Enforceable Security Policies

FRED B. SCHNEIDER
Cornell University

A precise characterization is given for the class of security policies enforceable with mechanisms that work by monitoring system execution, and automata are introduced for specifying exactly that class of security policies. Techniques to enforce security policies specified by such automata are also discussed.

Categories and Subject Descriptors: D.2.1 [**Software Engineering**]: Requirements/Specifications; D.2.9 [**Software Engineering**]: Management—*Software configuration management*; D.4.6 [**Operating Systems**]: Security and Protection; F.1.1 [**Computation by abstract devices**]: Models of Computation—*Automata* (e.g., finite, push-down, resource-bounded); F.3.1 [**Logics and Meanings of Programs**]: Specifying and Verifying and Reasoning about Programs—*Specification techniques*; K.4.4 [**Computers and Society**]: Electronic Commerce—*Security*; K.6.5 [**Management of Computing and Information Systems**]: Security and Protection

General Terms: Security

Additional Key Words and Phrases: security automata, safety properties, security policies, EM security policies, proof carrying code, SASI, inlined reference monitors

1. INTRODUCTION

A *security policy* defines execution that, for one reason or another, has been deemed unacceptable. For example, a security policy might concern

—*access control*, and restrict what operations principals can perform on objects,

—*information flow*, and restrict what principals can infer about objects from observing system behavior, or

Supported in part by ARPA/RADC grant F30602-96-1-0317, AFOSR grant F49620-94-1-0198, Defense Advanced Research Projects Agency (DARPA) and Air Force Research Laboratory, Air Force Material Command, USAF, under agreement number F30602-99-1-0533, National Science Foundation Grant 9703470, and a grant from Intel Corporation. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of these organizations or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright annotation thereon.
Author's address: Cornell University, Upson Hall, Ithaca, NY ; email: fbs@cs.cornell.edu.

—*availability*, and restrict principals from denying others the use of a resource.

To date, general-purpose security policies, like those above, have attracted the most attention. But application-dependent and special-purpose security policies are increasingly important.[1] A system to support mobile code, like Java [Gong 1997], might prevent information leakage by enforcing a security policy that bars messages from being sent after files have been read. To support electronic commerce, a security policy might prohibit executions in which a customer pays for a service but the seller does not provide that service. And finally, electronic storage and retrieval of intellectual property is governed by rights-management schemes that restrict not only the use of stored materials but also the use of any derivatives [Stefik 1996].

The value of application-dependent and special-purpose security policies is perhaps best explained in terms of the Principle of Least Privilege [Saltzer and Schroeder 1975], which holds that each principal be accorded the minimum access needed to accomplish its task. Clearly, richer notions of "minimum access" allow the Principle of Least Privilege to discriminate better between those actions that should and those that should not be allowed. Application-dependent security policies can depend on an application's state along with the semantics of that application's abstractions, so richer prescriptions for "minimum access" now become useful. In contrast, operating system abstractions—the traditional vocabulary for security policies—constitute a coarse basis for prescribing "minimum access," often forcing security policies to be approximations for what is desired.

The practicality of any security policy depends on whether that policy is enforceable and at what cost. This paper addresses those questions for the class of enforcement mechanisms that work by monitoring execution steps of some system, herein called the *target*, and terminating[2] the target's execution if it is about to violate the security policy being enforced. We call this class EM, for Execution Monitoring. EM includes security kernels, reference monitors, firewalls, and most other operating system and hardware-based enforcement mechanisms that have appeared in the literature. Our targets may be objects, modules, processes, subsystems, or entire systems; the execution steps monitored may range from fine-grained actions (such as memory accesses) to higher-level operations (such as method calls) to operations that change the security-configuration and thus restrict subsequent execution.

Mechanisms that use more information than would be available only from observing the steps of a target's execution are, by definition, excluded from EM. Information provided to an EM mechanism is thus insufficient to

[1]For example, see [Null and Wong 1992; Woo and Lam 1992; Jajodia et al. 1997; Edjlali et al. 1998; Pandey and Hashii 1998; Evans and Twyman 1999; Grimm and Bershad 1999].

[2]The case where instead of terminating the target, an attempt to violate the policy causes substitution of an acceptable execution step for an unacceptable one is not materially different. This is discussed in Section 4.

predict future steps the target might take, alternative possible executions, or all possible target executions. Therefore, compilers and theorem-provers, which analyze a static representation of a target to deduce information about all of its possible executions, are not EM mechanisms. The availability of information about future execution, about possible alternative executions, or about all possible target executions gives power to an enforcement mechanism. How much power remains an open question.

Also outside EM are mechanisms that modify a target before executing it. The modified target must be equivalent to the original, except for aborting executions that violate the security policy of interest. A definition for equivalent is thus required to analyze this class of mechanisms.

A formal characterization of what can and cannot be accomplished using mechanisms in EM has both practical and theoretical utility. Clearly, such a characterization can inform system builders' selections of enforcement mechanisms by circumscribing the intrinsic limits of reference monitors and derivative mechanisms. From a theoretical perspective, the characterization constitutes a first step toward a taxonomy of security policies that is based on a mathematical semantics of programs. Two other classes in that taxonomy might come from relaxing EM's defining restrictions: (i) a class of enforcement mechanisms that have access to some (perhaps incomplete) representation of the target, and (ii) a class of enforcement mechanisms that modify the target before execution.

We proceed as follows. In Section 2, a precise characterization is given for security policies that can be enforced using mechanisms in EM. An automata-based formalism for specifying those security policies is the subject of Section 3. Mechanisms in EM for enforcing security policies specified by automata are described in Section 4. Section 5 discusses some pragmatic issues related to specifying and enforcing security policies as well as the application of our enforcement mechanisms to safety-critical systems. The appendix contains a summary of the notation used in the paper.

2. CHARACTERIZING EM ENFORCEMENT MECHANISMS

We represent target executions by finite and infinite sequences, where Ψ denotes a universe of all possible finite and infinite sequences. The manner in which executions are represented is irrelevant here. Finite and infinite sequences of atomic actions, of higher-level system steps, of program states, or of state/action pairs are all plausible alternatives. A target S defines a subset Σ_S of Ψ corresponding to the executions of S.

