

Access Control By Tracking Shallow Execution History

Philip W. L. Fong
Department of Computer Science
University of Regina
Regina, Saskatchewan, Canada S4S 0A2
pwlffong@cs.uregina.ca

Abstract

Software execution environments like operating systems, mobile code platforms and scriptable applications must protect themselves against potential damages caused by malicious code. Monitoring the execution history of the latter provides an effective means for controlling the access pattern of system services. Several authors have recently proposed increasingly general automata models for characterizing various classes of security policies enforceable by execution monitoring. An open question raised by Bauer, Ligatti and Walker is whether one can further classify the space of security policies by constraining the capabilities of the execution monitor. This paper presents a novel information-based approach to address the research problem. Specifically, security policies are characterized by the information consumed by an enforcing execution monitor.

By restricting the execution monitor to track only a shallow history of previously granted access events, a precise characterization of a class of security policies enforceable by restricted access to information is identified. Although provably less expressive than the general class of policies enforceable by execution monitoring, this class does contain naturally occurring policies including Chinese Wall policy, low-water-mark policy, one-out-of- k authorization, assured pipelines, etc. Encouraged by this success, the technique is generalized to produce a lattice of policy classes. Within the lattice, policy classes are ordered by the information required for enforcing member policies. Such a fine-grained policy classification lays the semantic foundation for future studies on special-purpose policy languages.

1. Introduction

Software execution environments like operating systems, mobile code platforms and scriptable applications must protect themselves against potential damages caused by malicious code. Monitoring the execution history of the latter

provides an effective means for controlling the access pattern of system services. Execution monitoring (EM) can be implemented either by interposing a reference monitor between system service entry points and the code providing the services [15, 4, 12], or by injecting monitoring code into client programs at load time [13, 27, 28, 30, 31, 23, 24]. Schneider [26] proposed an automata-theoretic characterization of security policies enforceable by EM. Specifically, an EM-enforceable policy prescribes access event sequences recognized by a Büchi automaton [2]. It is observed that Büchi-like security automata can only enforce safety properties, but not liveness properties. Subsequently, Bauer, Ligatti and Walker [3, 18] proposed a characterization of increasingly general classes of security policies enforceable by insertion, suppression and editing automata. These policy classes are provably more expressive than EM-enforceable policies.

An open question raised by the work of Bauer *et al* is whether or not one can further classify the space of EM-enforceable policies by *constraining* the capability of the execution monitor. Not only does such a fine-grained classification help us understand the inherent complexity of security policies, it also has a number of practical engineering ramifications. In an environment in which users invest a high degree of trust on the formulation of security policies, and in which the complexity of the security policies increases with the complexity of the software environment, one has to face the reality that policy engineering shares many challenges once considered unique to software engineering. Characterizing security policies as members of well-understood policy classes can facilitate policy engineering in the following ways:

- Special-purpose policy languages can be designed for a policy class to facilitate the correct formulation of member policies.
- Efficient decision procedures may exist for verifying the correctness of policies belonging to policy classes with rich internal structure.

- Some policy classes may exhibit structural properties that render their member policies decomposable into more manageable policy components. Discovery of such structural properties enables the composition of complex policies from reusable components.

This paper presents a novel information-based approach to address the research problem raised by Bauer *et al.* Instead of following their proposal, and classifying security policies by constraining *computational resources* available to the execution monitor, this work classifies EM-enforceable security policies by the kind of *information* that needs to be tracked by the execution monitor. Such a fine-grained policy classification lays the semantic foundation for future studies on special-purpose policy languages.

Consider the Chinese Wall policy [7, 19, 25], a commercial policy for preventing accesses leading to conflict of interests. As observed by Brewer and Nash in their original formulation of the policy, successful enforcement of the Chinese Wall policy only requires the maintenance of a *shallow access history* of previously granted access events. Specifically, the decision on whether an access is to be granted is based solely on the *set*¹ of access events that have already been granted, and not on the actual *sequencing* of such access events. In contrast, a Büchi automaton could potentially have the full history of granted access events at its disposal when an access granting decision is made. This paper presents a characterization of security policies enforceable by tracking only the shallow access history of a system. Security policies in this class are recognizable by *shallow history automata* (SHA), expressiveness of which is provably more restrictive than that of Büchi-like security automata. Surprisingly, it is still possible to express a wide range of well-known and realistic security policies with SHA: Chinese Wall policy [7], low-water-mark policy [5], one-out-of- k authorization [11], assured pipelines [6, 32], etc. This demonstrates the feasibility of defining meaningful policy classes by constraining information accessible to execution monitors.

Motivated by the above success, the state abstraction techniques applied to characterize shallow access history tracking is generalized. A lattice of security policy classes is obtained as a result. At the top of the lattice is the class of policies enforceable by tracking the full history of access events, as in the case of Büchi-like security automata. As one moves down the lattice, one finds classes of policies that are enforceable by consuming less and less information, with the class of policies enforceable by SHA and that by stack inspection [30, 14] somewhere in the middle, incomparable to each other, and the class of memoryless policies at the bottom. This work has therefore laid the theoret-

ical groundwork for studying special-purpose subclasses of EM-enforceable security policies.

This paper is organized as follows. Related works are reviewed in Section 2. SHA are defined in Section 3 to provide an information-based characterization of security policies enforceable by tracking shallow access history. A number of naturally occurring security policies are shown to be SHA-enforceable in Section 4. The SHA configuration is generalized in Section 5 to yield a lattice of policy classes. Section 6 applies the lattice framework to study the relative expressiveness of some naturally occurring policy classes, including that of stack inspection. Discussions can be found in Section 7. Section 8 concludes the paper.

2. Related Works

Schneider [26] pioneered the characterization of security policies enforceable by execution monitoring (EM). Specifically, an EM-enforceable policy prescribes access event sequences recognized by a Büchi automaton [2]. It is observed that Büchi-like security automata can only enforce safety properties, but not liveness properties. Viswanathan [29] points out that any reasonable characterization on execution monitoring must involve a computability constraint. Subsequently, Bauer, Ligatti and Walker [3, 18] proposed a characterization of increasingly general classes of security policies enforceable by insertion, suppression and editing automata, while Hamlen, Morrisett and Schneider [16] offers a characterization of security policies enforceable by code rewriting. These policy classes are provably more expressive than EM-enforceable policies. The work reported in this paper is the first one to provide a fine-grained, information-based characterization of subclasses of EM-enforceable policies.

3. Access Control By Shallow History Tracking

This section introduces shallow history automata, the definition of which provides an information-based characterization of security policies enforceable by tracking shallow access history. The class of security policies expressible by such automata is proven to be a proper subset of the general class of EM-enforceable security policies, thereby confirming the claim that subclasses of EM-enforceable policies can be defined through the restriction of information accessible to the execution monitor. To fix thoughts, the notion of EM-enforceable policies and its characterization via security automata are reviewed in Section 3.1 and 3.2 respectively. Shallow history automata are then discussed in Section 3.3.

¹ Brewer and Nash used a history matrix to track this set.

3.1. EM-Enforceable Security Policies

Let Σ be a finite or countably infinite set of *access events*. A *policy* is a set $P \subseteq \Sigma^*$ of finite sequences of access events. An *EM-enforceable policy*² is a non-empty *prefix-closed* policy. A policy P is prefix-closed if it satisfies the following condition:

$$\forall u \in \Sigma^* : u \notin P \Rightarrow (\forall v \in \Sigma^* : uv \notin P)$$

Let $\text{prefix}(w)$ be the set of all prefixes of w , including ϵ and w itself. That is, $\text{prefix}(w) = \{u \in \Sigma^* \mid \exists v \in \Sigma^* \text{ s.t. } uv = w\}$. It is easy to see that the following is an equivalent characterization of prefix-closed policies:

$$\forall w \in \Sigma^* : w \in P \Rightarrow \text{prefix}(w) \subseteq P \quad (1)$$

In the following, we consider only EM-enforceable policies.

3.2. Security Automata

A variant of Büchi automata is defined here. A *security automaton (SA)* is a quadruple $\langle \Sigma, Q, q_0, \delta \rangle$, where

- Σ is a finite or countably infinite set of access events,
- Q is a finite or countably infinite set of *automaton states*,
- $q_0 \in Q$ is an *initial state*,
- $\delta : Q \times \Sigma \rightarrow Q$ is a (possibly partial) *transition function*³.

