

Edit Automata: Enforcement Mechanisms for Run-time Security Policies *

Jay Ligatti Lujo Bauer David Walker
Department of Computer Science
Princeton University
Princeton, NJ 08544

Technical Report TR-681-03

May 30, 2003

Abstract

We analyze the space of security policies that can be enforced by monitoring and modifying programs at run time. Our program monitors, called *edit automata*, are abstract machines that examine the sequence of application program actions and transform the sequence when it deviates from a specified policy. Edit automata have a rich set of transformational powers: They may terminate the application, thereby truncating the program action stream; they may suppress undesired or dangerous actions without necessarily terminating the program; and they may also insert additional actions into the event stream.

After providing a formal definition of edit automata, we develop a rigorous framework for reasoning about them and their cousins: *truncation automata* (which can only terminate applications), *suppression automata* (which can terminate applications and suppress individual actions), and *insertion automata* (which can terminate and insert). We give a set-theoretic characterization of the policies each sort of automaton can enforce and we provide examples of policies that can be enforced by one sort of automaton but not another.

1 Introduction

When designing a secure, extensible system such as an operating system that allows applications to download code into the kernel or a database that allows users to submit their own optimized queries, we must ask two important questions.

*This is a revised and extended version of “More Enforceable Security Policies,” a paper that first appeared in the Workshop on Foundations of Computer Security, June 2002 [BLW02].

1. What sorts of security policies can we expect our system to enforce?
2. What sorts of mechanisms do we need to enforce these policies?

Neither of these questions can be answered effectively without understanding the space of enforceable security policies and the power of various enforcement mechanisms.

The first significant effort to define the class of enforceable security policies is due to Schneider [Sch00]. He investigated the security properties that can be enforced by a specific type of program monitor. One of Schneider’s monitors can interpose itself between an untrusted program and the machine on which the program runs. The monitor can examine the sequence of security-relevant program actions one at a time and if it recognizes an action that will violate its policy, the monitor terminates the program. This mechanism is very general since decisions about whether or not to terminate the program can depend upon the entire history of the program’s execution. However, since these monitors can only terminate programs and cannot otherwise modify their behavior, it is possible to define still more powerful enforcement mechanisms.

In this paper, we re-examine the question of which security policies can be enforced at run time by monitoring untrusted programs. Our overall approach differs from Schneider’s, who also used automata to model program monitors, in that we view program monitors as transformers that *edit* the stream of actions produced by an untrusted application. This new viewpoint leads us to define a hierarchy of enforcement mechanisms, each with different transformational capabilities:

- A *truncation automaton* can recognize bad sequences of actions and halt program execution before the security policy is violated, but cannot otherwise modify program behavior. These automata are similar to Schneider’s original security monitors.
- A *suppression automaton*, in addition to being able to halt program execution, has the ability to suppress individual program actions without terminating the program outright.
- An *insertion automaton* is able to insert a sequence of actions into the program action stream as well as terminate the program.
- An *edit automaton* combines the powers of suppression and insertion automata. It is able to truncate action sequences and insert or suppress security-relevant actions at will.

We use the general term *security automaton* to refer to any automaton that is used to model a program monitor, including the automata mentioned above.¹

The main contribution of this article is the development of a robust theory for reasoning about these machines under a variety of different conditions. We

¹Previous authors [Sch00] have used the term to refer specifically to automata with powers similar to our truncation automata, which we discuss in Section 3.

use our theory to characterize the class of security policies that can be enforced by each sort of automaton, and we provide examples of security policies that lie in one class but not another.

More important than any particular result is our methodology, which gives rise to straightforward, rigorous proofs concerning the power of security mechanisms and the range of enforceable security policies. This overall methodology can be broken down into four main parts.

