) ∧ relevant(P, C, E2) ∧ l(E1) = cor(P, C)
 ⇒ prec(E1, E2) ∨ prec(E2, E1) ∨ E1 = E2.
relevant(P, C, E1) ∧ relevant(P, C, E2) ∧ l(E1) = rep(P, C)
 ⇒ prec(E1, E2) ∨ prec(E2, E1) ∨ E1 = E2.

% Introduction of ms_evt:
l(E) = msp(P2, M, P1, T) ⇒ ms_evt(E).
```

Figure 6: The Chase theory for adversary-ordered event systems.

The theory of any given phrase will also contain a formula stating that if E is a measurement event, it must be one of the events generated by the Copland semantics. This enforces Def. 8 ensuring models found by Chase are adversary enrichments of the given Copland semantics. That formula is necessarily dependent on the Copland phrase being analyzed, so it is not included as part of the general theory.

Axiomatizing the theory of saturated queries is not as straightforward. The

theory as presented in Section 4 relies on the derived predicates $\mathsf{Detects}_E$ and Cor_E. These have properties that are not preserved under model homomorphisms (i.e. adversary-enrichments). That is, if we know $E \leq E'$ we cannot necessarily conclude that $\mathsf{Cor}_E \leq \mathsf{Cor}_{E'}$. As a case in point, consider some execution E in which $\mathsf{Cor}_E(p.c, e)$ holds at some event e. If we enrich E to E' by adding a repair event for $p.c$ just prior to e, then $\neg\, \mathsf{Cor}_{E'}(p.c, e)$. When $E \leq E'$, we similarly cannot conclude $\mathsf{Cor}_{E'} \leq \mathsf{Cor}_E$, nor can we conclude anything about the relation between $\mathsf{Detects}_E$ and $\mathsf{Detects}_{E'}$. Since formulas in finitary special geometric form are precisely those that are preserved under homomorphism, and since the $\leq$ relation on the derived predicates Cor_E and $\mathsf{Detects}_E$ are not preserved under homomorphism, we cannot directly reference $\mathsf{Detects}_E$ or Cor_E in our axiomatization.

We therefore need to find an equivalent set of formulas in finitary special geometric form that do not reference any facts not preserved by homomorphism. We use the four formulas depicted in Fig. 7. It is far from obvious that these four formulas are equivalent to the theory of saturated queries. The following theorem demonstrates that they are indeed equivalent.

```
% Rule 1
l(E) = msp(P2, M, P1, T) ∧ phi(P1, T, E)
  ⇒ phi(P2, M, E) ∨ dep(P2, C, M) ∧ phi(P2, C, E).

% Rule 2
phi(P, C, E1)
 ⇒ prec(E0, E1) ∧ l(E0) = cor(P, C).

% Rule 3
prec(E1, E2) ∧ phi(P, C, E2) ∧ l(E1) = rep(P, C)
  ⇒ prec(E1, E3) ∧ prec(E3, E2) ∧ l(E3) = cor(P, C).

% Rule 4
l(E1) = cor(P2, C) ∧ ms_evt(E2) ∧ prec(E1, E2) ∧
relevant(P2, C, E2)
  ⇒ phi(P2, C, E2)
    ∨ prec(E1, E3) ∧ prec(E3, E2) ∧ l(E3) = rep(P2, C).
```

Figure 7: The Chase theory for saturated queries.

Theorem 18. The query (E, φ) is saturated if and only if none of Rules 1–4 applies.

Proof. ($\Rightarrow$): We proceed by contrapositive, assuming one of the four rules does apply, and conclude that (E, φ) is unsaturated.