The notion of acceptance for a SA is different from that for a regular finite state machine, in which a final state is explicitly identified. An access event sequence is accepted by a SA if a transition is defined for every event in the sequence. The notion is formalized as follows. Given a SA $M = \langle \Sigma, Q, q_0, \delta \rangle$, the following notations are defined for $q, q' \in Q, a \in \Sigma$ and $w \in \Sigma^*$:

$$\begin{aligned} q &\xrightarrow{a}_M q' && \text{if } \delta(q, a) = q' \\ q &\xrightarrow{\epsilon}_M q && \\ q &\xrightarrow{wa}_M q' && \text{if } \exists q'' \in Q . q \xrightarrow{w}_M q'' \wedge q'' \xrightarrow{a}_M q' \end{aligned}$$

2 For the purpose of this work, the definition of an EM-enforceable policy as adopted here is different from the one used by Schneider [26] in the following ways: (1) Schneider differentiates between general security policies and a special class of policies that he calls security properties. Only those policies that are properties are considered in this paper. (2) Schneider considers infinite sequences of access events. Following the practice of Bauer *et al* [3], only finite sequences of access events are considered in this paper.

3 The formulation of security automata as given here differs from that of Bauer *et al*: As our focus is information rather than resource constraints, no tractability restriction is imposed on the transition function. See, however, [29] for the need of such constraints in the general cases.

We say that M accepts an access event sequence w if $q_0 \xrightarrow{w}_M q$ for some $q \in Q$. The policy $\mathcal{P}(M)$ *recognized* by the SA M is then defined as the set of all sequences accepted by M :

$$\{w \in \Sigma^* \mid \exists q \in Q : q_0 \xrightarrow{w}_M q\}.$$

It is easy to see that such a set is always non-empty and prefix-closed, that is, $\mathcal{P}(M)$ contains ϵ and satisfies condition (1). Conversely, given any non-empty prefix-closed policy P , there is a SA M so that $P = \mathcal{P}(M)$. To see this, consider the SA $\langle \Sigma, \Sigma^*, \epsilon, \delta_P \rangle$, where $\delta_P(w, a)$ is defined to be wa if $w, wa \in P$, and is otherwise undefined. Such a SA recognizes P . Consequently, the class of EM-enforceable policies coincides with the class of policies recognized by a SA. We call the above SA constructed to recognize P the *canonical SA* for policy P , and denote it by $\text{SA}(P)$.

Intuitively, the state of a SA represents the information that is tracked by the corresponding execution monitor. It represents the internal data structure maintained by the execution monitor across subsequent access granting decisions. The image of the transition function captures the updating procedure of the internal data structure, while domain of the transition function captures the logic of access granting decisions. Notice that the canonical SA tracks the full history of previously granted access events.

3.3. Shallow History Automata

Let $\mathcal{F}(S)$ be the set of all *finite* subsets of a set S . A *shallow access history* (or simply *shallow history*) is a finite subset of Σ , that is, a member of $\mathcal{F}(\Sigma)$. Our goal is to define a class of automata that track only the shallow history of previously granted access events.

A *shallow history automaton (SHA)* is a SA of the form $\langle \Sigma, \mathcal{F}(\Sigma), H_0, \delta \rangle$, where

- Σ is a finite or countably infinite set of access events,
- The state set $\mathcal{F}(\Sigma)$ contains all possible shallow access histories.
- $H_0 \in \mathcal{F}(\Sigma)$ is an *initial access history*, and
- The transition function δ is such that $\delta(H, a) = H \cup \{a\}$ if δ is defined at $\langle H, a \rangle$.

Intuitively, a SHA tracks only a shallow access history, and bases its access granting decisions solely on this information. Sequencing information about previously granted access events are not retained for subsequent access granting decisions.

Notice that the image of δ is always $H \cup \{a\}$ if it is defined at $\langle H, a \rangle$. Consequently, defining δ amounts to specifying its domain as a subset of $\mathcal{F}(\Sigma) \times \Sigma$. That is, a SHA transition function is uniquely specified by listing all the points at which it is defined.

A policy recognized by some SHA is said to be *SHA-enforceable*. We write EM_{SHA} to denote the class of all SHA-enforceable policies.

As expected, SHA is strictly less expressive than SA:

Theorem 1 *Fixing the set Σ of access events, there is a SA M so that no SHA N is such that $\mathcal{P}(M) = \mathcal{P}(N)$.*

Proof: Let $\Sigma = \{a, b, c, d\}$. Consider the policy $P = \text{prefix}(abcd) \cup \text{prefix}(badc)$. The policy P is non-empty and prefix-closed by construction, and is thus recognizable by its canonical SA. Suppose that P is recognized by a SHA M . Let H_0 be the initial state of M . The following transitions are valid:

$$\begin{aligned} H_0 &\xrightarrow{a}_M \{a\} \cup H_0 \xrightarrow{b}_M \{a, b\} \cup H_0 \\ &\xrightarrow{c}_M \{a, b, c\} \cup H_0 \xrightarrow{d}_M \{a, b, c, d\} \cup H_0 \\ H_0 &\xrightarrow{b}_M \{b\} \cup H_0 \xrightarrow{a}_M \{a, b\} \cup H_0 \\ &\xrightarrow{d}_M \{a, b, d\} \cup H_0 \xrightarrow{c}_M \{a, b, c, d\} \cup H_0 \end{aligned}$$

However, with the above transitions, M also accepts $abdc$ and $bacd$:

$$\begin{aligned} H_0 &\xrightarrow{a}_M \{a\} \cup H_0 \xrightarrow{b}_M \{a, b\} \cup H_0 \\ &\xrightarrow{d}_M \{a, b, d\} \cup H_0 \xrightarrow{c}_M \{a, b, c, d\} \cup H_0 \\ H_0 &\xrightarrow{b}_M \{b\} \cup H_0 \xrightarrow{a}_M \{a, b\} \cup H_0 \\ &\xrightarrow{c}_M \{a, b, c\} \cup H_0 \xrightarrow{d}_M \{a, b, c, d\} \cup H_0 \end{aligned}$$

By way of contradiction, P is not SHA-enforceable.

Note that the use of a non-empty initial history H_0 in the specification of a SHA is merely an optional convenience, as the following proposition entails.

Proposition 2 *For every SHA $M = \langle \Sigma, \mathcal{F}(\Sigma), H_0, \delta \rangle$, there exists a SHA $M_0 = \langle \Sigma, \mathcal{F}(\Sigma), \emptyset, \delta_0 \rangle$ such that $\mathcal{P}(M) = \mathcal{P}(M_0)$.*

Proof: To establish the statement, it suffices to construct δ_0 so that for all $H \in \mathcal{F}(\Sigma)$ and $a \in \Sigma$, $\delta_0(H, a)$ is defined if $\delta(H \cup H_0, a)$ is defined.

We call a SHA with \emptyset as the initial state a *normalized SHA*. Proposition 2 shows that every policy recognizable by a SHA is recognizable by a normalized SHA.

3.4. Summary

A subclass of security automata, SHA, is successfully defined to capture the notion of execution monitoring by tracking only shallow access history. The separation result in Theorem 1 confirms that SHA are strictly less expressive than general SA.

4. Security Policies Enforceable By Tracking Shallow History

Defining artificial subclasses of security policies is pointless if the classes do not correspond to security policies found in real-life applications. This section demonstrates that the class of SHA-enforceable policies does include nontrivial security policies such as the Chinese Wall policy (Section 4.1), the low-water-mark policy (Section 4.2), one-out-of- k authorization (Section 4.3), and assured pipelines (Section 4.4). The goal is to show that the definition of SHA yields a class of naturally occurring security policies.

4.1. Chinese Wall Policy

Set in a commercial context, in which a consultant shall not advise clients whom he or she has insider knowledge of a competitor, the Chinese Wall security policy is designed to avoid any conflict of interest that may arise due to the unchecked flow of information across datasets belonging to competing parties. Let O be a set of data objects, S a set of subjects, G a set of company datasets, and T a set of conflict of interest classes. Associated with each data object $o \in O$ is a permanent label $group[o] \in G$ describing the company dataset in which o belongs. Similarly, a permanent label $type[g] \in T$ is assigned to each company dataset $g \in G$; the label describes the conflict of interest class in which a company dataset belongs. A subject s may access a data object o only if one of the following holds:

- Subject s has already accessed another object o' belonging to the same company dataset of o , that is, $group[o] = group[o']$.
- Every object o' that subject s has accessed so far belongs to a company dataset whose conflict of interest class is different from that of the company dataset in which o belongs, that is, $type[group[o]] \neq type[group[o']]$.