A characterization of EM enforceable security policies is interesting only if the definition being used for security policy is broad enough so that it does not exclude things usually considered security policies.[3] Also, the definition must be independent of how EM is defined, for otherwise the characterization of EM-enforceable security policies would be a tautology, hence uninteresting. We therefore adopt the following.

[3]However, there is no harm in being liberal about what is considered a security policy.

Definition of Security Policy: A *security policy* is specified by giving a predicate on sets of executions. A target S *satisfies* security policy $\mathcal{P}$ if and only if $\mathcal{P}(\Sigma_S)$ equals $true$.

These definitions are broad[4] (giving at least as much power for defining computations that are disallowed by security policies as for specifying the computations that are possible by targets) and correspond to the intuition that security policies rule out target executions that are deemed unacceptable.

Given a security policy $\mathcal{P}$ and sets Σ and Π of executions, note we do not require that if Σ satisfies $\mathcal{P}$ and $\Pi \subset \Sigma$ holds, then Π satisfies $\mathcal{P}$. Imposing such a requirement on security policies disqualifies interesting candidates. For instance, the requirement precludes information flow (as defined informally in Section 1) from being considered a security policy—universe Ψ of all finite and infinite state sequences satisfies information flow (because, for this set of sequences, the value of no state component is correlated with others), but a subset Π containing only those executions in which the value of a variable x in each execution is correlated with the value of y (say) might violate an information flow policy.

Safety Properties and EM Enforceability

By definition, enforcement mechanisms in EM work by monitoring execution of the target. Thus, any security policy $\mathcal{P}$ that can be enforced using a mechanism from EM must be specified by a predicate of the form

$$\mathcal{P}(\Pi)\colon (\forall \sigma \in \Pi\colon \hat{\mathcal{P}}(\sigma)) \tag{1}$$

where $\hat{\mathcal{P}}$ is a predicate on (individual) executions. $\hat{\mathcal{P}}$ formalizes the criteria used by the enforcement mechanism for deciding whether or not to terminate an execution that would otherwise violate the policy being enforced. In Alpern and Schneider [1985] and the literature on linear-time concurrent program verification, a set of executions is called a *property* if set membership is determined by each element alone and not by other members of the set. Using that terminology, we conclude from (1) that a security policy must be a property in order for that policy to have an enforcement mechanism in EM.

Not every security policy is a property. Some security policies cannot be defined using the criteria that individual executions must each satisfy in isolation. For example, the information flow policy discussed above characterizes sets that are not properties (as proved in McLean [1994][5]). Whether information flows from variable x to y in a given execution depends, in

[4]The definitions clearly subsume the noninterference-based definition of security policy in Goguen and Meseguer [1982].

[5]McLean acknowledged James Gray III as pointing out this limitation for dealing with security in frameworks based on our property abstraction.

part, on what values y takes in other possible executions (and whether those values are correlated with the value of x). A predicate to specify such sets of executions cannot be constructed using only predicates defined on single executions in isolation.

Not every property is EM enforceable. Enforcement mechanisms in EM cannot base decisions on possible future execution, since that information is, by definition, not available to such a mechanism, and this further restricts what security policies can be enforced by EM mechanisms. Consider security policy $\mathcal{P}$ of (1), and suppose σ' is the prefix of some finite or infinite execution σ where $\hat{\mathcal{P}}(\sigma) = \mathit{true}$ and $\hat{\mathcal{P}}(\sigma') = \mathit{false}$ hold. Because execution of a target might terminate before σ' is extended into σ, an enforcement mechanism for $\mathcal{P}$ must prohibit σ' (even though supersequence σ satisfies $\hat{\mathcal{P}}$).

We can formalize this requirement as follows. For σ a finite or infinite execution having i or more steps, and τ' a finite execution, let

$\sigma[..i]$ denote the prefix of σ involving its first i steps

$\tau'\sigma$ denote execution τ' followed by execution σ

and define Π^- to be the set of all finite prefixes of elements in set Π of finite and/or infinite sequences. Then, the above requirement for $\mathcal{P}$—that $\mathcal{P}$ is *prefix closed*—is:

$$(\forall \tau' \in \Psi^- : \neg\hat{\mathcal{P}}(\tau') \Rightarrow (\forall \sigma \in \Psi : \neg\hat{\mathcal{P}}(\tau'\sigma))) \tag{2}$$

Finally, note that any execution rejected by an enforcement mechanism must be rejected after a finite period. This is formalized by:

$$(\forall \sigma \in \Psi : \neg\hat{\mathcal{P}}(\sigma) \Rightarrow (\exists i : \neg\hat{\mathcal{P}}(\sigma[..i]))) \tag{3}$$

Security policies satisfying (1), (2), and (3) are *safety properties* [Lamport 1977], properties stipulating that no "bad thing" happens during any execution. Formally, a property Γ is defined in Lamport [1985] to be a safety property if and only if, for any finite or infinite execution σ,

$$\sigma \notin \Gamma \Rightarrow (\exists i : (\forall \tau \in \Psi : \sigma[..i]\tau \notin \Gamma)) \tag{4}$$

holds. This means that Γ is a safety property if and only if Γ can be characterized using a set of finite executions that are prefixes of all executions excluded from Γ. Clearly, a security policy $\mathcal{P}$ satisfying (1), (2), and (3) has such a set of finite prefixes—the set of prefixes $\tau' \in \Psi^-$ such that $\neg\hat{\mathcal{P}}(\tau')$ holds—so $\mathcal{P}$ is satisfied by sets that are safety properties according to (4).

The above analysis of enforcement mechanisms in EM has established:

Non EM-Enforceable Security Policies: If the set of executions for a security policy $\mathcal{P}$ is not a safety property, then an enforcement mechanism from EM does not exist for $\mathcal{P}$.

Obviously, the contrapositive holds as well: EM enforcement mechanisms enforce security policies that are safety properties. But, as discussed later in Section 4, the converse—that all safety properties have EM enforcement mechanisms—does not hold.