Rule 1: Assume Rule 1 applies. Then there is some measurement event e in which component $p.m$ measures component $q.t$, and $\varphi(q.t, e)$ but $\neg\varphi(p.m, e)$ and for every $p.c$ in the context of $p.m$, $\neg\varphi(p.c, e)$. If $\mathsf{Cor}_E \neq \varphi$ then, by

definition, (E, φ) is unsaturated which is what we aim to show. So we may assume $\mathsf{Cor}_E = \varphi$. Then $\neg\, \mathsf{Cor}_E(p.m, e)$ and for every $p.c$ in the context of $p.m$, $\neg\, \mathsf{Cor}_E(p.c, e)$. Then by the definition of $\mathsf{Detects}_E$ as described in Table 1, $\mathsf{Detects}_E(e)$ and so is not empty. Thus, by definition, (E, φ) is unsaturated.

Rule 2: Assume Rule 2 applies. Then for some $p.c$ and e, $\varphi(p.c, e)$, but there is no $e' \prec e$ with $\ell(e') = \mathsf{cor}(p.c)$. Then $\neg\, \mathsf{Cor}_E(p.c, e)$, and hence $\mathsf{Cor}_E \neq \varphi$. So, by definition, (E, φ) is unsaturated.

Rule 3: Assume Rule 3 applies. Then there is some repair event $\ell(e_1) = \mathsf{rep}(p.c)$ with $e_1 \prec e_2$ such that $\varphi(p.c, e_2)$, but there is no intervening corruption event $e_1 \prec e' \prec e_2$ such that $\ell(e') = \mathsf{cor}(p.c)$. Thus, the last adversary event for $p.c$ before e_2 is a repair event. This means $\neg\, \mathsf{Cor}_E(p.c, e_2)$. But $\varphi(p.c, e_2)$, so $\mathsf{Cor}_E \neq \varphi$, making (E, φ) unsaturated.

Rule 4: Assume Rule 4 applies. Then there is some corruption event $\ell(e_1) = \mathsf{cor}(p.c)$ preceding a measurement event e_2 to which $p.c$ is relevant where $\neg\varphi(p.c, e_2)$ and there is no intervening repair event $e_1 \prec e' \prec e_2$ with $\ell(e') = \mathsf{rep}(p.c)$. Since there is no intervening repair event, the last adversary event for $p.c$ before e_2 is a corruption event. This means that $\mathsf{Cor}_E(p.c, e_2)$. But $\neg\varphi(p.c, e_2)$, so $\mathsf{Cor}_E \neq \varphi$, making (E, φ) unsaturated.

This concludes the proof of one direction.

($\Leftarrow$): We suppose that (E, φ) is unsaturated and we demonstrate that one of Rules 1–4 applies. We take cases on whether $\mathsf{Cor}_E = \varphi$.

$\mathsf{Cor}_E = \varphi$: In this case, since (E, φ) is unsaturated, there must be an event e for which $\mathsf{Detects}_E(e)$. By examining Table 1, the only possibility is for $\neg\, \mathsf{Cor}_E(p.m, e)$ and for all $p.c$ in the context of $p.m$, $\neg\, \mathsf{Cor}_E(p.c, e)$, whereas the target of measurement $q.t$ satisfies $\mathsf{Cor}_E(q.t, e)$. Since $\mathsf{Cor}_E = \varphi$, $\varphi(q.t, e)$ but $\neg\varphi(p.m, e)$ and $\neg\varphi(p.c, e)$. Thus Rule 1 applies.

$\mathsf{Cor}_E \neq \varphi$: So there is some event e and some component $p.c$ such that either $\neg\, \mathsf{Cor}_E(p.c, e)$ and $\varphi(p.c, e)$, or $\mathsf{Cor}_E(p.c, e)$ and $\neg\varphi(p.c, e)$.

Suppose first that $\neg\, \mathsf{Cor}_E(p.c, e)$ and $\varphi(p.c, e)$. From $\neg\, \mathsf{Cor}_E(p.c, e)$ we know that either there are no adversary events for $p.c$ before e, or, if there are some, the most recent one has the label $\mathsf{rep}(p.c)$. In the first case, Rule 2 applies, because $\varphi(p.c, e)$ holds without a prior corruption event. In the second case, Rule 3 applies because $\varphi(p.c, e)$ holds with a prior repair event for $p.c$, but without an intervening corruption event.