A reference monitor enforcing the Chinese Wall policy may do so by keeping track of the set of objects previously accessed by each subject. The policy is therefore SHA-enforceable. This can be demonstrated formally by the following construction. Let the set of access events be $\Sigma = S \times O$, so that a pair $\langle s, o \rangle$ refers to the event of subject s accessing object o . Define SHA $N = \langle \Sigma, \mathcal{F}(\Sigma), \emptyset, \delta_{group, type} \rangle$, where the SHA transition function $\delta_{group, type}$ is defined at $\langle H, \langle s, o \rangle \rangle$ whenever the following holds:

- there exists $\langle s, o' \rangle \in H$ s.t. $group[o] = group[o']$, or
- for all $\langle s, o' \rangle \in H$, $type[group[o]] \neq type[group[o']]$

By construction, the SHA N enforces the Chinese Wall policy.

With refinement to the above construction, it is also possible to show that the variation of Chinese Wall policy as proposed by Lin [19] is SHA-enforceable.

4.2. Low-Water-Mark Policy (for Subjects)

The low-water-mark policy is one of the three lattice-based integrity policies proposed by Biba [5]. Defined in Biba's security model are a set S of subjects, a set O of objects and a set L of integrity levels, which are partially ordered by a binary relation \leq . At any point of time, a label $l[s] \in L$ is assigned to every subject $s \in S$, and likewise $l[o] \in L$ is assigned to every object $o \in O$. Intuitively, the label $l[\cdot]$ describes the *trustworthiness* of a subject or an object. Three kinds of access events are defined: **write**(s, o), **read**(s, o) and **exec**(s, s'). The low-water-mark policy grants accesses according to the following rules:

1. **read**(s, o) is always permitted, with the side effect of $l[s] \leftarrow l[s] \wedge l[o]$, where \wedge denotes the greatest lower bound between integrity levels.
2. **write**(s, o) is permitted if $l[o] \leq l[s]$.
3. **exec**(s, s') is permitted if $l[s'] \leq l[s]$.

Intuitively, a subject is not permitted to write to objects with integrity levels higher than itself, whether directly (Rule 2) or indirectly (Rule 3). In addition, reading low integrity objects downgrades the integrity level of a subject (Rule 1).

To construct a SHA for enforcing the low-water-mark policy, notice that the label of a subject is a function of the objects it has read, while that of an object is permanent. Let Σ be the set of all access events. Given an initial label assignment $l[\cdot]$, assume, without loss of generality, that for every $l \in L$, there is at least one $o \in O$ such that $l[o] = l$ (this can be easily achieved by introducing auxiliary objects into O). Define a SHA $N_l = \langle \Sigma, \mathcal{F}(\Sigma), H_l, \delta_l \rangle$, where the initial history H_l is further specified as follows:

$$H_l = \{\mathbf{read}(s, o) \mid s \in S, o \in O, l[s] = l[o]\},$$

and the SHA transition function δ_l is defined at exactly the following points:

- $\langle H, \mathbf{read}(s, o) \rangle$ for all $H \in \mathcal{F}(\Sigma)$, $s \in S$ and $o \in O$, and
- $\langle H, \mathbf{write}(s, o) \rangle$ for all $H \in \mathcal{F}(\Sigma)$, $s \in S$ and $o \in O$ such that for all $\mathbf{read}(s, o') \in H$, $l[o] \leq l[o']$, and
- $\langle H, \mathbf{exec}(s, s') \rangle$ for all $H \in \mathcal{F}(\Sigma)$, $s, s' \in S$ such that for all $\mathbf{read}(s, o) \in H$, there is a $\mathbf{read}(s', o') \in H$, $l[o'] \leq l[o]$.

By construction, the SHA N_l enforces the low-water-mark policy.

4.3. One-Out-Of- k Authorization

One-out-of- k authorization [11] classifies applications into equivalence classes based on the kind of access rights required to complete tasks. For example, the following application classification is given in [11]:

- A **browser** is a program that connects to remote sites (**network-connection**), access per-process temporary files (**access-tmp-files**), and displays them to the user (**console-io**).
- An **editor** is a program that access user files (**access-user-files**), access per-process temporary files (**access-tmp-files**), and interacts with the user (**console-io**).
- A **shell** is a program that interacts with the user (**console-io**) and creates subprocesses (**create-subprocess**).

The goal is to dynamically classify an executing program into one of the application classes based on the access requests it makes. Once classified into an application class, a program is only allowed to exercise access rights granted to the class. For instance, once an application has opened a network socket, it is classified as a browser, and will subsequently not be allowed to read user files.

The one-out-of- k constraint can be easily enforced by a SHA. Let Σ be the set of all possible accesses made by an application. An application class i is completely characterized by the finite or countably infinite set $C_i \subseteq \Sigma$ of permitted accesses. Let the set $\mathcal{C} = \{C_1, C_2, \dots, C_k\}$ be the set of all application classes. Define SHA $N_{\mathcal{C}} = \langle \Sigma, \mathcal{F}(\Sigma), \emptyset, \delta_{\mathcal{C}} \rangle$ s.t. the SHA transition function $\delta_{\mathcal{C}}$ is defined at $\langle H, a \rangle$ iff $H \cup \{a\} \subseteq C_i$ for some $C_i \in \mathcal{C}$. By construction, $N_{\mathcal{C}}$ enforces one-out-of- k authorization.

4.4. Assured Pipelines

Assured pipelines [6, 32] arise in the context of ensuring data integrity when data objects are processed by (usually linear) pipelines of transformation procedures. Let O be a set of data objects. Let S be a set of transformation procedures. Assume that S contains a distinguished member **create**. Define the set of access events to be $\Sigma = S \times O$, where the member $\langle s, o \rangle$ denotes the application of transformation procedure s to data object o . An assured pipeline policy is specified as an *enabling relation* $e \subseteq S \times S$, with the following restrictions:

1. No circularity: the binary relation defines a directed acyclic graph (DAG).
2. No pair of the form $\langle s, \mathbf{create} \rangle$ may be included: **create** is the sole source node of the acyclic graph.

The intention is that $\langle s, s' \rangle \in e$ means that access $\langle s', o \rangle$ is granted if $\langle s, o \rangle$ is the last access on object o . In other

words, application of transformation procedure s on an object is a sufficient condition for the subsequent application of transformation procedure s' on the same object. Because of the “no circularity” clause, each event $\langle s, o \rangle$ may occur at most once.

An assured pipeline policy can be enforced by a SHA $N_e = \langle \Sigma, \mathcal{F}(\Sigma), \emptyset, \delta_e \rangle$, where the SHA transition function δ_e is defined at exactly the points $\langle H, \langle s, o \rangle \rangle$ satisfying one of the following conditions:

- $s = \text{create}$ and $\langle s, o \rangle \notin H$, or
- for some $s' \in S$, all of the following hold:
 - $\langle s', o \rangle \in H$, and
 - $\langle s', s \rangle \in e$, and
 - there is no $s'' \in S$ s.t. $\langle s'', o \rangle \in H$ and $\langle s', s'' \rangle \in e$.

By construction N_e enforces the assured pipeline policy specified by enabling relation e .

4.5. Summary

Four naturally occurring security policies have been shown to be SHA-enforceable, thereby demonstrating that information restriction could indeed be employed to classify real-life security policies.

5. Obtaining Policy Classes By Abstraction

Given that it is possible to define a meaningful subclass of EM-enforceable security policies, the next question is: can the above technique be generalized to obtain other subclasses of EM-enforceable policies? Section 5.1 offers an affirmative answer to the question, with the help of concepts borrowed from automata theory [8, 17]. The structural properties of EM-enforceable policy classes are then studied in Section 5.2 and 5.3.

5.1. Abstraction By Homomorphism

Given an arbitrary restriction on the information accessible to an execution monitor, the following procedure can be systematically carried out to obtain an automata-theoretic characterization of the policies thus enforceable:

1. An information restriction constraint is specified as a set \mathcal{A} of *abstract states*, which represents the kind of information that the execution monitor is allowed to track (e.g., shallow histories as finite subsets of access events).
2. An interpretation of the abstract states in \mathcal{A} is defined, so they refer to states of the canonical SA in a consistent manner (e.g., each shallow history documents

the set of previously granted access events in an execution sequence).