One consequence of our Non EM-Enforceable Security Policies result is that ruling-out additional executions never causes an EM-enforceable policy to be violated, since ruling-out executions never invalidates a safety property. Thus, an EM enforcement mechanism for any security policy $\mathcal{P}'$ satisfying $\mathcal{P}' \Rightarrow \mathcal{P}$ also enforces security policy $\mathcal{P}$. However, a stronger policy $\mathcal{P}'$ might proscribe executions that do not violate $\mathcal{P}$, so using $\mathcal{P}'$ is not without potentially significant adverse consequences. The limit case, where $\mathcal{P}'$ specifies the empty set, illustrates this problem.

Second, our Non EM-Enforceable Security Policies result implies that EM mechanisms compose in a natural way. When multiple EM mechanisms are used in tandem, the policy enforced by the aggregate is the conjunction of the policies that are enforced by each mechanism in isolation. This is attractive because it enables complex policies to be decomposed into conjuncts, with a separate mechanism used to enforce each of the component policies.

Revisiting the three application-independent security policies described in Section 1, we find:

—Access control defines safety properties. The set of proscribed partial executions contains those partial executions ending with an unacceptable operation being attempted.

—Information flow does not define sets that are properties (as argued above), so it does not define sets that are safety properties. Not being safety properties, there are no EM enforcement mechanisms for exactly this policy.[6]

—Availability, if taken to mean that no principal is forever denied use of some given resource, is not a safety property—any partial execution can be extended in a way that allows a principal to access the resource, so the defining set of proscribed partial executions that every safety property must have is absent. In Gligor [1984], availability is defined to rule out all denials in excess of *MWT* seconds (for some predefined Maximum Waiting Time parameter *MWT*). This is a safety property; the defining set of partial executions contains prefixes ending in intervals that exceed *MWT* seconds during which a principal is denied use of a resource.

[6]Mechanisms from EM purporting to prevent information flow do so by enforcing a security policy that implies, but is not equivalent to, the absence of information flow. And, there do exist security policies that both imply restrictions on information flow and define sets that are safety properties.

3. SECURITY AUTOMATA

Enforcement mechanisms in EM work by terminating target execution that is described by a finite prefix σ' such that $\neg\hat{\mathcal{P}}(\sigma')$ holds, for a predicate $\hat{\mathcal{P}}$ defined by the policy being enforced. In addition, we established in Section 2 that the set of executions satisfying $\hat{\mathcal{P}}$ must be a safety property. Those being the only constraints on $\hat{\mathcal{P}}$, we conclude that recognizers for sets of executions that are safety properties can serve as the basis for enforcement mechanisms in EM.

A class of Büchi automata [Eilenberg 1974] that accept safety properties was introduced (although not named) in Alpern and Schneider [1987]. We shall herein refer to these recognizers as *security automata*; they are similar to ordinary non-deterministic finite-state automata [Hopcroft and Ullman 1969]. Formally, a security automaton is defined by:

—a countable set Q of *automaton states*,

—a countable set $Q_0 \subseteq Q$ of *initial automaton states*,

—a countable set I of *input symbols*, and

—a *transition function*[7], δ: $(Q \times I) \rightarrow 2^Q$.

Set I of input symbols is dictated by the security policy being enforced and the manner in which target executions are being represented; the symbols in I might correspond to system states, atomic actions, higher-level actions of the system, or state/action pairs.

To process a sequence $s_1 s_2 \cdots$ of input symbols, the *current state* Q' of the security automaton starts equal to Q_0 and the sequence is read one input symbol at a time. As each input symbol s_i is read, the security automaton changes Q' to

$$\bigcup_{q \in Q'} \delta(q, s_i).$$

If Q' is ever the empty set, then the input is rejected; otherwise the input is accepted. Notice that this acceptance criterion means that a security automaton can accept sequences that have infinite length as well as those having finite length.

Figure 1 depicts a security automaton for a security policy that prohibits execution of *Send* operations after a *FileRead* has been executed. In this diagram, the automaton states are represented by the two nodes labeled q_{nfr} (for "no file read") and q_{fr} (for "file read"). Initial states of the automaton are represented in the diagram by nodes with unlabeled incoming edges, so automaton state q_{nfr} is the only initial automaton state. Transition function δ is specified in terms of edges labeled by *transition*

[7]Notation 2^Q denotes the power set for set Q.

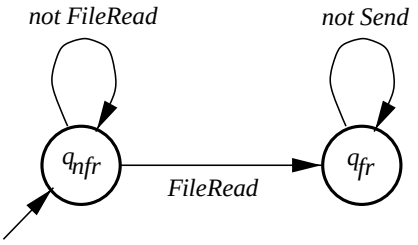

Fig. 1. No *Send* after *FileRead*.

predicates, which are Boolean-valued effectively computable total functions with domain I. Let p_{ij} denote the predicate that labels the edge from node q_i to node q_j. Then, the security automaton, upon reading an input symbol s, changes Q' to

$$\{q_j | q_i \in Q' \wedge p_{ij}(s)\}$$

where $p_{ij}(s)$ is true if and only if input symbol s satisfies predicate p_{ij}.

In Figure 1, transition predicate *not FileRead* is assumed to be satisfied by input symbols (system execution steps) that are not file read operations, and transition predicate *not Send* is assumed to be satisfied by input symbols that are not message-send operations. Since no transition is defined from q_{fr} for input symbols corresponding to message-send execution steps, the security automaton in Figure 1 rejects inputs in which a *Send* follows a *FileRead*.

Diagrams like Figure 1 are impractical to draw and hard to understand if set Q of automaton states is large or transition function δ is complex. We can avoid these difficulties by encoding current state Q' for an automaton in multiple variables and by using *guarded commands* [Dijkstra 1975] to describe the transition function for the security automaton. Guarded command

$$B \rightarrow S \tag{5}$$

specifies that the state transition defined by program fragment S occurs whenever predicate B is satisfied by the current input symbol and the current state of the automaton. In (5), B is called the *guard*, and it is a predicate that can refer only to the current input symbol and to the variables encoding the current state of the automaton; S is called the *command*, and it is a computation that updates (only) the variables encoding the current state of the automaton.