Finally, suppose $\mathsf{Cor}_E(p.c, e)$ and $\neg\varphi(p.c, e)$. From $\mathsf{Cor}_E(p.c, e)$ we know that there is a corruption event for $p.c$ prior to e without any intervening repair event. But since $\neg\varphi(p.c, e)$, this describes the conditions for Rule 4 to apply. □

Theorem 18 tells us that models found with Chase using the theories of Figs. 6 and 7 will be saturated queries. As Chase is designed to produce a set of support, if we feed Chase with a formula representing an initial query (E, φ), its output will include all the minimal saturated queries (E', φ') such that $(E, \varphi) \leq (E', \varphi')$. So by Thm 14 and Lemma 12, we can project these queries onto their executions E' to obtain the minimal elements of the denotation $[\![\, (E, \varphi) \,]\!]$. These minimal executions characterize all the ways an adversary can avoid detection

consistent with the original query (E, φ). For any other (non-minimal) execution in the denotation, there is a minimal one in which the adversary performs less work (i.e. strictly fewer actions or strictly weaker orderings). An analyst can then inspect the minimal executions to determine if the work required of an adversary is sufficiently difficult. We demonstrate such analyses in the next section.

7 Analysis of Copland Phrases

In this section we build on the examples from Section 2 to show how to use Chase and our input theory to analyze the trust properties of Copland phrases. We start by returning to Example 2 from Section 2:

$$*\mathsf{bank} : @_{\mathsf{ks}}\,[\mathsf{av\ us\ bmon}] \sim @_{\mathsf{us}}\,[\mathsf{bmon\ us\ exts}]$$

This phrase is designed to check the list of extension exts installed in the browser. In our analysis, therefore, we aim to discover the ways in which the adversary might evade detection assuming exts is corrupt (i.e. contains an unapproved extension) when it is measured. Our initial query, therefore, will be (E, φ) in which E is the event system produced by the Copland semantics, and φ only holds for $\varphi(\mathsf{us.bmon}, \mathsf{exts})$.

We have implemented a preprocessing pipeline that converts a raw Copland phrase into a Chase formula representing its event system E. This phrase is a logical conjunction of facts that correspond to the existence of events (labeled as measurement events) and facts expressing the precedence order among them. The result when applied to Example 2 is the following.

```
l(e2) = msp(ks, av, us, bmon)
  ∧ l(e5) = msp(us, bmon, us, exts).
```

The non-consecutive event numbers reflect the presence of non-measurement events in the full Copland semantics described in Section 3. We represent φ explicitly with the following rule.

```
l(E) = msp(us, bmon, us, exts)
  ⇒ phi(us, exts, E).
```

We then submit to Chase a theory consisting of

- the formulas in Figs. 6 and 7,
- a single formula expressing the fact that any measurement event is one of the ones implied by the Copland semantics,
- the two formulas above representing the initial query, and
- a set of auxiliary formulas expressing lower level facts such as the injectivity of function symbols.