- Step 1. Define the underlying computational framework and the range of security policies that will be considered. In this paper, we define the software systems and sorts of policies under consideration in Section 2.1 through Section 2.3.
- Step 2. Specify what it means to enforce a security policy. As we will see in Section 2.4, there are several choices to be made in this definition. One must be sure that the enforcement model accurately reflects the desires of the system implementer and the environment in which the monitor operates. Section 2.5 explains some of the limitations induced by our decisions in steps 1 and 2.
- Step 3. Formally specify the operational behavior of the enforcement mechanism in question. Sections 3, 4, 5 and 6 define the operational semantics of four different sorts of monitors and provides examples of the policies that they can enforce.
- Step 4. Prove from the previous definitions that the security mechanism in question is able to enforce the desired properties. Sections 3, 4, 5 and 6 state theorems concerning the security policies that each type of monitor can enforce. The formal proofs can be found in the Appendix A.

After completing our analysis of edit automata and related machines, we discuss related work (Section 7). Finally, Section 8 concludes the paper with a taxonomy of security policies and a discussion of some unanswered questions and our continuing research.

2 Security Policies and Enforcement Mechanisms

In this section, we define the overarching structure of the secure systems we intend to explore. We also define what it means to be a security policy, and what it means to enforce a security policy. Finally, we give a generic definition of a security automaton as an action sequence transformer.

2.1 Systems, Executions, and Policies

We specify software systems at a high level of abstraction. A *system* $S = (\mathcal{A}, \Sigma)$ is specified via a set of *program actions* \mathcal{A} (also referred to as program events) and a set of possible executions Σ . An *execution* σ is simply a finite sequence of actions a_1, a_2, \dots, a_n . Previous authors have considered infinite executions as well as finite ones [Sch00]. Some of the applications on which we might want to enforce policies (such as web servers or operating systems) are often considered to run infinitely, but in practice their executions will always eventually terminate. Although we allow \mathcal{A} and Σ to be countably infinite, in this paper we restrict ourselves to finite but arbitrarily long executions to simplify our analysis. We use the metavariables σ and τ to range over finite sequences of actions.

The symbol \cdot denotes the empty sequence. We use the notation $\sigma[i]$ to denote the i^{th} action in the sequence (beginning the count at 0). The notation $\sigma[..i]$ denotes the subsequence of σ involving the actions $\sigma[0]$ through $\sigma[i]$, and $\sigma[i+1..]$ denotes the subsequence of σ involving all other actions. We use the notation $\tau;\sigma$ to denote the concatenation of two sequences. When τ is a prefix of σ we write $\tau \preceq \sigma$. Given a set of executions Σ , $pre(\Sigma)$ is the set of all prefixes of all executions in Σ .

In this work, it will be important to distinguish between uniform systems and nonuniform systems. (\mathcal{A}, Σ) is a *uniform* system if $\Sigma = \mathcal{A}^*$ where \mathcal{A}^* is the set of all finite sequences of symbols from \mathcal{A} . Conversely, (\mathcal{A}, Σ) is a *nonuniform* system if $\Sigma \subset \mathcal{A}^*$. Uniform systems arise naturally when a program is completely unconstrained; unconstrained programs may execute operations in any order. However, an effective security system will often combine static program analysis and preprocessing with run-time security monitoring. Such is the case in Java virtual machines, for example, which combine type checking with stack inspection. Program analysis, preprocessing, model checking, control- or data-flow analysis, program instrumentation, type checking, and proof-carrying code can also give rise to nonuniform systems.

A *security policy* is a predicate P on sets of executions. A set of executions Σ satisfies a policy P if and only if $P(\Sigma)$. Most common extensional program properties fall under this definition of security policy, including the following.

- *Access Control* policies specify that no execution may operate on certain resources such as files or sockets, or invoke certain system operations.
- *Availability* policies specify that if a program acquires a resource during an execution, then it must release that resource at some (arbitrary) later point in the execution.
- *Bounded Availability* policies specify that if a program acquires a resource during an execution, then it must release that resource by some fixed point later in the execution. For example, the resource must be released in at most ten steps or after some system invariant holds. We call the condition that demands release of the resource the *bound* for the policy.

- An *Information Flow* policy concerning inputs s_1 and outputs s_2 might specify that if $s_2 = f(s_1)$ in one execution (for some function f) then there must exist another execution in which $s_2 \neq f(s_1)$.