To illustrate this alternative notation for security automata, Figure 2 gives a specification for the same security policy as given in Figure 1. The **state vars** section of this specification introduces the variables that encode the current state of the security automaton. The **transitions** section gives a list of guarded commands that define the transition function. In Figure 2,

$$
\begin{array}{ll}
\textbf{state vars} & state \ : \ \{0,1\} \ \ \textbf{initial}\ 0 \\
\textbf{transitions} & \textbf{not}\ FileRead \wedge state = 0 \longrightarrow \textbf{skip} \\
 & FileRead \wedge state = 0 \longrightarrow state := 1 \\
 & \textbf{not}\ Send \wedge state = 1 \longrightarrow \textbf{skip}
\end{array}
$$

Fig. 2. Alternative specification for policy: No *Send* after *FileRead*.

state—a two-valued variable with initial value 0—encodes the current state of the security automaton, and each of the three guarded commands corresponds to a single edge in the diagram of Figure 1.

As a second example, Figure 3 specifies a security automaton for a simple model of access control introduced in Lampson [1974]. *PRINS* is a universe of principals, *OBJS* is a universe of objects, and *RIGHTS* is a universe of rights. The current state of this security automaton is characterized by a set P of principals, a set O of objects, and a set A of the rights that principals have to objects. A principal p has right r to object o if and only if $\langle p, o, r \rangle \in A$ holds.

Transitions in the security automaton of Figure 3 are specified using predicates defined on the input symbols that correspond to a next step of the target's execution:

Oper(p, o, r): Principal p invoked an operation involving object o and requiring right r to that object.

AddRight(p, p', r', o'): Principal p invoked an operation to add right r' to object o' for principal p'.

RmvRight(p, p', r', o'): Principal p invoked an operation to remove right r' to object o' for principal p'.

AddP(p, p'): Principal p invoked an operation to create a principal named p'.

RmvP(p, p'): Principal p invoked an operation to delete the principal named p'.

AddO(p, o'): Principal p invoked an operation to create an object named o'.

RmvO(p, o'): Principal p invoked an operation to delete object o'.

The transitions specify whether the next step of the target's execution is permitted by the access control policy being defined. The first guarded command asserts that a principal must have the necessary right in order to invoke an operation involving an object implemented by the target. The second and third guarded commands specify a (simplified) policy for granting and revoking rights to principals for objects—the second (third) guarded command asserts that only principals having the cntrl right for an

state vars P : **set of** $PRINS$ **initial** $\emptyset$
O : **set of** $OBJS$ **initial** $\emptyset$
A : **set of** $\langle s : PRINS;\ o : OBJS;\ r : RIGHTS\rangle$ **initial** $\emptyset$

transitions $Oper(p, o, r) \ \wedge \ \langle p, o, r\rangle \in A \longrightarrow$ **skip**

$$AddRight(p, p', r', o') \ \wedge \ \langle p, o', \text{cntrl}\rangle \in A \ \longrightarrow \ A := A \cup \{\langle p', o', r'\rangle\}$$

$$RmvRight(p, p', r', o') \ \wedge \ \langle p, o', \text{cntrl}\rangle \in A \ \longrightarrow \ A := A - \{\langle p', o', r'\rangle\}$$

$$AddP(p, p') \longrightarrow P := P \cup \{p'\}$$
$$O := O \cup \{p'\}$$
$$A := A \cup \{\langle p, p', \text{cntrl}\rangle\}$$

$$RmvP(p, p') \ \wedge \ \langle p, p', \text{cntrl}\rangle \in A \longrightarrow P := P - \{p'\}$$
$$O := O - \{p'\}$$
$$A := A - \{\langle p', \hat{o}, \hat{r}\rangle \mid \hat{o} \in O \wedge \hat{r} \in RIGHTS\}$$

$$AddO(p, o') \longrightarrow O := O \cup \{o'\}$$
$$A := A \cup \{\langle p, o', \text{cntrl}\rangle\}$$

$$RmvO(p, o') \ \wedge \ \langle p, o', \text{cntrl}\rangle \in A \longrightarrow O := O - \{o'\}$$
$$A := A - \{\langle \hat{p}, o', \hat{r}\rangle \mid \hat{p} \in P \wedge \hat{r} \in RIGHTS\}$$

Fig. 3. Access control.

object o can grant (remove) rights to other principals for accessing o. The remaining four guarded commands specify a policy regarding creation and deletion of principals and objects:

—Every principal is also an object.

—The principal that creates a principal (object) is given the cntrl right for that principal (object).

—A principal must have the corresponding cntrl right in order to delete a principal or object.

It ought to be clear that more realistic policies are easily accommodated by modifying the guarded commands of Figure 3.

Two things are worth noting about this access control example. First, leverage results from employing a suitable representation (namely sets P, O, and A which together encode an access control matrix) for the current state of the automaton. Imagine how awkward it would be to try and describe changes to access rights in terms of a flat set of uninterpreted automaton states. Second, our method of specifying security policies allows—but does not force—a distinction between security-configuration changes (i.e., changing A when access rights are added and deleted) and accesses to objects implemented by the target system. And, we would argue that there is no value in forcing such a distinction, although this view is not universally held [Gligor et al. 1998].

state vars $state : \{0, 1\}$ **initial** 0

transitions **not** $Pay(C) \wedge state = 0 \longrightarrow$ **skip**

$Pay(C) \wedge state = 0 \longrightarrow state := 1$

$Serve(C) \wedge state = 1 \longrightarrow state := 0$

Fig. 4. Security automaton for fair transaction.

As a final illustration, we turn to electronic commerce. We might, for example, desire that a service provider be prevented from engaging in actions other than delivering service for which a customer has paid. This requirement is a security policy; it can be formalized in terms of the following predicates on input symbols, if input symbols represent operation executions:

$pay(C)$: customer C requests and pays for service

$serve(C)$: customer C is rendered service

The security policy of interest proscribes executions in which the service provider executes an operation that does not satisfy $serve(C)$ after having engaged in an operation that satisfies $pay(C)$. A security automaton for this policy is defined in Figure 4.

Notice, the security automaton of Figure 4 does not stipulate that payment guarantees service—it only limits what the service provider can do once a customer has made payment. In particular, the security policy that is specified allows a service provider to stop executing (i.e., stop producing input symbols) rather than rendering a paid-for service. We cannot specify the stronger security policy (that service be guaranteed after payment) because that is not a safety property—there is no defining set of proscribed partial executions.