3. Under the above interpretation, a subclass of SA is defined so that member automata behave consistently according to the interpretation (e.g., shallow history automata).

This plan is executed as follows.

Let \mathcal{A} be a finite or countably infinite set of abstract states. A function $\alpha : \Sigma^* \rightarrow \mathcal{A}$ is an *abstraction* if it satisfies the following *compatibility property*:

$$\text{For all } w, w' \in \Sigma^* \text{ and } a \in \Sigma, \\ \alpha(w) = \alpha(w') \Rightarrow \alpha(wa) = \alpha(w'a) \quad (2)$$

An abstraction mashes distinct states of the canonical SA into a single abstract state, rendering them indistinguishable. The compatibility property guarantees that the loss of information does not introduce confusion.

The *homomorphic image* $\text{SA}_\alpha(P)$ of the canonical security automaton $\text{SA}(P)$ induced by abstraction α is the security automaton $\langle \Sigma, \mathcal{A}, \alpha(\epsilon), \delta_{P/\alpha} \rangle$, where

$$\delta_{P/\alpha}(\alpha(w), a) = \alpha(wa) \quad \text{whenever} \quad \delta_P(w, a) = wa$$

and δ_P is the transition function of $\text{SA}(P)$. Since α satisfies the compatibility property, $\delta_{P/\alpha}$ is a well-defined (partial) function.

Notice that not every policy P is recognized by $\text{SA}_\alpha(P)$ — in general, $P \subseteq \mathcal{P}(\text{SA}_\alpha(P))$ (see Corollary 16). Those policies P recognized by $\text{SA}_\alpha(P)$ are the policies that can be enforced by consuming only information left behind by the abstraction α . A policy P is said to be *enforceable* by abstraction α iff P is recognized by the security automaton $\text{SA}_\alpha(P)$, that is, iff $P = \mathcal{P}(\text{SA}_\alpha(P))$. Fixing the set Σ of access events, the class of all policies enforceable by abstraction α is denoted by EM_α . In summary, given any abstraction α , it is now possible to define exactly the class of security policies enforceable by tracking information permitted by the abstraction.

The policy class EM_α is defined as the set of all solutions to a fixed point equation. Although it facilitates the development of theory (see the appendix), such a definition is not easy to apply. This motivates an alternative characterization of EM_α . Given an abstraction function α , an α -SA is a SA $\langle \Sigma, \mathcal{A}, \alpha(\epsilon), \delta \rangle$ such that for all $w \in \Sigma^*$ and $a \in \Sigma$, either $\delta(\alpha(w), a) = \alpha(wa)$ or δ is not defined at $\langle \alpha(w), a \rangle$ at all. Because α satisfies the compatibility property, the transition function δ is well defined. An α -SA has the following useful property.

Proposition 3 *Given an α -SA M and $w \in \Sigma^*$, $w \in \mathcal{P}(M)$ iff $\alpha(\epsilon) \xrightarrow{w}_M \alpha(w)$.*

The following proposition asserts that α -SA provides an alternative characterization of policies enforceable by an abstraction α .

Proposition 4 *EM $_{\alpha}$ contains precisely the class of policies recognizable by α -SA.*

Proof: Notice that every homomorphic image $\text{SA}_{\alpha}(\cdot)$ induced by α is an α -SA, and thus every policy enforceable by abstraction α is recognizable by an α -SA. The converse is also true: given an α -SA M , $\mathcal{P}(M) = \mathcal{P}(\text{SA}_{\alpha}(\mathcal{P}(M)))$. To demonstrate this, it suffices to show that

$$\mathcal{P}(\text{SA}_{\alpha}(\mathcal{P}(M))) \subseteq \mathcal{P}(M),$$

as Corollary 16 supplies the other direction of containment. The above statement can be easily established by induction with the help of the following lemma.

Lemma: Let M be an α -SA $\langle \Sigma, \mathcal{A}, \alpha(\epsilon), \delta \rangle$. For all $q, q' \in \mathcal{A}$ and $a \in \Sigma$, if $q \xrightarrow{a}_{\text{SA}_{\alpha}(\mathcal{P}(M))} q'$ then $q \xrightarrow{a}_M q'$.

Proof: Suppose the antecedent holds, then by definition of $\text{SA}_{\alpha}(\mathcal{P}(M))$ there exists $w \in \Sigma^*$ so that $q = \alpha(w)$, $q' = \alpha(wa)$ and $w \xrightarrow{a}_{\text{SA}(\mathcal{P}(M))} wa$, meaning that both w and wa are in $\mathcal{P}(M)$. By Proposition 3, we have $\alpha(\epsilon) \xrightarrow{w}_M q \xrightarrow{a}_M q'$.

The following example shows that access control by tracking shallow history is simply an instance of abstraction.

Example 5 *There exists an abstraction function α such that $EM_{\alpha} = EM_{SHA}$. To see this, define $\alpha : \Sigma^* \rightarrow \mathcal{F}(\Sigma)$ such that for $w \in \Sigma^*$, $\alpha(w)$ is the set of access events in Σ that occur in w . The function α satisfies the compatibility property: if $w, w' \in \Sigma^*$ are such that $\alpha(w) = \alpha(w')$, then w and w' contain the same set $H \in \mathcal{F}(\Sigma)$ of access events, and thus $\alpha(wa) = \alpha(w'a) = H \cup \{a\}$ for any $a \in \Sigma$. Consequently, α is an abstraction function. Under this definition of α , every normalized SHA is an α -SA and vice versa. The claim follows from Propositions 2 and 4.*

5.2. Abstractions As Congruence Relations

It turns out that the notion of information abstraction is very robust: it can be defined independent of the choice of abstract states.

Every abstraction α induces an equivalence relation \equiv_{α} as follows:

$$\text{For all } w, w' \in \Sigma^*, w \equiv_{\alpha} w' \Leftrightarrow \alpha(w) = \alpha(w').$$

It can be shown that such an equivalence relation satisfies the following *substitution property*:

$$\text{For all } w, w' \in \Sigma^* \text{ and } a \in \Sigma, \\ w \equiv w' \Rightarrow wa \equiv w'a \quad (3)$$

Conversely, given any equivalence relation \equiv over Σ^* that satisfies the substitution property, it can be shown that the mapping $\alpha_{\equiv} : \Sigma^* \rightarrow \Sigma^*/\equiv$ defined as follows:

$$\alpha_{\equiv}(w) = [w]_{\equiv}.$$

is in fact an abstraction. That is, α_{\equiv} satisfies the compatibility property. Consequently, the notion of information abstraction can be characterized either by an abstract state mapping satisfying the compatibility property, or by a partition of the access sequence space with an equivalence relation satisfying the substitution property. The latter characterization has a clear advantage over the former — it is independent of the choice of abstract states. In fact, abstraction functions with entirely different codomains may induce the same equivalence relation, making the latter a better means of capturing the essence of the notion of information abstraction.

Let us call an equivalence relation satisfying the substitution property a *congruence relation*. We write $\text{SA}_{\equiv}(P)$ as a shorthand for $\text{SA}_{\alpha_{\equiv}}(P)$. A policy P is *enforceable* by a congruence relation \equiv if P is recognized by $\text{SA}_{\equiv}(P)$. We also write EM_{\equiv} as a shorthand for $\text{EM}_{\alpha_{\equiv}}$.

Defining abstractions in terms of congruence relations gives us a simple way of comparing their information complexities. Intuitively, more policies are enforceable by \equiv_1 than by \equiv_2 if \equiv_1 is more differentiating than \equiv_2 .

Theorem 6 *Let \equiv_1 and \equiv_2 be two congruence relations over Σ^* . If $\equiv_1 \subseteq \equiv_2$, then $\text{EM}_{\equiv_2} \subseteq \text{EM}_{\equiv_1}$. Moreover, the latter containment is proper if the former is proper.*

Consult the appendix for a proof of this theorem.