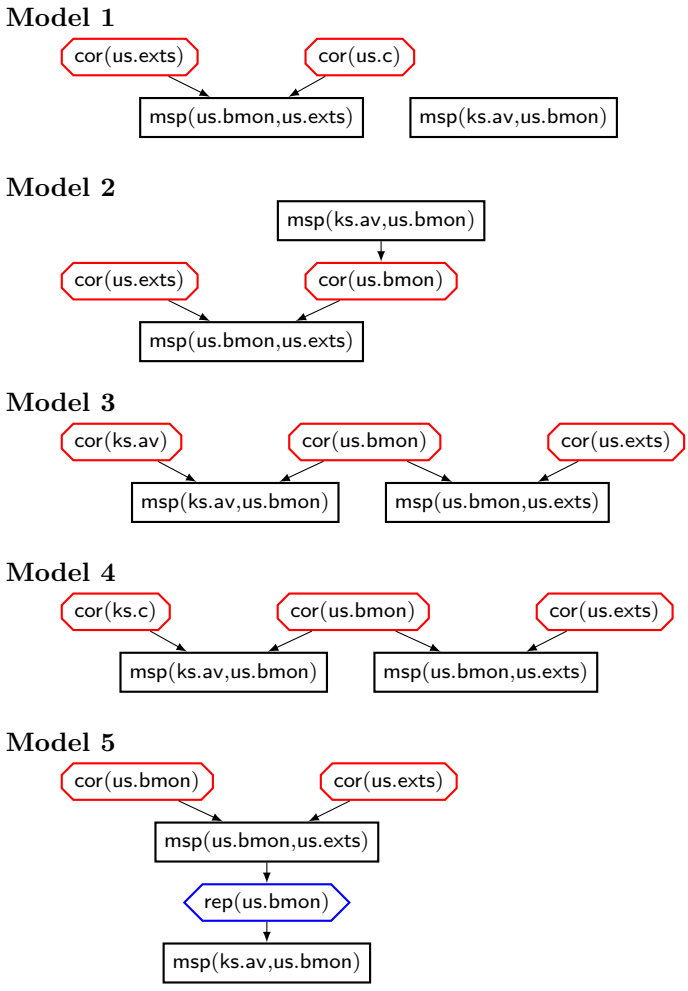

Figure 8: Executions of Example 2

The Chase performs its search as described in Section 5, and Theorem 18 ensures that models found by Chase are saturated queries compatible with the semantics of the given Copland phrase. Example 2 results in 5 different saturated queries. Fig. 8 depicts the executions corresponding to those saturated queries. The fact that the adversary can avoid detection is not surprising because there is always a way for an adversary to succeed by corrupting enough components or by corrupting components in the intervals between when they are measured and when they perform measurements. Indeed two of the models (Models 2 and 3) represent executions in which the adversary utilizes exactly those two strategies. In Model 3, ks.av, us.bmon, and us.exts are all corrupted before the start of the attestation. In Model 2, only us.bmon and us.exts are corrupted, but the former is corrupted *after* it is measured, but *before* it performs it measurement. Two other models (Models 1 and 4) are variants on these in which the measurement components themselves are not corrupted but Chase posits some other component on which they depend might be corrupted instead. The Chase cannot identify what such a component might be. It simply recognizes that the existence of such a dependency (which is not itself measured) could lead to the adversary avoiding detection. So, for example, us.av probably relies on the kernel to function correctly. If the kernel is corrupted with a root kit, this could cause the av not to discover a corruption of us.bmon. Finally, in the fifth model (Model 5), the adversary succeeds by corrupting both us.bmon and us.exts before the attestation begins. But in order to avoid detection at the measurement of us.bmon by ks.av the adversary relies on the possibility that that this measurement could happen *after* the other measurement, providing an opportunity to repair us.bmon after it performs its measurement, but before it gets measured itself.

These 5 models characterize where the remaining risk lies after performing the layered attestation specified by the Copland phrase. Put another way, they help identify assumptions that must hold in order for the attestation to guarantee detection of corruption. For example, the adversary must not corrupt ks.av or any component it depends on. Similarly, the adversary must not corrupt us.bmon after it has been measured. A relying party may have reason to believe that the target system is capable of ensuring these assumptions are met. For example, the relying party may trust that kernel level protections will protect ks.av from corruption, or that, say, address space layout randomization makes a successful runtime corruption of us.bmon extremely unlikely. Alternatively, the relying party may simply be willing to take these assumption on blind faith. If the attestation is meant to support an interaction that is not overly sensitive, the relying party might believe the adversary is capable of performing one of these actions, but also be willing to take such a risk.

One advantage of our analysis approach is that we can often explicitly express such assumptions as additional formulas input to the Chase. For example, we can write

```
depends(us, bmon, C) ⇒ false.
```

to express the assumption that the browser monitor does not depend in any way that matters on any other component. If we add this formula to our Chase theory we are essentially asking Chase not to show us any models in which it discovers such a dependency. Indeed, when we run Chase with this extra formula, it identifies 4 models instead of 5 because only one of the original 5 satisfied the antecedent of the formula (and hence was discarded).