4. USING SECURITY AUTOMATA FOR ENFORCEMENT

Any security automaton can serve as the basis for an enforcement mechanism in EM. The target is executed in tandem with a simulation of the security automaton.[8] In particular, initialization or creation of the target causes an initialized instance of the security automaton simulation to be created. And, each step that the target is about to take generates an input symbol, which is sent to that simulation:

(i) If the automaton can make a transition on that input symbol, then the target is allowed to perform that step and the automaton state is changed according to its transition function.

[8]A similar approach—developed independently—for integrating software components whose behaviors need to be reconciled is outlined in Marchukov and Sullivan [1999].

(ii) If the automaton cannot make a transition on that input symbol, then the target is terminated (for having attempted to violate the security policy).

Implicit in this approach are some assumptions.

Bounded Memory. The memory that can be devoted to simulating a security automaton will, of necessity, be finite—real computers have finite memories. Recall from Section 3 that our security automata can have an infinite (countable) number of automaton states.

Infinite sets of automaton states are necessary for recognizing certain safety properties, because whether a given prefix should be rejected might depend on all of the input symbols in that prefix. The ever-larger prefixes produced as execution proceeds thus require ever-larger sets of states to encode needed information about the past. For example, a safety property stipulating that, at each step of execution, the value of some target variable x equals the sum of its values in preceding states requires (to store the sum of the past values of x) a state variable that grows without bound.

Security policies of concern in real systems do not seem to require large amounts of storage and, in fact, are enforced today using mechanisms that use only modest amounts of storage; a security automaton to specify such a policy would also require only a modest-sized set of automaton states. We see no reason to expect application-specific or special-purpose security policies to be different. So, restricting the **state vars** for a security automaton to a finite amount of storage is not, in practice, a limitation.

Target Control. Implicit in (ii) is the assumption that the target can be terminated by the enforcement mechanism. Specifically, we assume that the enforcement mechanism has sufficient control over the target to stop further automaton input symbols from being produced. This control requirement is subtle and makes certain security policies—even though they characterize sets that are safety properties—unenforceable using mechanisms from EM.

For example, recall from Section 2 the definition of availability in Gligor [1998]:

> **Real-Time Availability:** One principal cannot be denied use of a resource for more than *MWT* seconds.

Sets satisfying Real-Time Availability are safety properties—the "bad thing" is an interval of execution spanning more than *MWT* seconds during which some principal is denied the resource. The input symbols of a security automaton for Real-Time Availability will therefore encode time, and a new input symbol is produced whenever time increases.

While individual clocks might be stopped, the passage of time cannot be stopped. So the target cannot be stopped from producing input symbols as time passes. Real-Time Availability simply cannot be enforced by running an automaton simulation in tandem with a target, because targets cannot provide the necessary controls to the enforcement mechanism. And since

the other mechanisms in EM are no more powerful, we conclude that Real-Time Availability cannot be enforced using any mechanism in EM. Change the specification from "*MWT* seconds" to "*MWT* execution steps" and the target can be prevented from violating the policy by stopping execution, resulting in an EM-enforceable security policy.

Enforcement Mechanism Integrity. A target that corrupts a security automaton simulation can subvert an enforcement mechanism built on that simulation. Input to the enforcement mechanism must correspond to target execution, and state transitions must follow the automaton's transition function. Ensuring that input to the enforcement mechanism is both correct and complete is a question of target instrumentation and monitoring. The "complete mediation" requirement associated with reference monitors is one way to discharge this assumption. Ensuring that the target does not interfere with automaton transitions is a matter of isolation—the enforcement mechanism must be isolated from the target. Isolation of our enforcement mechanism is accomplished if, for example, the **state vars** and **transitions** for the security automaton are not writable by the target.

Automaton Simulation Pragmatics

Two mechanisms are involved in the above security-automaton based implementation of an enforcement mechanism.

Automaton Input Read: A mechanism to determine that an input symbol has been produced by the target and then to forward that symbol to the security automaton simulation.

Automaton Transition: A mechanism to determine whether the security automaton can make a transition on a given input and then to perform that transition.

How these are implemented determines the cost of the enforcement mechanism. For example, when the automaton's input symbols are the set of target states and its transition predicates are arbitrary state predicates, a new input symbol is produced each time any component of the target's state changes. Since the program counter is a state component and it changes each time a machine-language instruction is executed or an interrupt occurs, the enforcement mechanism must be involved in executing each target instruction. That could be quite costly.

For security policies where the target's production of automaton input symbols coincides with occurrences of hardware traps, an automaton-based enforcement mechanism can be supported quite cheaply by incorporating it into the trap handler. One example is implementing an enforcement mechanism for access control policies on operating system objects, such as files. Here, the target is a file and the production of input symbols coincides with invocations of system operations (i.e., file access operations). The production of input symbols now coincides with occurrences of system-call traps.

A second example where hardware traps can be exploited arises in implementing memory protection. A security automaton for a typical memory protection policy would expect an input symbol for each memory reference. But most of these input symbols would cause no change to the security automaton's state. Input symbols that do not cause automaton state transitions need not be forwarded to the automaton, and that justifies the following optimization of Automaton Input Read:

> **Automaton Input Read Optimization:** Input symbols are not forwarded to the security automaton if the state of the automaton just after the transition would be the same as it was before the transition.

Given this optimization, the production of automaton input symbols for memory protection can be made to coincide with occurrences of traps. The target's memory-protection hardware—base/bounds registers or page and segment tables—is initialized so that a trap occurs when an input symbol should be forwarded to the memory protection automaton. Memory references that do not cause traps are effectively filtered and thus never cause a state transition or undefined transition by the automaton. Note, however, if this optimization is used, then a target can subvert the enforcement mechanism by corrupting the filter that selects whether to forward an input symbol to the security automaton.

Finally, inexpensive implementation of our automata-based enforcement mechanisms is also possible when programs are executed by a software-implemented virtual machine. The virtual machine instruction-processing cycle is augmented so that it produces input symbols and makes automaton transitions according to either an internal or an externally specified security automaton. For example, the Java virtual machine [Lindholm and Yellin 1997] could easily be augmented to implement the Automaton Input Read and Automaton Transition Mechanisms for input symbols that correspond to method invocations.