5.3. Lattice of Policy Classes

Define the binary operator \sqcup , the *join*, on the space of all congruence relations over Σ^* :

$$\equiv_1 \sqcup \equiv_2 = \equiv_1 \cap \equiv_2$$

That is, the join operator combines the differentiating power of its operands. Similarly, define the binary operator \sqcap , the *meet*, as follows:

$$\equiv_1 \sqcap \equiv_2 = (\equiv_1 \cup \equiv_2)^+$$

that is, the transitive closure of the union of the two operand congruence relations. In other words, a meet captures the common differentiating power of the two operands. The following theorem is a well-known result in automata theory:

Theorem 7 *The binary operators \sqcup and \sqcap define a lattice on the space of all congruence relations over Σ^* . The lattice has both a top element \equiv_{\top} and a bottom element \equiv_{\perp} .*

Proof: It is mechanical to check that \sqcup and \sqcap produce congruence relations and define a lattice. The top element \equiv_{\top} is the congruence relation \emptyset in which all members of Σ^* are distinct, while the bottom element \equiv_{\perp} is the congruence relation $\Sigma^* \times \Sigma^*$, in which all members of Σ^* are equivalent.

Intuitively, the top element \equiv_{\top} induces the class $EM_{\equiv_{\top}}$ of all EM-enforceable security policies. In contrast, the bottom element \equiv_{\perp} induces the class $EM_{\equiv_{\perp}}$ of security policies enforceable by SA with only one state. Such SA are *memory-less*: they do not track historical information at all, and grant access in a static manner. The rest of the congruence relations are ordered in decreasing differentiating power as we move down the lattice. As a consequence of Theorem 6, this lattice of congruence relations induces a poset of EM-enforceable policy classes ordered by class containment.

Example 8 *The abstraction function α defined in Example 5 induces a congruence \equiv_{α} so that, for $w, w' \in \Sigma^*$, $w \equiv_{\alpha} w'$ iff w and w' contain the same set of access events. It is obvious that $\equiv_{\top} \subset \equiv_{\alpha} \subset \equiv_{\perp}$. According to Theorem 6, $EM_{\equiv_{\perp}} \subset EM_{SHA} \subset EM_{\equiv_{\top}}$.*

5.4. Summary

This section generalizes the construction of SHA to obtain a lattice of policy classes. The notion of an information-based characterization of security policies is shown to be more general than shallow access history tracking. New policy classes may be characterized through the specification of either an abstraction mapping or a congruence relation. Theorems 6 and 7 also confirm the following intuition: the more information the execution monitor is allowed to track, the more security policies it is able to enforce.

6. Comparing the Differentiating Power of Abstractions

The policy classes EM_{SHA} , $EM_{\equiv_{\top}}$ and $EM_{\equiv_{\perp}}$ are first examples of naturally occurring policy classes in the policy class lattice. It is natural to ask if there are other members of the lattice that correspond to naturally occurring policy classes. This section explores policy classes induced by stack inspection and a forgetful variant of SHA. The former is shown to be equivalent to a simple abstraction, while the latter corresponds to a family of abstractions. The expressiveness of the two protection mechanisms are then compared to that of SHA and SA. The utility of the policy class

lattice and a number of useful proof techniques are also highlighted.

6.1. Generalized Stack Inspection

Stack inspection [30, 14] is an access control mechanism that has been implemented in safe runtime environments such as the Java Virtual Machine (JVM) [20] and the Common Language Runtime (CLR) [10]. The runtime rights of a piece of code is a function of the content of the execution stack. Stack inspection has been criticized as being unable to enforce certain history-based policies [1]. The following analysis adds another data point into the debate by formally establishing the limitation of stack inspection.

Let $\Sigma = \Sigma_{\oplus} \cup \Sigma_{\ominus}$, where $\Sigma_{\oplus} \cap \Sigma_{\ominus} = \emptyset$. Suppose further that there is a bijection of the type $\Sigma_{\oplus} \rightarrow \Sigma_{\ominus}$. Denote the image of the bijection by \bar{a} for $a \in \Sigma_{\oplus}$. Intuitively, $a \in \Sigma_{\oplus}$ represents a subprogram (e.g., function, method, procedure, etc) invocation event, while \bar{a} represents the corresponding return event. An execution stack can then be represented as a sequence in Σ_{\oplus}^* , with the top of the stack corresponding to the end of the sequence. A Stack Inspection Automaton (SIA) is an α -SA $\langle \Sigma, \Sigma_{\oplus}^* \cup \{\perp\}, \epsilon, \delta \rangle$, so that $\perp \notin \Sigma_{\oplus}^*$, and the abstraction function α is defined as follows:

$$\begin{aligned} \alpha(\epsilon) &= \epsilon \\ \alpha(wa) &= \begin{cases} \alpha(w)a & \text{if } \alpha(w) \neq \perp \\ \perp & \text{otherwise} \end{cases} \\ \alpha(w\bar{a}) &= \begin{cases} w' & \text{if } \alpha(w) = w'a \text{ for } w' \in \Sigma_{\oplus}^* \\ \perp & \text{otherwise} \end{cases} \end{aligned}$$

where $w \in \Sigma^*$ and $a \in \Sigma_{\oplus}$. Intuitively, a SIA maintains a stack of invocation records, and bases its access granting decisions solely on the content of the stack. The error state \perp is for handling invocation sequences that are not well-formed. Denote EM_{α} by EM_{SIA} . Since \equiv_{α} is not empty, it follows from Theorem 6 that the SIA is strictly less expressive than the SA.

Theorem 9 $EM_{SIA} \subset EM$.

The expressiveness of the SIA and that of the SHA, however, are not comparable.

Theorem 10 EM_{SIA} and EM_{SHA} are incomparable.

Proof: Let $\Sigma_{\oplus} = \{a, b\}$, and thus $\Sigma = \{a, \bar{a}, b, \bar{b}\}$. Consider the SIA-enforceable policy “*b should be disallowed if a is in the execution stack.*” An enforcing SIA accepts $a\bar{a}b$ and $a\bar{a}a$, but rejects $a\bar{a}b$.

Suppose the policy is enforceable by a SHA M . Then M accepts $a\bar{a}b$ and $a\bar{a}a$.

$$\begin{aligned} \emptyset &\xrightarrow{a}_M \{a\} \xrightarrow{\bar{a}}_M \{a, \bar{a}\} \xrightarrow{b}_M \{a, \bar{a}, b\} \\ \emptyset &\xrightarrow{a}_M \{a\} \xrightarrow{\bar{a}}_M \{a, \bar{a}\} \xrightarrow{a}_M \{a, \bar{a}\} \end{aligned}$$

But then it means M also accepts $a\bar{a}ab$, a contradiction.

$$\begin{aligned} \emptyset &\xrightarrow{a}_M \{a\} \xrightarrow{\bar{a}}_M \{a, \bar{a}\} \\ &\xrightarrow{a}_M \{a, \bar{a}\} \xrightarrow{b}_M \{a, \bar{a}, b\} \end{aligned}$$

Therefore, $\text{EM}_{\text{SIA}} \not\subseteq \text{EM}_{\text{SHA}}$.

Now, consider the SHA-enforceable policy “ b should be disallowed if a has already occurred.” An enforcing SHA accepts $b\bar{b}a\bar{a}$ but rejects $a\bar{a}b$.

Suppose the policy is enforceable by a SIA N . N accepts $b\bar{b}a\bar{a}$.

$$\epsilon \xrightarrow{b}_N b \xrightarrow{\bar{b}}_N \epsilon \xrightarrow{a}_N a \xrightarrow{\bar{a}}_N \epsilon$$

But then it means N also accepts $a\bar{a}b$, a contradiction.

$$\epsilon \xrightarrow{a}_N a \xrightarrow{\bar{a}}_N \epsilon \xrightarrow{b}_N b$$

Therefore, $\text{EM}_{\text{SHA}} \not\subseteq \text{EM}_{\text{SIA}}$.