Recent or deep corruptions have been identified as winning strategies for an adversary seeking to avoid detection given this adversary model [18]. However, as we saw above, there may be reasons to believe such corruptions are not likely. In either case, it is useful to discover if a Copland phrase *only* admits such strategies, or if there are other ways for the adversary to succeed. We can express the lack of deep corruptions directly by identifying the components which are considered "deep enough" to have strong protection and writing, for example:

$$\texttt{l(E) = cor(ks, av)} \Rightarrow \texttt{false.}$$

This excludes models in which ks.av is ever corrupted. We can similarly exclude *all* recent corruptions by disallowing corruption events that occur after some measurement event:

$$\texttt{prec(E, E2)} \wedge \texttt{l(E2) = cor(P,C)} \wedge \texttt{ms\_evt(E)} \Rightarrow \texttt{false.}$$

For the Copland phrase in Example 2, the addition of the two formulas above precluding deep or recent corruptions, together with assumptions that there are no unaccounted for dependency relationships excludes all but one execution. This is Model 5 in Fig. 8 in which a corrupted us.bmon performs its measurement, then repairs itself before getting measured. While this attack may require some skill (or possibly luck) in ensuring the two measurements happen in the required order, a stronger Copland phrase would guarantee the impossibility of this order of measurement. The Copland phrase from Example 3 specifies that ks.av must perform its measurement of us.bmon before the latter performs its measurement. This measurement is "bottom-up" in the sense that components higher up in the layered structure of the target system are measured later than the lower layers. Bottom-up measurement orderings are known to guarantee that any corruptions are recent or deep [18]. It is easy to verify this result for the particular case of Example 3 by submitting a query that excludes recent or deep corruptions. Indeed, Chase finds no executions for Example 3 when recent and deep corruptions are excluded.

With Copland phrases as simple as those from Examples 2 and 3, it seems feasible to perform an analysis by hand. However, as soon we begin to consider more complicated phrases involving numerous components with various dependencies and measurements that can be ordered in many ways, the analysis becomes much more difficult. Consider, for example, the following more complex version of the bank's attestation problem. Instead of using a special-purpose measurer to simply list installed browser extensions, the bank is willing

to accept a list generated by the browser's extension manager us.extmgr. However, the bank is interested in gaining trust in the extension manager which relies on part of the core browser code us.bser to function properly. Thus, the bank's special-purpose browser monitor us.bmon would now be responsible for measuring core parts of the browser code as well as the extension manager. For instance, it could hash elements of core functions needed to properly enumerate the list of extensions. The general-purpose antivirus software us.av would still be responsible for scanning for malware affecting us.bmon. For extra assurance, the bank will also request a runtime measurement of the operating system kernel. Tools such as LKIM [12] inspect the structure of the memory of a Linux kernel to detect violations of invariants that commonly occur when rootkits attempt to hide. WinKIM is a similar tool for the Windows kernel, and it could be run in a separate VM supported by Microsoft's Hyper-V virtualization technology. So it may request the kernel integrity monitor in Hyper-V hv.kim to measure the kernel ks.ker. For completeness, we imagine the target system has a way to measure us.av itself from another component hv.avm living in a Hyper-V VM. The following Copland phrase accomplishes the above attestation.

$$
\begin{aligned}
&*\mathsf{bank} : @_{\mathsf{hv}}[(\mathsf{kim\ ks\ ker} +\!\sim\!+ \mathsf{avm\ ks\ av}) \\
&\qquad +\!<\!+\ @_{\mathsf{ks}}[\mathsf{av\ us\ bmon} \\
&\qquad\qquad +\!<\!+\ @_{\mathsf{us}}[(\mathsf{bmon\ us\ extmgr} +\!\sim\!+ \mathsf{bmon\ us\ bser}) \\
&\qquad\qquad\qquad +\!<\!+\ \mathsf{extmgr\ us\ exts}]]]
\end{aligned}
$$

The value of automation provided by Chase for analyzing such a phrase becomes quickly apparent. In analyzing this phrase, we assume two dependency relations exist. Specifically, we assume both Depends(us.extmgr,us.bser) and Depends(ks.av,ks.ker).