Beyond EM Enforcement Mechanisms

Response to Violations. Termination of a target that is about to violate a security policy might seem draconian. Yet, by definition, this is how an EM mechanism responds to an attempted violation. Why not simply notify the target that an erroneous execution step has been attempted? The target could then substitute another step and its execution might then continue.

In terms of our security automata framework, notifying a target is equivalent to having the security automaton extend that target's execution (rather than truncating that execution). And some—but not all—security policies do allow input prefixes to be extended in this manner. A security policy that does not enjoy this attribute is the variant of Real-Time Availability given in Section 2 where *MWT* bounds the number of execution steps (not seconds) that elapse before an action is taken. Various other safety properties also do not allow execution prefixes to be extended,

although their practical significance as security policies is an open question.

EM was defined to truncate execution for generality. Expanding EM to include enforcement mechanisms that handle violations by notifying the target or by truncating its execution would not change the set of security policies that are EM enforceable. Modifying EM to require enforcement mechanisms that handle violations by necessarily notifying the target would shrink the set of security policies that are EM enforceable, and with no apparent gain.

Program Modification. The overhead of enforcement can be reduced by merging the enforcement mechanism into the target. One such scheme is *software-based fault isolation* (SFI), also known as "sandboxing" [Wahbe et al. 1993; Small 1997]. SFI implements memory protection but does so without hardware assistance. Instead, a program is edited before it is executed, and only such edited programs are executed by the target. (Usually, it is the object code that is edited.) The edits insert instructions to check and/or modify the values of operands, so that illegal memory references are never attempted.

SFI is not in EM because SFI involves modifying the target, and such modifications are not permitted of enforcement mechanisms in EM. But viewed in our framework, the inserted instructions for SFI can be seen to implement Automaton Input Read by copying code for Automaton Transition in line before each target instruction that produces an input symbol. Generalizing, nothing prevents the SFI approach from being used with arbitrary security automata, thereby enforcing any EM-enforceable security policy. Trust must be placed in the tools used to modify the target, however.

Our SASI (Security Automata SFI Implementation) prototypes for Intel's x86 object code and SUN's JVM (Java Virtual Machine) explored the use of an SFI-like approach for EM-enforceable policies [Erlingsson and Schneider 1999]. Each of our prototypes merges the simulation of a security automaton into the object code for the program that is the target. New variables—accessible only to the code added for SASI—represent the current state of a security automaton, and new code—that cannot be circumvented—simulates automaton state transitions. The new code also causes the target system to halt whenever the automaton rejects its input (because the current automaton state does not allow a transition for the next target instruction). Analysis of a target allows simplification of the code inserted for simulating a security automaton. Each inserted copy of the automaton simulation is a candidate for simplification based on the context in which that code appears. By using partial evaluation [Jones et al. 1993] on the guards as well as by using the automaton structure, irrelevant tests and updates to the security automaton state can be removed.

Program Analysis. There is no need for any run-time enforcement mechanism if the target can be analyzed and proved not to violate the security policy of interest. This approach has been employed for a security policy

like what SFI was originally intended to address in *proof carrying code* (PCC) [Necula 1997]. With PCC, a proof is supplied along with a program, and this proof comes in a form that can be checked mechanically before running that program. The security policy will not be violated if, before the program is executed, the accompanying proof is checked and found to be correct. The original formulation of PCC required that proofs be constructed by hand. This restriction can be relaxed. For certain security policies, a compiler can automatically produce PCC from programs written in high-level, type-safe programming languages [Morrisett et al. 1998; Necula and Lee 1998].

To extend PCC for security policies that are specified by arbitrary security automata, a method is needed to extract proof obligations for establishing that a program satisfies the property given by such an automaton. Such a method does exist—it is described in Alpern and Schneider [1989].

5. DISCUSSION

The utility of a formalism partly depends on the ease with which it can be read and written. Users of the formalism must be able to translate informal requirements into the formalism. With security automata, establishing the correspondence between transition predicates and informal requirements on system behavior is crucial and can require a detailed understanding of the target. The automaton of Figure 1, for example, only captures the informal requirement that messages are not sent after a file is read if it is impossible to send a message unless transition predicate *Send* is *true* and it is impossible to read a file unless transition predicate *FileRead* is *true*. There might be many ways to send messages—some obvious and others buried deep within the bowels of the target. All must be identified and included in the definition of *Send*; a similar obligation accompanies transition predicate *FileRead*.

The general problem of establishing the correspondence between informal requirements and some purported formalization of those requirements is not new to software engineers. The usual solution is to analyze the formalization, being alert to inconsistencies between the results of the analysis and the informal requirements. We might use a formal logic to derive consequences from the formalization; we might use partial evaluation to analyze what the formalization implies about one or another scenario, a form of testing; or, we might (manually or automatically) transform the formalization into a prototype and observe its behavior in various scenarios.

Success with proving, testing, or prototyping as a way to gain confidence in a formalization depends upon two things. The first is deciding what aspects of a formalization to check, and this is largely independent of the formalism. But the second, having the means to do those checks, not only depends on the formalism but largely determines the usability of that formalism. To do proving, we require a logic whose language includes the

formalism; to do testing, we require a means of evaluating a formalization in one or another scenario; and to do prototyping, we must have some way to transform a formalization into a computational form.

As it happens, a rich set of analytical tools does exist for security automata, because security automata are a class of Büchi automata that are widely used in computer-aided program verification tools. Existing formal methods based either on model checking or on theorem proving can be employed to analyze a security policy that has been specified as a security automaton. And, testing or prototyping a security policy that is specified by a security automaton is just a matter of running the automaton.

Guidelines for Structuring Security Automata

Real system security policies are best given as collections of simpler policies, a single large monolithic policy being difficult to comprehend. The system's security policy is then the result of composing the simpler policies in the collection by taking their conjunction. To employ such a separation of concerns when security policies are specified by security automata, we must be able to compose security automata in an analogous fashion. Given a collection of security automata, we must be able to construct a single *conjunction security automaton* for the conjunction of the security policies specified by the automata in the collection. That construction is not difficult: An execution is rejected by the conjunction security automaton if and only if it is rejected by any automaton in the collection.