2.2 Security Properties

Alpern and Schneider [AS87] distinguish between *properties* and more general policies as follows. A security policy P is deemed to be a (computable) *property* when it has the following form.

$$P(\Sigma) = \forall \sigma \in \Sigma. \hat{P}(\sigma) \quad (\text{PROPERTY})$$

where \hat{P} is a computable predicate on \mathcal{A}^* .

Hence, a property is defined exclusively in terms of individual executions. It may not specify a relationship between possible executions of the program. Information flow, for example, which can only be specified as a condition on the set of possible executions of a program, is not a property. The other example policies provided in the previous section are all security properties.

We assume that the empty sequence is contained in any property. This describes the idea that an untrusted program that has not started executing is not yet in violation of any property. From a technical perspective, this decision allows us to avoid repeatedly considering the empty sequence as a special case of an execution sequence in future definitions of enforceable properties.

Given some set of actions \mathcal{A} , a predicate \hat{P} over \mathcal{A}^* induces the security property $P(\Sigma) = \forall \sigma \in \Sigma. \hat{P}(\sigma)$. We often use the symbol \hat{P} interchangeably as a predicate over execution sequences and as the induced property. Normally, the context will make clear which meaning we intend.

Safety Properties Properties that specify that “nothing bad ever happens” are called *safety properties* [Lam77]. We can make this definition precise as follows. Predicate \hat{P} induces a safety property if and only if,

$$\forall \sigma \in pre(\Sigma). \neg \hat{P}(\sigma) \Rightarrow \forall \sigma' \in \Sigma. (\sigma \preceq \sigma' \Rightarrow \neg \hat{P}(\sigma')) \quad (\text{SAFETY})$$

Informally, this definition states that once a bad action has taken place, thereby excluding the initial segment of an execution from the property, there is no extension of that segment that can remedy the situation. For example, access-control policies are safety properties since once a restricted resource has been accessed, the policy is broken. There is no way to “un-access” the resource and fix the situation afterwards.

Our definition of safety differs slightly from that of previous authors. Since we wish to consider nonuniform systems, σ ranges over $pre(\Sigma)$ rather than Σ . On uniform systems $\Sigma = \mathcal{A}^*$ and therefore $pre(\Sigma) = \Sigma$; consequently, the definition we give corresponds exactly to previous work. On nonuniform systems $pre(\Sigma)$ is a superset of Σ . In our definition of safety, this implies that a sequence may become irremediably bad at a point that does not correspond to a full execution sequence.

Liveness Properties In contrast to safety properties, *liveness properties* [AS85] state that nothing exceptionally bad can happen in any finite amount of time. Any finite sequence of actions can always be extended so that it lies within the property. Formally, predicate \hat{P} induces a liveness property if and only if,

$$\forall \sigma \in pre(\Sigma). \exists \sigma' \in \Sigma. (\sigma \preceq \sigma' \wedge \hat{P}(\sigma')) \quad (\text{LIVENESS})$$

On uniform systems availability is a liveness property. If the program has acquired a resource, we can always extend its execution so that it releases the resource in the next step. As with safety, considering nonuniformity requires us to use a modified definition of liveness that lets σ range over $pre(\Sigma)$ rather than Σ .

Because we are limiting ourselves to considering only finite sequences, our definition of liveness excludes many interesting liveness properties, such as termination. We present the definition of liveness for the sake of comparison with safety.

Other Properties Alpern and Schneider [AS87] show that any property can be decomposed into a combination of a safety property and a liveness property. Bounded availability, for example, is a property that combines safety and liveness on uniform systems. Suppose our bounded-availability policy states that every resource that is acquired must be released and must be released at most ten steps after it is acquired. This property contains an element of safety because there is a bad thing that may occur (e.g., taking eleven steps without releasing the resource). It is not purely a safety property because there are sequences that are not in the property (e.g., we have taken eight steps without releasing the resource) that may be extended to sequences that are in the property (e.g., we release the resource on the ninth step).