The Copland semantics for this phrase is depicted in Fig. 9. If we submit a query to Chase in which we only stipulate that φ holds for us.exts when it is measured assuming only that there are no unaccounted for dependency relationships, we discover 40 distinct ways for the adversary to avoid detection. By submitting more constrained queries that make stronger assumptions, we can develop an understanding of what hoops an adversary is forced to jump through in those 40 possibilities. For example, if we additionally assume that neither of the components protected by Hyper-V are corrupted, there are only 24 possibilities. If, instead, we assumed only that no corruptions occur during the attestation (but allow components in Hyper-V to be corrupted) there are 12 possibilities. If we assume both that components in Hyper-V are not corrupted and that there are no other corruptions during the attestation, then the adversary cannot succeed.

While we believe constraining the adversary to performing recent or deep corruptions is often an acceptable risk, there may be circumstances where this assumption is unrealistic. Perhaps the signature file for the anti-virus is not integrity protected. Modification of this file would affect the outcome of the virus scan. We would want to enrich the Copland phrase above to include

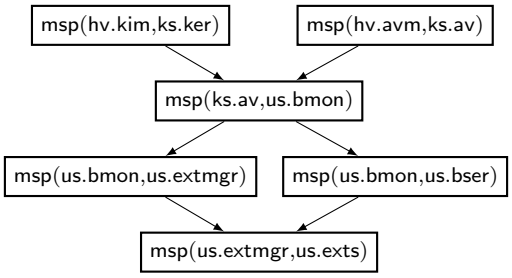

Figure 9: Copland semantics for complex attestation.

a measurement of us.sigfile. However, it could be very easy for malware to remove its own signature from this unprotected list during the attestation. In our analysis, therefore, we would preclude all corruptions during the attestation except for corruptions of the signature. We would write:

```
prec(E, E2) ∧ l(E2) = cor(P,C) ∧ ms_evt(E)
 ⇒ P = us ∧ C = sigfile.
```

In this way, Chase run with our layered attestation theory provides an interactive method for exploring the trust consequences of Copland phrases. There are often many ways to collect any given set of measurements. By running Chase on the Copland phrases representing the variety of measurement strategies, we can understand the relative strengths and weaknesses among them. The fewer assumptions that need to be made in order to guarantee successful detection of any corruptions, the stronger the Copland phrase. Since there will never be a single solution to fit all use cases, we believe this exploratory approach to analyzing the trustworthiness of layered attestation strategies is an essential capability in designing attestation systems & protocols and in selecting sets of Copland phrases suitable for given situations.

8 Related work

Too much has been written about remote attestation to perform a comprehensive literature review, so we contrast our work with the most relevant examples of related efforts in the literature.

One central tenet of the approach we take is that any fixed set of attestation solutions will be insufficient to accommodate the inevitable variety encountered in the type and architecture of target devices and the contexts in which the trust decisions they support will be made. This observation is not new. This was implicit in [3] which develops a set of principles that should guide the design and implementation of remote attestation systems. More recently, the IETF has formed a working group focused on Remote ATtestation procedureS

(RATS). Their draft architecture document [2] explicitly acknowledges the need for flexibility of mechanisms in implementing remote attestations. They also similarly envision the use of layered architectures to further enhance the quality of attestations. Similarly, Copland was designed to enable flexible specifications of complex attestations [17]. A complementary body of work seeks to develop flexible implementations that can be adapted to a wide range of scenarios [13, 15, 14]. The flexibility allowed for by both design and implementation leaves room for undesirable configurations that can yield untrustworthy results. We believe that our analysis methodology complements such work by offering a means for assessing the risk of any given option by understanding what an adversary must do to defeat it.