Beyond comprehensibility, there are other advantages to specifying system security policies as collections of security automata. First, having a collection allows different enforcement mechanisms to be used for the different automata (hence the different security policies) in the collection. Second, security policies specified by distinct automata can be enforced by distinct system components, something that is attractive when all of a given security automaton's input symbols correspond to events at a single system component. Benefits that accrue from having the source of all of an automaton's input symbols be a single component include:

—Enforcement of a component's security policy involves trusting only that component.

—The overhead of an enforcement mechanism is lower because communication between components can be reduced.

For example, the security policy for a distributed system might be specified by giving a separate security automaton for each system host. Then, each host would itself implement the Automaton Input Read and Automaton Transitions mechanisms for only the security automata concerning that host.

Application to Safety-Critical Systems

The idea that security kernels might have application in safety-critical systems is eloquently justified in Rushby [1989] and continues to interest

researchers such as Wika and Knight [1995]. Safety-critical systems are, for the most part, concerned with enforcing properties that are safety properties (in the sense of Lamport [1985]), so it is natural to expect an enforcement mechanism for safety properties to have application in this class of systems. And, we see no impediments to using security automata or our security-automata based enforcement mechanisms for enforcing safety properties in safety-critical systems.

The justification given in Rushby [1989] for using security kernels in safety-critical systems involves a characterization of what types of properties can be enforced by a security kernel. As do we in this paper, Rushby [1989] concludes that safety properties but not liveness properties[9] are enforceable. However, the arguments given in Rushby [1989] are informal and are coupled to the semantics of kernel-supported operations. The essential attributes of enforceability, which we isolate and formalize by equations (1), (2), and (3), are neither identified nor shown to imply that only safety properties can be enforced.

In addition, because Rushby [1989] concerns kernelized systems, the notion of property is restricted there to being sequences of kernel-provided functions. By allowing security automata to have arbitrary sets of input symbols, our results can be seen as generalizing those of Rushby [1989]. And the generalization is a useful one, because it applies to enforcement mechanisms that are not part of a kernel. Thus, we can now extend the central thesis of Rushby [1989], that kernelized systems have application beyond implementing security policies, to justify the use of enforcement mechanisms from EM when building safety-critical systems.

APPENDIX: SUMMARY OF NOTATION

Ψ: The set of all finite and infinite sequences.

S: A target.

Σ_S: The set of executions possible by target S.

$\mathcal{P}$: A predicate specifying a security policy.

Σ: A set of executions.

Π: A set of executions.

$\hat{\mathcal{P}}$: A predicate on executions used in defining security policy $\mathcal{P}$.

σ: A finite or infinite execution.

σ': A finite execution.

τ: A finite or infinite execution

[9]A *liveness property* is a property that stipulates some "good thing" happens during any execution. See Alpern and Schneider [1987] for a formal definition.

τ': A finite execution.

$\sigma[..i]$: The prefix of σ involving its first i steps.

$\tau'\sigma$: Finite execution τ' followed by execution σ.

Π^-: The set of all finite prefixes of elements in set Π.

Γ: A set of executions that is a safety property.

Q: The set of automaton states.

Q_0: The set of initial automaton states.

I: The set of automaton input symbols.

δ: The automaton next-state transition function.

Q': The current state of a security automaton.

ACKNOWLEDGMENTS

I am grateful to Robbert van Renesse, Greg Morrisett, Úlfar Erlingsson, Yaron Minsky, and Lidong Zhou for helpful feedback on the use and implementation of security automata and for comments on previous drafts of this paper. Helpful comments on earlier drafts of this paper were also provided by Earl Boebert, Dave Evans, Li Gong, Robert Grimm, Keith Marzullo, Andrew Myers, John Rushby, and Chris Small. John McLean served as a valuable sounding board for these ideas as I developed them. Feedback from Martin Abadi helped to sharpen the formalism. Virgil Gligor wrote a lengthy review that helped in refining my arguments. Finally, the University of Tromso was a hospitable setting and a compelling excuse for performing some of the work reported herein.

REFERENCES

ALPERN, B. AND SCHNEIDER, F. B. 1985. Defining liveness. *Inf. Process. Lett. 21*, 4 (Oct.), 181–185.

ALPERN, B. AND SCHNEIDER, F. B. 1987. Recognizing safety and liveness. *Distrib. Comput. 2*, 117–126.

ALPERN, B. AND SCHNEIDER, F. B. 1989. Verifying temporal properties without temporal logic. *ACM Trans. Program. Lang. Syst. 11*, 1 (Jan. 1989), 147–167.

DIJKSTRA, E. W. 1975. Guarded commands, nondeterminacy, and formal derivation of programs. *Commun. ACM 18*, 8 (Aug.), 453–475.

EDJLALI, G., ACHARYA, A., AND CHAUDHARY, V. 1998. History-based access control for mobile code. In *Proceedings of the 5th ACM Conference on Computer and Communications Security* (CCS '98, San Francisco, CA, Nov. 3–5), L. Gong and M. Reiter, Eds. ACM Press, New York, NY, 38–48.

EILENBERG, S. 1974. *Automata, Languages, and Machines, Volume A*. Academic Press, Inc., New York, NY.

ERLINGSSON, U. AND SCHNEIDER, F. B. 1999. SASI enforcement of security policies: A retrospective. In *Proceedings of the on New Security Paradigms Workshop* (Ontario, Canada, Sept.), ACM Press, New York, NY, 87–95.

EVANS, D. AND TWYMAN, A. 1999. Policy-directed code safety. In *Proceedings of the 1999 IEEE Computer Society Symposium on Research in Security and Privacy* (Oakland, CA, May), IEEE Computer Society Press, Los Alamitos, CA, 32–45.

GLIGOR, V. D. 1984. A note on denial-of-service in operating systems. *IEEE Trans. Softw. Eng. SE-10*, 3 (May 1984), 320–324.