2.3 Security Automata

One way of enforcing security properties is with a *monitor* that runs in parallel with a *target* program [Sch00]. Whenever the target program wishes to execute a security-relevant operation, the monitor first checks its policy to determine whether or not that operation is allowed. If the target program's execution sequence σ is not in the property, the monitor transforms it into a sequence σ' that obeys the property.

A program monitor can be formally modeled by a *security automaton* A , which is a deterministic finite or infinite state machine (Q, q_0, \mathcal{T}) that is defined with respect to some system (\mathcal{A}, Σ) . Q specifies the possible automaton states and q_0 is the initial state. Each variety of automata will have a slightly different sort of complete, deterministic transition function \mathcal{T} , and these differences account for the variations in their expressive power. The exact specification of \mathcal{T} is part of the definition of each kind of automaton.

We will specify the execution of each different kind of security automaton A using a labeled operational semantics. The basic single-step judgment will have

the form $(\sigma, q) \xrightarrow{\tau}_A (\sigma', q')$ where σ denotes the sequence of actions that the target program wants to execute; q denotes the current state of the automaton; σ' and q' denote the action sequence and state after the automaton takes a single step; and τ denotes the sequence of actions produced by the automaton. The input sequence σ is not observable to the outside world whereas the outputs, τ , are observable. When we consider nonuniform systems, we will constrain σ rather than τ (i.e., τ may be any element of \mathcal{A}^*).

We generalize the single-step judgment to a multi-step judgment as follows. The basic multi-step judgment has the form $(\sigma, q) \xRightarrow{\tau}_A (\sigma', q')$.

$$\boxed{(\sigma, q) \xRightarrow{\tau}_A (\sigma', q')}$$

$$\frac{}{(\sigma, q) \xRightarrow{\tau}_A (\sigma, q)} \quad (\text{A-REFLEX})$$

$$\frac{(\sigma, q) \xrightarrow{\tau_1}_A (\sigma'', q'') \quad (\sigma'', q'') \xrightarrow{\tau_2}_A (\sigma', q')}{(\sigma, q) \xRightarrow{\tau_1; \tau_2}_A (\sigma', q')} \quad (\text{A-TRANS})$$

In this paper we consider security automata whose transition functions take the current state and input action (the first of the actions that the target program wants to execute) and returns a new automaton state and the sequence of actions the automaton outputs. Our automata may or may not consume the current input action while making a transition.

2.4 Enforceable Properties

To be able to compare and contrast the power of various security mechanisms, it is crucial to have a rigorous definition of what it means for a mechanism to enforce a security property. We believe that enforcement mechanisms can only accomplish their task effectively when they obey the following two abstract principles.

(Soundness) An enforcement mechanism must ensure that all observable outputs obey the property in question.

(Transparency) An enforcement mechanism must preserve the semantics of executions that already obey the property in question.

The first criterion requires that our automata transform bad program executions in such a way that their observable outputs obey the property. A security mechanism that allows outsiders to observe executions that do not satisfy the given policy is an unsound and inadequate mechanism. The second criterion requires that enforcement mechanisms operate invisibly on executions that obey the property in question. In other words, any editing operations that they perform must not change the meaning of a valid execution.

Conservative Enforcement. When an enforcement mechanism is able to obey the first criterion but not necessarily the second, we say that the mechanism enforces the property *conservatively*.

Definition 1 (Conservative Enforcement)

An automaton A with starting state q_0 conservatively enforces a property \hat{P} on the system (\mathcal{A}, Σ) if and only if $\forall \sigma \in \Sigma \exists q' \exists \sigma' \in \mathcal{A}^*$.

1. $(\sigma, q_0) \xrightarrow{\sigma'}_A (\cdot, q')$, and
2. $\hat{P}(\sigma')$

Conservative enforcement gives the automaton great latitude. The relationship between the automaton's output and the output of the target is not constrained, so the automaton may choose to enforce policies simply by always outputting a particular stream of actions (or nothing at all), regardless of the behavior of the target. Such an automaton is able to enforce any property, but often not in a useful manner. The existential quantification of q' and σ' ensures that the automaton eventually halts on all inputs.