6.2. Shallow History Automata with Forgetting

The formulation of SHA in Section 3 mandates that the entire set of previously granted access events be remembered. Unable to forget previous access events, the SHA is only capable of handling assured pipelines with an acyclic enabling relation. This motivates the introduction of forgetting into the SHA. Given a set Σ of access events, and a function $f : \mathcal{F}(\Sigma) \times \Sigma \rightarrow \mathcal{F}(\Sigma)$, a *forgetful shallow history automaton with forgetting function f* , or a f -SHA, is a SA $\langle \Sigma, \mathcal{F}(\Sigma), \emptyset, \delta \rangle$ such that $\delta(H, a) = H \cup \{a\} \setminus f(H, a)$ whenever δ is defined at $\langle H, a \rangle \in \mathcal{F}(\Sigma) \times \Sigma$. Intuitively, f specifies the subset of shallow history that should be erased from memory. For a fixed f , a f -SHA is merely an α_f -SA with the abstraction function $\alpha_f : \Sigma^* \rightarrow \mathcal{F}(\Sigma)$ defined as follows.

$$\begin{aligned} \alpha_f(\epsilon) &= \emptyset \\ \alpha_f(wa) &= \alpha_f(w) \cup \{a\} \setminus f(\alpha_f(w), a) \end{aligned}$$

The following invariant can be established by induction:

$$\alpha_f(w) \subseteq \{a \in \Sigma \mid a \text{ occurs in } w\} \quad (4)$$

Example 11 Suppose the “no circularity” clause is removed from the definition of an enabling function. An assured pipeline policy with a (possibly cyclic) enabling relation e can be enforced by a f -SHA, where

$$f(H, \langle s, o \rangle) = \{\langle s', o \rangle \in H \mid s' \neq s\}$$

Specifically, the transition function δ_e of the f -SHA is defined at exactly the points $\langle H, \langle s, o \rangle \rangle$ satisfying one of the following conditions:

- $s = \text{create}$ and $\langle s, o \rangle \notin H$, or
- for some $s' \in S$, all of the following hold:
 - $\langle s', o \rangle \in H$, and
 - $\langle s', s \rangle \in e$, and

By construction N_e enforces the assured pipeline policy specified by enabling relation e .

With an appropriately chosen forgetting function f , a f -SHA can enforce policies not enforceable by a regular SHA. In fact, a regular SHA is nothing more than a forgetful SHA with a constant forgetting function. In short, forgetting increases the expressiveness of a SHA. It is therefore natural to ask if all EM-enforceable policies can be recognized by some forgetful SHA.

Let us begin with a weaker result. By a counting argument in the style of the pumping lemma, one can show that, for any fixed forgetting function f , the f -SHA is strictly less expressive than the SA.

Theorem 12 Given any forgetting function f , $\text{EM}_{\alpha_f} \subset \text{EM}_{\equiv_{\top}}$.

Proof: Let E be a finite subset of Σ . By (4), $\alpha_f(w) \in \mathcal{F}(E)$ for any $w \in E^*$. The size of $\mathcal{F}(E)$ is finite, while E^* contains infinitely many members. By the pigeonhole principle, the congruence relation \equiv_{α_f} is clearly non-empty. It follows from Theorem 6 that $\text{EM}_{\alpha_f} \subset \text{EM}$.

The above result can be further strengthened: the class of all policies enforceable by forgetful SHA does not cover all of the EM-enforceable policies.

Theorem 13 $\bigcup_f \text{EM}_{\alpha_f} \subset \text{EM}$.

Proof: We construct a non-empty prefix-closed policy P so that $P \not\subseteq \text{EM}_{\alpha_f}$ for every forgetting function f . The case for $|\Sigma| = 1$ is trivial to deal with. Let $E = \{a, b\}$ be a size-2 subset of Σ . Let $r : E^3 \rightarrow \mathcal{N}$ be an injective ranking function for length-3 sequences in E^3 . Let $P = \bigcup_{w \in E^3} \text{prefix}(wa^{r(w)})$. Assume there is an α_f -SA M such that $P = \mathcal{P}(M)$. By (4), $\alpha_f(w) \in \mathcal{F}(E)$ for any $w \in E^*$. The size of $\mathcal{F}(E)$ is 4, while E^3 contains 8 members. By the pigeonhole principle, there exist $u, v \in E^3$ such that $\alpha_f(u) = \alpha_f(v)$. Without loss of generality, assume that $r(u) > r(v)$. Proposition 3 therefore guarantees the following.

$$\begin{aligned} \alpha(\epsilon) &\xrightarrow{u}_M \alpha(u) \xrightarrow{a^{r(u)}}_M \alpha(ua^{r(u)}) \\ \alpha(\epsilon) &\xrightarrow{v}_M \alpha(v) \xrightarrow{a^{r(v)}}_M \alpha(va^{r(v)}) \end{aligned}$$

By construction, $\alpha(u) = \alpha(v)$, and thus the following transitions are legitimate.

$$\alpha(\epsilon) \xrightarrow{v}_M \alpha(v) \xrightarrow{a^{r(u)}}_M \alpha(va^{r(u)})$$

But $va^{r(u)} \notin P$ by construction, a contradiction.

Notice that the proof makes use of a diagonalization-style argument, in which a policy is constructed to diagonalize against *congruences* induced by α_f .

6.3. Summary

Classes of naturally occurring policies, including those induced by stack inspection and forgetful SHA have been shown to result from an appropriately constructed abstraction or family of abstractions. Several proof techniques for comparing the differentiating power of abstractions have been highlighted.

7. Discussion

The notion of information complexity is in fact very sensitive to the choice of access events. We have seen that, with a *fixed* set Σ of access events, there are EM-enforceable policies that SHA cannot enforce. Now, if we *instrument* the access events so that an event is a pair of the form $\langle a, i \rangle$, in which i is the time index of access a in the event sequence, then a SHA can enforce every policy that a SA could enforce without the instrumentation. Consequently, one must fix the set of access events to get a fair comparison of the expressiveness of different abstractions. Also, one should attempt to pick the most “*natural*” formulation of access events.

To simplify discussion, all event sequences from Σ^* are treated as being equally plausible. In reality, the execution environment may generate only sequences belonging to a subset of Σ^* . For example, in the case of (linear) assured pipelines, the requirement that every access event occurs only once is in fact a part of the behavior characteristics of the execution environment, and not a part of the security policy per se. To focus on the essence of the security policy rather than the idiosyncrasy of the execution environment, one could have defined a security automaton alternatively as a 5-tuple $\langle \Sigma, \Psi, Q, q_0, \delta \rangle$, in which the newly introduced second component $\Psi \subseteq \Sigma^*$ is the set of all event sequences that could be generated by the execution environment. For an elaboration of this treatment, consult the works of Schneider [26] and Bauer *et al* [3].

A number of theoretical apparatus, including that of automata homomorphism, congruence relations, compatibility property, substitution property, and lattices of congruence relations, are all borrowed from Büchi’s algebraic approach to automata theory [8, 17]. Instead of using these tools to

study the minimization and decomposition of individual automata, they are applied to obtain a novel, information-based characterization of security policies:

1. Automata homomorphisms and their counterparts, congruence relations, are used for defining automata classes with restricted access to history information.
2. Each automata class is then used to characterize a class of security policies enforceable with the corresponding information constraint.
3. The mapping from the lattice of congruence relations to their corresponding policy classes is then shown to be order-preserving.

An abstraction is basically a syntactic means for defining the data structure tracked by the execution monitor. The corresponding congruence is the semantics of this syntax. Seen in this light, the present theoretical framework offers a very precise semantic infrastructure for defining syntactic constructs that represent the data structure tracked by an execution monitor. A future direction is to employ the current framework to define special-purpose policy languages for mobile code systems. By bringing in automata theory into the study, it is hoped that automata decomposition results from the theory could facilitate the development of reusable policy components. Specifically, Corollary 16, Proposition 17 and Proposition 18 imply that, given a congruence \equiv , $\mathcal{P}(\text{SA}_{\equiv}(\cdot))$ is a closure operator [9], and thus policies enforceable by \equiv form a complete lattice. This constitutes a starting point from which one could provide a semantic foundation for policy decomposition.

The notion of abstraction has close resemblance to that of abstraction interpretation [22]. A future direction is to apply the current framework to study security policies in the context of programming languages, and explore interactions with the theory of abstract interpretation.

McLean [21] proposed a framework for Mandatory Access Control (MAC) models that support changes in security levels. Models of this framework form a Boolean algebra. Further enrichment of the framework to handle n -person rules results in a distributive lattice of security models. An interesting further work is to explore if McLean’s lattice can be embedded in the lattice of EM-enforceable policies.

8. Conclusion

A novel approach has been proposed to address the open question raised by Bauer *et al*. The space of EM-enforceable security policies is classified according to the information consumed by the execution monitor. The feasibility of this approach has been demonstrated by the characterization of security policies enforceable by tracking shal-

low execution history. Although the class is provably less expressive than the general class of EM-enforceable policies, it nevertheless contains a number of naturally occurring security policies. Generalization of the technique allows one to define a lattice of security policy classes, in which member classes are ordered by the amount of information that must be tracked by an enforcing execution monitor. The expressiveness of access control mechanisms such as stack inspection has been studied in this lattice framework.