GLIGOR, V. D., GAVRILA, S., AND FERRAIOLO, D. 1998. On the formal definition of separation-of-duty policies and their composition. In *Proceedings of the 1998 IEEE Computer Society Symposium on Research in Security and Privacy* (Oakland, CA, May), IEEE Computer Society Press, Los Alamitos, CA, 172–183.

GOGUEN, J. A. AND MESEGUER, J. 1982. Security policies and security models. In *Proceedings of the 1982 IEEE Computer Society Symposium on Research in Security and Privacy* (Oakland, CA, May), IEEE Computer Society Press, Los Alamitos, CA, 11–20.

GONG, L. 1997. Java security: Present and near future. *IEEE Micro 17*, 3 (May/June), 14–19.

GRIMM, R. AND BERSHAD, B. N. 1999. Providing policy-neutral and transparent access control in extensible systems. In *Secure Internet Programming: Security Issues for Mobile and Distributed Objects*, J. Vitek and C. Jensen, Eds. Lecture Notes in Computer Science, vol. 1603. Springer-Verlag, New York, NY, 317–338.

HOPCROFT, J. AND ULLMAN, J. 1969. *Formal Languages and Their Relation to Automata*. Addison-Wesley, Reading, MA.

JAJODIA, S., SAMARATI, P., AND SUBRAHMANIAN, V. S. 1997. A logical language for expressing authorizations. In *Proceedings of the 1997 IEEE Computer Society Symposium on Reasearch in Security and Privacy* (Oakland, CA, May), IEEE Computer Society Press, Los Alamitos, CA, 31–42.

JONES, N. D., GOMARD, C. K., AND SESTOFT, P. 1993. *Partial Evaluation and Automatic Program Generation*. Prentice-Hall International Series in Computer Science. Prentice-Hall, Inc., Upper Saddle River, NJ.

LAMPORT, L. 1977. Proving the correctness of multiprocess programs. *IEEE Trans. Softw. Eng. 3*, 2 (Mar.), 125–143.

LAMPORT, L. 1985. Logical foundation. In *Distributed systems: Methods and tools for specification. An advanced course*, M. W. Alford, J. P. Ansart, G. Hommel, L. Lamport, B. Liskov, G. P. Mullery, F. B. Schneider, M. Paul, and H. J. Siegert, Eds, Springer Lecture Notes in Computer Science, vol. 190. Springer-Verlag, New York, NY, 119–130.

LAMPSON, B. 1974. Protection. In *Proceedings of the 5th Symposium on Information Sciences and Systems* (Princeton, NJ, Mar.), 437–443.

LINDHOLM, T. AND YELLIN, F. 1997. *The Java Virtual Machine Specification*. Addison-Wesley, Reading, MA.

MARCHUKOV, M. AND SULLIVAN, K. 1999. Reconciling behavioral mismatch through component restriction. CS 99-22. Department of Computer Science, University of Virginia, Charlottesville, VA. Technical Report

MCLEAN, J. 1994. A general theory of composition for trace sets closed under selective interleaving functions. In *Proceedings of the 1994 IEEE Computer Society Symposium on Research in Security and Privacy* (Oakland, CA, May), IEEE Computer Society Press, Los Alamitos, CA, 79–93.

MORRISETT, G., WALKER, D., CRARY, K., AND GLEW, N. 1998. From system F to typed assembly language. In *Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (POPL '98, San Diego, CA, Jan. 19–21), D. B. MacQueen and L. Cardelli, Eds. ACM Press, New York, NY, 85–97.

NECULA, G. C. 1997. Proof-carrying code. In *Proceedings of the 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (POPL '97, Paris, France, Jan. 15–17, 1997), P. Lee, Ed. ACM Press, New York, NY, 106–119.

NECULA, G. C. AND LEE, P. 1998. The design and implementation of a certifying compiler. In *Proceedings of the ACM SIGPLAN '98 Conference on Programming Language Design and Implementation* (Montreal, Que., Canada, June), *SIGPLAN Not. 33*, 5, 333–344.

NULL, L. M. AND WONG, J. 1992. The DIAMOND security policy for object-oriented databases. In *Proceedings of the 1992 ACM Computer Science 20th Annual Conference on*

Communications (CSC '92, Kansas City, MO, Mar. 3–5), J. P. Agrawal, V. Kumar, and V. Wallentine, Eds. ACM Press, New York, NY, 49–56.

PANDEY, R. AND HASHII, B. 1998. Providing fine grained access control for mobile programs through binary editing. TR98 08. Department of Computer Science, University of California at Davis, Davis, CA.

RUSHBY, J. 1989. Kernels for safety? In *Safe and Secure Computing Systems*, T. Anderson, Ed. Blackwell Scientific Publications, Ltd., Oxford, UK, 210–220.

SALTZER, J. H. AND SCHROEDER, M. D. 1975. The protection of information in computer systems. *Proc. IEEE 63*, 9 (Sept.), 1278–1308.

SMALL, C. 1997. Misfit: A tool for constructing safe extensible C++ systems. In *Proceedings of the 3rd USENIX Conference on Object-Oriented Technologies* (Portland, OR, June), USENIX Assoc., Berkeley, CA, 38–48.

STEFIK, M. 1996. Letting loose the light: Igniting commerce in electronic publication. In *Internet Dreams*, M. Stefik, Ed. MIT Press, Cambridge, MA.

WAHBE, R., LUCCO, S., ANDERSON, T. E., AND GRAHAM, S. L. 1993. Efficient software-based fault isolation. In *Proceedings of the 14th ACM Symposium on Operating Systems Principles* (Asheville, NC, Dec. 5–8), A. P. Black and B. Liskov, Eds. ACM Press, New York, NY, 202–216.

WIKA, K. G. AND KNIGHT, J. C. 1995. On the enforcement of software safety policies. In *Proceedings of the 10th Annual IEEE Conference on Computer Assurance* (COMPASS '95, Gaithersburg, MD, June), IEEE Computer Society Press, Los Alamitos, CA.

WOO, T. Y. C. AND LAM, S. S. 1992. Authorization in distributed systems: A formal approach. In *Proceedings of the ACM/IEEE Symposium on Research in Security and Privacy* (Oakland, CA, May), IEEE Computer Society Press, Los Alamitos, CA, 33–50.

Received: January 1998; revised: July 1999; accepted: October 1999