Precise Enforcement. In order to formalize the second enforcement criterion (TRANSPARENCY), we must be specific about what it means to preserve the semantics of a program execution. However, since this abstract notion may differ substantially from one computational context to the next, we offer a couple of different definitions. The simplest reasonable definition requires that the automaton in question outputs program actions in lockstep with the target program's action stream.

Definition 2 (Precise Enforcement)

An automaton A with starting state q_0 precisely enforces a property \hat{P} on the system (\mathcal{A}, Σ) if and only if $\forall \sigma \in \Sigma \exists q' \exists \sigma' \in \mathcal{A}^*$.

1. $(\sigma, q_0) \xrightarrow{\sigma'}_A (\cdot, q')$,
2. $\hat{P}(\sigma')$, and
3. $\hat{P}(\sigma) \Rightarrow \forall i \exists q''. (\sigma, q_0) \xrightarrow{\sigma[1..i]}_A (\sigma[i+1..], q'')$

An automaton precisely enforces a property if and only if it conservatively enforces the property and outputs actions in lockstep with the target application. The automaton must not interfere with program execution in any way when the input obeys the property; every action in a valid input sequence must be accepted before the next action is considered. Precise enforcement is often the right correctness criterion for security monitors when any delay in outputting program actions will change the semantics of the application. This is usually the case in interactive applications that require an action to return a value before continuing.

Effective Enforcement. Our definition of precise enforcement does not recognize the fact that it is often the case that two syntactically different action sequences may be semantically equivalent. This often happens in three sorts of ways:

1. A sequence containing unnecessary actions or a series of idempotent actions may be equivalent to a sequence in which one or more of these actions are removed. In some systems, for example, closing a file twice in succession is equivalent to closing it just once.
2. A sequence which replaces some action with a different action that has the same semantics may be considered equivalent to the original sequence. It may also be permissible to replace several actions with a single action that subsumes their behavior. For example, an action that opens a socket followed by an action that sends data on the socket is semantically equivalent to a macro instruction that does both.
3. A sequence in which two or more independent actions are permuted may be equivalent to the original sequence. For example, if we want to open two different files, the order in which we do this may not matter.

Our final definition of enforcement uses a system-specific equivalence relation (\cong) on executions to take these possibilities into account. We require that the relation be reflexive, symmetric, and transitive. Moreover, any property that we might consider should not distinguish equivalent sequences:

$$\sigma \cong \sigma' \Rightarrow \hat{P}(\sigma) \Leftrightarrow \hat{P}(\sigma') \quad (\text{EQUIVALENCE})$$

With our equivalence relation in hand, we can give a concrete definition of what it means for a security automaton to *effectively enforce* a security property.

Definition 3 (Effective Enforcement)

An automaton A with starting state q_0 effectively enforces a property \hat{P} on the system (\mathcal{A}, Σ) if and only if $\forall \sigma \in \Sigma \exists q' \exists \sigma' \in \mathcal{A}^*$.

1. $(\sigma, q_0) \xrightarrow{\sigma'}_A (\cdot, q')$,
2. $\hat{P}(\sigma')$, and
3. $\hat{P}(\sigma) \Rightarrow \sigma \cong \sigma'$

Informally, an automaton effectively enforces a property whenever it conservatively enforces the property and obeys the principle of transparency.

2.5 Limitations

As always when reasoning about security, one must be acutely aware of situations in which formal results do not apply and of attacks that can come from outside the model. In our case, there are several realistic situations in which security automata fail to enforce a desired policy.

- The policy is not a predicate on execution sequences. If the policy depends upon some external environmental factors, such as the value of some secret that has been hidden from the monitor, the monitor may not be able to enforce the policy properly. Moreover, because monitors only see individual sequences of actions and cannot in general base decisions on other executions of the target application, security automata are enforcers of *properties* rather than general *policies*.
- The monitor is unable to effectively manipulate (insert, suppress, etc.) certain security-relevant actions. In a real-time system, some properties might not be enforceable because the monitor simply cannot perform the expected computation in the necessary real-time window. Alternatively, there may be some data (secret keys, passwords, etc.) that the monitor cannot effectively synthesize. This might mean that the monitor is unable to insert certain program actions and enforce certain properties.
- The monitor is unable to observe certain security-relevant actions. If the application has direct access to some hardware device and the monitor is unable to interpose itself between the application and the device, it may be unable to enforce certain properties.
- The monitor is compromised by the untrusted program. If the application is able to corrupt the monitoring code or data then the monitor will not be able to enforce any meaningful properties. In practice, this means that any software monitor must be isolated from the application using safe language technology or operating system support.