Acknowledgments

I would like to thank John McLean and the anonymous reviewers for their constructive comments. The quality of this paper has been greatly improved by their feedback.

References

- [1] M. Abadi and C. Fournet. Access control based on execution history. In *Proceedings of the 10th Annual Network and Distributed System Security Symposium*, San Diego, California, Feb. 2003.
- [2] B. Alpern and F. B. Schneider. Recognizing safety and liveness. *Distributed Computing*, 2(3):117–126, 1987.
- [3] L. Bauer, J. Ligatti, and D. Walker. More enforceable security policies. In *Proceedings of the Workshop on Foundations of Computer Security (FCS'02)*, Copenhagen, Denmark, July 2002.
- [4] M. Bernaschi, E. Gabrielli, and L. V. Mancini. REMUS: A security-enhanced operating system. *ACM Transactions on Information and System Security*, 5(1):36–61, Feb. 2002.
- [5] K. Biba. Integrity considerations for secure computer systems. Technical Report 76–372, U. S. Air Force Electronic Systems Division, 1977.
- [6] W. E. Boebert and R. Y. Kain. A practical alternative to hierarchical integrity policies. In *Proceedings of the 8th National Computer Security Conference*, pages 18–27, Oct. 1985.
- [7] D. F. C. Brewer and M. J. Nash. The Chinese Wall security policy. In *Proceedings of the IEEE Symposium on Research in Security and Privacy*, pages 206–214, Oakland, California, May 1989.
- [8] J. R. Büchi. *Finite Automata, Their Algebra and Grammars*. Springer-Verlag, 1988.
- [9] S. Burris and H. P. Sankappanavar. *A Course in Universal Algebra*. Springer-Verlag, 1981.
- [10] ECMA. *Standard ECMA-335: Common Language Infrastructure (CLI)*, 2nd edition, Dec. 2002.
- [11] G. Edjladi, A. Acharya, and V. Chaudhary. History-based access control for mobile code. In *Proceedings of the 5th ACM Conference on Computer and Communications Security*, San Francisco, CA, USA, 1998.
- [12] A. Edwards, T. Jaeger, and X. Zhang. Runtime verification of authorization hook placement for the Linux security modules framework. In *Proceedings of the 9th ACM Conference on Computer and Communications Security*, pages 225–234, Washington, DC, Nov. 2002.
- [13] D. Evans and A. Twyman. Flexible policy-directed code safety. In *Proceedings of the 1999 IEEE Symposium on Security and Privacy*, pages 32–45, Oakland, California, May 1999.
- [14] C. Fournet and A. D. Gordon. Stack inspection: Theory and variants. *ACM Transactions on Programming Languages and Systems*, 25(3):360–399, May 2003.
- [15] T. Fraser. LOMAC: Low water-mark integrity protection for COTS environments. In *Proceedings of the 2000 IEEE Symposium on Security and Privacy*, Oakland, California, 2000.
- [16] K. W. Hamlen, G. Morrisett, and F. B. Schneider. Computability classes for enforcement mechanisms. Technical Report TR 2003-1908, Computer Science Department, Cornell University, Aug. 2003.
- [17] M. A. Harrison. *Introduction to Switching and Automata Theory*. McGraw-Hill, 1965.
- [18] J. Ligatti, L. Bauer, and D. Walker. Edit automata: Enforcement mechanisms for run-time security policies. *International Journal of Information Security*, 2003. To appear.
- [19] T. Y. Lin. Chinese Wall security policy — an aggressive model. In *Proceedings of the Fifth Annual Computer Security Applications Conference*, pages 282–289, Dec. 1989.
- [20] T. Lindholm and F. Yellin. *The Java Virtual Machine Specification*. Addison Wesley, 2nd edition, 1999.
- [21] J. McLean. The algebra of security. In *Proceedings of the 1988 IEEE Symposium on Security and Privacy*, pages 2–7, Oakland, CA, Apr. 1988.
- [22] F. Nielson, H. R. Nielson, and C. Hankin. *Principles of Program Analysis*. Springer, 1999.
- [23] A. Rudys and D. S. Wallach. Termination in language-based systems. *ACM Transactions on Information and System Security*, 5(2):138–168, May 2002.
- [24] A. Rudys and D. S. Wallach. Transactional rollback for language-based systems. In *Proceedings of the International Conference on Dependable Systems and Networks (DSN'02)*, pages 439–448, Washington, D.C., June 2002.
- [25] R. S. Sandhu. A lattice interpretation of the Chinese Wall policy. In *Proceedings of the 15th NIST-NCSC National Computer Security Conference*, pages 221–235, Baltimore, MD, Oct. 1992.
- [26] F. B. Schneider. Enforceable security policies. *ACM Transactions on Information and System Security*, 3(1):30–50, Feb. 2000.
- [27] Úlfar Erlingsson and F. B. Schneider. SASI enforcement of security policies: A retrospective. In *Proceedings of the 1999 New Security Paradigms Workshop*, pages 87–95, Caledon Hills, Ontario, Canada, Sept. 1999.
- [28] Úlfar Erlingsson and F. B. Schneider. IRM enforcement of Java stack inspection. In *Proceedings of the 2000 IEEE Symposium on Security and Privacy*, pages 246–255, Berkeley, California, May 2000.
- [29] M. Viswanathan. *Foundations for the Run-time Analysis of Software systems*. PhD thesis, University of Pennsylvania, Dec. 2000.

- [30] D. S. Wallach, A. W. Appel, and E. W. Felten. SAFKASI: A security mechanism for language-based systems. *ACM Transactions on Software Engineering and Methodology*, 9(4):341–378, Oct. 2000.
- [31] I. Welch and R. J. Stroud. Using reflection as a mechanism for enforcing security policies on compiled code. *Journal of Computer Security*, 10(4):399–432, 2002.
- [32] W. D. Young, P. A. Telega, and W. E. Boebert. A verified labler for the secure Ada target. In *Proceedings of the 9th National Computer Security Conference*, pages 55–61, Sept. 1986.

Appendix: Proof of Theorem 6

Observation 14 Let P be a non-empty prefix-closed policy, and \equiv be a congruence relation. Then, for any $w \in \Sigma^*$ and $a \in \Sigma$, we have:

$$w \xrightarrow{a}_{SA(P)} wa \Rightarrow [w]_{\equiv} \xrightarrow{a}_{SA_{\equiv}(P)} [wa]_{\equiv} \quad (5)$$

and

$$[w]_{\equiv} \xrightarrow{a}_{SA_{\equiv}(P)} [wa]_{\equiv} \Rightarrow w' \xrightarrow{a}_{SA(P)} w'a \quad (6)$$

for some $w' \in [w]_{\equiv}$.

Proposition 15 Let P be a non-empty prefix-closed policy, and \equiv_1 and \equiv_2 congruence relations such that $\equiv_1 \subseteq \equiv_2$. Then

$$\mathcal{P}(SA_{\equiv_1}(P)) \subseteq \mathcal{P}(SA_{\equiv_2}(P)).$$

Proof: The statement can be easily demonstrated by induction with the help of the following lemma.

Lemma: Denote $SA_{\equiv_1}(P)$ and $SA_{\equiv_2}(P)$ by M_1 and M_2 respectively. For any $w \in \Sigma^*$ and $a \in \Sigma$, we have:

$$[w]_{\equiv_1} \xrightarrow{a}_{M_1} [wa]_{\equiv_1} \Rightarrow [w]_{\equiv_2} \xrightarrow{a}_{M_2} [wa]_{\equiv_2}.$$

Proof:

$$\begin{aligned} & [w]_{\equiv_1} \xrightarrow{a}_{M_1} [wa]_{\equiv_1} \\ \Rightarrow & w' \xrightarrow{a}_{SA(P)} w'a && w' \in [w]_{\equiv_1}, \text{ by (6)} \\ \Rightarrow & [w']_{\equiv_2} \xrightarrow{a}_{M_2} [w'a]_{\equiv_2} && \text{by (5)} \\ \Rightarrow & [w]_{\equiv_2} \xrightarrow{a}_{M_2} [wa]_{\equiv_2} && \because \equiv_1 \subseteq \equiv_2 \text{ and} \\ & && w' \equiv_1 w \text{ and (3)} \end{aligned}$$

Corollary 16 Let P be a non-empty prefix-closed policy, and \equiv a congruence relation. Then

$$P \subseteq \mathcal{P}(SA_{\equiv}(P)).$$

Proof: Let $\equiv_{\top} = \emptyset$ be the congruence relation in which every sequence belong to a distinct equivalence class. Then $P = \mathcal{P}(SA(P)) = \mathcal{P}(SA_{\equiv_{\top}}(P))$. The result follows from Proposition 15.