Of course, our work assumes that the target system has certain measurement capabilities and that they are reasonably effective at detecting corruption of subcomponents. Plenty has been written about particular approaches to measuring individual components ranging from flexible and programmable solutions [8, 20] to fixed solutions designed for particular architectures or components [12, 5, 11]. That line of research lies below the level of abstraction used in our execution model. However, it would be interesting to investigate whether the ideas we develop here could be suitably adapted to analyzing the measurement strategies for particular components. In particular, it may be possible to combine Copland [17] with the MSRR specification language for measurements [8] to obtain a top-to-bottom specification. By adapting the methods presented here, one would then have an accompanying top-to-bottom trust analysis as well.

Although many approaches take into account the possibility of an adversary interfering, few explicitly consider a runtime adversary that can interfere *during* an attestation. There are a few exceptions, most notably [18] which is the basis of the current work. The adversary model was enriched in [19] to account for manipulation of evidence in transit, a possibility not explored by the current work. Another analytical framework for analyzing attestations can be found in [4] in which they account for the effects of corruption on an attestation. The primary difference is that runtime corruptions are not considered. Additionally, the analysis is focused on Intel's "late launch" capability. None of these analytic frameworks have any automated support. To the best of our knowledge the current work is the first automated analysis methodology for remote attestation.

9 Conclusion

In this paper we introduced an approach to automated trust analysis of layered attestation protocols. We introduce a tool chain that ingests the specification of a layered attestation protocol written in the Copland language together with a configurable set of system assumptions, and produces a characterization of all the ways an active adversary might avoid detection by the attestation protocol. Our methodology presents opportunities for designers and implementers alike in developing attestation protocols and understanding how they help buy down risk when used in support of trust decisions.

Our approach is to compile a Copland specification into a partially ordered set of measurement events and apply our general-purpose model finder for geometric logic, called Chase, to characterize all the ways an active adversary can avoid detection under a given set of assumptions. This strategy requires us to equip Chase with a first-order logical theory that axiomatizes our execution model in the presence of an active adversary. We develop the theory of *saturated queries* and demonstrate that this theory correctly captures our intended denotation (Theorem 14). We present an axiomatization of the theory of saturated queries and prove that this axiomatization is correct (Theorem 18).

While the adversary model used in this paper is a useful one, it does not capture all the ways to interfere with the expected functioning of a layered attestation. We believe a fruitful line of future research would be to expand the adversary model to account for the ability to tamper with the integrity of evidence and to alter the control flow of an attestation. Such capabilities are reminiscent of the abilities of a network adversary assumed in the analysis of cryptographic protocols. A promising direction would be to determine how to integrate analyses using Chase with analyses using cryptographic protocol analysis tools. This would allow for a more complete understanding of the true benefits provided by well-designed layered attestations.

References

[1] Paul Barham, Boris Dragovic, Keir Fraser, Steven Hand, Tim Harris, Alex Ho, Rolf Neugebauer, Ian Pratt, and Andrew Warfield. Xen and the art of virtualization. *SIGOPS Oper. Syst. Rev.*, 37(5):164–177, October 2003.

[2] H. Birkholz, D. Thaler, M. Richardson, N. Smith, and W. Pan. Remote attestation procedures architecture, 2021. `https://datatracker.ietf.org/doc/draft-ietf-rats-architecture/` (Accessed 22-Feb-2021).

[3] George Coker, Joshua D. Guttman, Peter Loscocco, Amy L. Herzog, Jonathan K. Millen, Brian O'Hanlon, John D. Ramsdell, Ariel Segall, Justin Sheehy, and Brian T. Sniffen. Principles of remote attestation. *Int. J. Inf. Sec.*, 10(2):63–81, 2011.

[4] Anupam Datta, Jason Franklin, Deepak Garg, and Dilsun Kirli Kaynar. A logic of secure systems and its application to trusted computing. In *30th IEEE Symposium on Security and Privacy (S&P 2009), 17-20 May 2009, Oakland, California, USA*, pages 221–236, 2009.

[5] Lucas Davi, Ahmad-Reza Sadeghi, and Marcel Winandy. Dynamic integrity measurement and attestation: towards defense against return-oriented programming attacks. In *Proceedings of the 4th ACM Workshop on Scalable Trusted Computing, STC 2009, Chicago, Illinois, USA, November 13, 2009*, pages 49–54, 2009.

[6] Herbert B. Enderton. *A mathematical introduction to logic*. Academic Press, 2001.

[7] John Fisher and Marc Bezem. *Geolog and Skolem Machines*. California State Polytechnic and University of Bergen, 2007. `https://www.cpp.edu/~jrfisher/www/prolog_tutorial/logic_topics/geolog/index.html`.

[8] Jason Gevargizian and Prasad Kulkarni. MSRR: measurement framework for remote attestation. In *2018 IEEE 16th Intl Conf on Dependable, Autonomic and Secure Computing, 16th Intl Conf on Pervasive Intelligence and Computing, 4th Intl Conf on Big Data Intelligence and Computing and Cyber Science and Technology Congress, DASC/PiCom/DataCom/CyberSciTech 2018, Athens, Greece, August 12-15, 2018*, pages 748–753. IEEE Computer Society, 2018.

[9] Trusted Computing Group. TPM Main Specification Level 2 version 1.2, Parts 1–3, Revision 116, 2011. `https://trustedcomputinggroup.org/resource/tpm-main-specification/`.

[10] Intel. Intel® Software Guard Extensions (Intel® SGX), 2016. `https://software.intel.com/en-us/sgx`.

[11] Chongkyung Kil, Emre Can Sezer, Ahmed M. Azab, Peng Ning, and Xiaolan Zhang. Remote attestation to dynamic system properties: Towards providing complete system integrity evidence. In *Proceedings of the 2009 IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2009, Estoril, Lisbon, Portugal, June 29 - July 2, 2009*, pages 115–124, 2009.

[12] Peter Loscocco, Perry W. Wilson, J. Aaron Pendergrass, and C. Durward McDonell. Linux kernel integrity measurement using contextual inspection. In *Proceedings of the 2nd ACM Workshop on Scalable Trusted Computing, STC 2007, Alexandria, VA, USA, November 2, 2007*, pages 21–29, 2007.

[13] Aaron Pendergrass, Sarah Helble, John Clemens, and Peter Loscocco. A platform service for remote integrity measurement and attestation. In *Military Communications Conference (MILCOM) 2018*, October 2018.

[14] Adam Petz. An infrastructure for faithful execution of remote attestation protocols. In Perry Alexander, Drew Davidson, and Baek-Young Choi, editors, *Proceedings of the 7th Annual Symposium on Hot Topics in the Science of Security, HotSoS 2020, Lawrence, Kansas, USA, September 22-24, 2020*, page 17:1. ACM, 2020.

[15] Adam Petz and Perry Alexander. A copland attestation manager. In Xenofon D. Koutsoukos, Alvaro A. Cárdenas, and Ehab Al-Shaer, editors, *Proceedings of the 6th Annual Symposium on Hot Topics in the Science of Security, HotSoS 2019, Nashville, TN, USA, April 1-3, 2019*, pages 6:1–6:10. ACM, 2019.

[16] John D. Ramsdell. *Chase Source Repository*. The MITRE Corporation, 2019. `https://github.com/ramsdell/chase`, install with `opam install chase`.

[17] John D. Ramsdell, Paul D. Rowe, Perry Alexander, Sarah C. Helble, Peter Loscocco, J. Aaron Pendergrass, and Adam Petz. Orchestrating layered attestation. *LNCS*, 11426:197–221, 2019.

[18] Paul D. Rowe. Confining adversary actions via measurement. *Third International Workshop on Graphical Models for Security*, pages 150–166, 2016.

[19] Paul D. Rowe. Bundling evidence for layered attestation. In *Trust and Trustworthy Computing - 9th International Conference, TRUST 2016*, In press.

[20] Mark Thober, J. Aaron Pendergrass, and Andrew D. Jurik. JMF: Java measurement framework. In *Proceedings of the Seventh ACM Workshop on Scalable Trusted Computing*, pages 21–32. ACM, 2012.

[21] Anthony Velte and Toby Velte. *Microsoft virtualization with Hyper-V*. McGraw-Hill, Inc., 2009.