Proposition 17 Let P_1 and P_2 be non-empty prefix-closed policies so that $P_1 \subseteq P_2$, and let \equiv be a congruence relation. Then

$$\mathcal{P}(SA_{\equiv}(P_1)) \subseteq \mathcal{P}(SA_{\equiv}(P_2)).$$

Proof: The statement can be demonstrated easily by induction with the help of the following lemma.

Lemma: Denote $SA_{\equiv}(P_1)$ and $SA_{\equiv}(P_2)$ by M_1 and M_2 respectively. Then for $w \in \Sigma^*$ and $a \in \Sigma$, we have

$$[w]_{\equiv} \xrightarrow{a}_{M_1} [wa]_{\equiv} \Rightarrow [w]_{\equiv} \xrightarrow{a}_{M_2} [wa]_{\equiv}$$

Proof:

$$\begin{aligned} & [w]_{\equiv} \xrightarrow{a}_{M_1} [wa]_{\equiv} \\ \Rightarrow & w' \xrightarrow{a}_{SA(P_1)} w'a && w' \in [w]_{\equiv}, \text{ by (6)} \\ \Rightarrow & w' \xrightarrow{a}_{SA(P_2)} w'a && \because P_1 \subseteq P_2 \\ \Rightarrow & [w']_{\equiv} \xrightarrow{a}_{M_2} [w'a]_{\equiv} && \text{by (5)} \\ \Rightarrow & [w]_{\equiv} \xrightarrow{a}_{M_2} [wa]_{\equiv} && \because w \equiv w' \text{ and (3)} \end{aligned}$$

Proposition 18 Let P be a non-empty prefix-closed policy and \equiv a congruence relation. Then

$$\mathcal{P}(SA_{\equiv}(\mathcal{P}(SA_{\equiv}(P)))) = \mathcal{P}(SA_{\equiv}(P)).$$

Therefore, $\mathcal{P}(SA_{\equiv}(P)) \in EM_{\equiv}$.

Proof: By Corollary 16 and Proposition 17, we already have $\mathcal{P}(SA_{\equiv}(P)) \subseteq \mathcal{P}(SA_{\equiv}(\mathcal{P}(SA_{\equiv}(P))))$. It therefore suffices to show that $\mathcal{P}(SA_{\equiv}(\mathcal{P}(SA_{\equiv}(P)))) \subseteq \mathcal{P}(SA_{\equiv}(P))$. The inclusion can be demonstrated easily by induction with the help of the following lemma:

Lemma: Denote $SA_{\equiv}(P)$ and $SA_{\equiv}(\mathcal{P}(SA_{\equiv}(P)))$ by M_1 and M_2 respectively. Then for any $w \in \Sigma^*$ and $a \in \Sigma$, we have:

$$[w]_{\equiv} \xrightarrow{a}_{M_2} [wa]_{\equiv} \Rightarrow [w]_{\equiv} \xrightarrow{a}_{M_1} [wa]_{\equiv}.$$

Proof:

$$\begin{aligned} & [w]_{\equiv} \xrightarrow{a}_{M_2} [wa]_{\equiv} \\ \Rightarrow & w' \xrightarrow{a}_{SA(\mathcal{P}(M_1))} w'a && w' \equiv w, \text{ by (6)} \\ \Rightarrow & w', w'a \in \mathcal{P}(M_1) && \text{by def of } SA(\cdot) \\ \Rightarrow & [w']_{\equiv} \xrightarrow{a}_{M_1} [w'a]_{\equiv} && \text{by Proposition 3} \\ \Rightarrow & [w]_{\equiv} \xrightarrow{a}_{M_1} [wa]_{\equiv} && \because w \equiv w' \text{ and (3)} \end{aligned}$$

Proof of Theorem 6

Suppose $\equiv_1 \subseteq \equiv_2$. From Proposition 15 and Corollary 16, we have:

$$P \subseteq \mathcal{P}(\text{SA}_{\equiv_1}(P)) \subseteq \mathcal{P}(\text{SA}_{\equiv_2}(P)).$$

It follows that $P = \mathcal{P}(\text{SA}_{\equiv_2}(P))$ implies $P = \mathcal{P}(\text{SA}_{\equiv_1}(P))$. Therefore, $\text{EM}_{\equiv_2} \subseteq \text{EM}_{\equiv_1}$.

Now, suppose further that $\equiv_1 \subset \equiv_2$. We want to show that there is a policy in EM_{\equiv_1} that is not in EM_{\equiv_2} .

Let w be a shortest sequence in Σ^* such that $[w]_{\equiv_1} \subset [w]_{\equiv_2}$. Let w' be a shortest sequence in $[w]_{\equiv_2} \setminus [w]_{\equiv_1}$. We then also have $[w']_{\equiv_1} \subset [w']_{\equiv_2}$.

Let $a \in \Sigma$ be an arbitrary access event. Define non-empty prefix-closed policy $P = \text{prefix}(wa) \cup \text{prefix}(w')$. Define also non-empty prefix-closed policy $P' = \mathcal{P}(\text{SA}_{\equiv_1}(P))$. We claim that, although P' is obviously a member of EM_{\equiv_1} (Proposition 18), it does not belong to EM_{\equiv_2} , that is, $P' \neq \mathcal{P}(\text{SA}_{\equiv_2}(P'))$. To demonstrate this, we show that (1) $w'a \notin P'$, but (2) $w'a \in \mathcal{P}(\text{SA}_{\equiv_2}(P'))$.

1. $w'a \notin P'$: If $w'a \in \mathcal{P}(\text{SA}_{\equiv_1}(P))$, then there must be some $u, ua \in P$, so that $u \in [w']_{\equiv_1}$ and $ua \in [w'a]_{\equiv_1}$.

The following case analysis demonstrates that this is impossible.

- (a) $u \notin \text{prefix}(w') \setminus \{w'\}$, for otherwise, $u \equiv_1 w' \not\equiv_1 w$ but $u \equiv_2 w' \equiv_2 w$ (because $\equiv_1 \subset \equiv_2$), and thus u would be a shorter candidate for w' .
- (b) $u \notin \text{prefix}(w) \setminus \{w\}$, for otherwise, by the same argument as the previous case, u would be a shorter candidate for w .
- (c) $u \neq w$, for otherwise $w \equiv_1 w'$, contradicting the definition of w' .
- (d) $u \neq wa$, for otherwise either $waa \notin P$ or $u \in \text{prefix}(w') \setminus \{w'\}$.
- (e) $u \neq w'$, for otherwise either $w'a \notin P$ or $u = w$ or $u \in \text{prefix}(w) \setminus \{w\}$.

2. $w'a \in \mathcal{P}(\text{SA}_{\equiv_2}(P'))$: Notice that, as Corollary 16 implies $P \subseteq P'$, it then follows from Proposition 17 that $\mathcal{P}(\text{SA}_{\equiv_2}(P)) \subseteq \mathcal{P}(\text{SA}_{\equiv_2}(P'))$. Consequently, it suffices to show $w'a \in \mathcal{P}(\text{SA}_{\equiv_2}(P))$.

As $w, wa \in P$, we have $[w]_{\equiv_2} \xrightarrow{a}_{\text{SA}_{\equiv_2}(P)} [wa]_{\equiv_2}$. By Proposition 3, $w' \in P$ implies that $[\epsilon]_{\equiv_2} \xrightarrow{w'}_{\text{SA}_{\equiv_2}(P)} [w']_{\equiv_2}$. Since $w \equiv_2 w'$, $[w]_{\equiv_2}$ and $[w']_{\equiv_2}$ are the same equivalence class, and, by the substitution property, so are $[wa]_{\equiv_2}$ and $[w'a]_{\equiv_2}$. Therefore, $[\epsilon]_{\equiv_2} \xrightarrow{w'}_{\text{SA}_{\equiv_2}(P)} [w']_{\equiv_2} \xrightarrow{a}_{\text{SA}_{\equiv_2}(P)} [w'a]_{\equiv_2}$.