3 Truncation Automata

The first and most limited sort of security automaton that we will investigate is the *truncation automaton*. A truncation automaton does not change the target's output in any way unless the target attempts to invoke a forbidden operation. If this occurs, the truncation automaton halts the target program. This sort of automaton is the primary focus of Schneider's work [Sch00]. However, he viewed truncation automata as sequence recognizers. To fit this idea into our framework, we recast these machines as sequence transformers.

3.1 Definition

A truncation automaton T is a finite or infinite state machine (Q, q_0, δ) that is defined with respect to some system (\mathcal{A}, Σ) . Q specifies the possible automaton states and q_0 is the initial state. The partial function $\delta : \mathcal{A} \times Q \rightarrow Q$ specifies the transition function for the automaton that indicates that the automaton

should accept the current input action and move to a new state. On inputs for which δ is not defined, the automaton halts the target program.

The operational semantics of truncation automata is specified below.

$$\boxed{(\sigma, q) \xrightarrow{\tau}_T (\sigma', q')}$$

$$\begin{array}{ll}
 (\sigma, q) \xrightarrow{a}_T (\sigma', q') & \text{(T-STEP)} \\
 \text{if } \sigma = a; \sigma' & \\
 \text{and } \delta(a, q) = q' & \\
 (\sigma, q) \xrightarrow{\cdot}_T (\cdot, q) & \text{(T-STOP)} \\
 \text{otherwise} &
 \end{array}$$

As described in Section 2.3, we extend the single-step relation to a multi-step relation using reflexivity and transitivity rules.

3.2 Precisely Enforceable Properties

Erlingsson and Schneider [UES99, UES00] demonstrate that mechanisms similar to truncation automata can precisely enforce important access-control properties including software fault isolation and Java stack inspection. Interestingly, although the definition of truncation automata is independent of the definition of safety, when we limit our attention to uniform systems we find that truncation automata can precisely enforce exactly the set of safety properties.

Theorem 1 (Uniform Precise T-Enforcement)

A property \hat{P} on the uniform system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be precisely enforced by a truncation automaton if and only if $\forall \sigma \in \text{pre}(\Sigma). \neg \hat{P}(\sigma) \Rightarrow \forall \sigma' \in \Sigma. (\sigma \preceq \sigma' \Rightarrow \neg \hat{P}(\sigma'))$.

Proof Please see Appendix A for the proof. We omit it here due to its length. ■

As will be the case for all security automata introduced in this paper, we can expand the set of properties that truncation automata can precisely enforce by considering how they operate in nonuniform systems. For example, suppose there is a policy requiring that if a program exits then all open resources must be closed, and suppose that we are working in a system in which all programs are instrumented by inserting resource-close actions at every exit point of the program. Although this is not a safety property because there is an illegal sequence (**open**) that may lead to a legal sequence (**open; close**), the instrumenting ensures that all illegal executions will eventually become legal. A truncation automaton can use this fact to accept potentially illegal sequences.

This example demonstrates that on nonuniform systems truncation automata can precisely enforce more than just safety properties. It also illustrates, however, that they can do this only when the nonuniform system rules out many

illegal input sequences on which it would be impossible for truncation automata to precisely enforce the desired property. The extent of a truncation automaton's power in nonuniform systems is captured by the following theorem.

Theorem 2 (Nonuniform Precise T-Enforcement)

A property \hat{P} on the system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be precisely enforced by some truncation automaton if and only if there exists a decidable predicate D such that for all executions $\sigma \in \text{pre}(\Sigma)$, if $\neg\hat{P}(\sigma)$ then

1. $D(\sigma) \Rightarrow (\forall \sigma' \in \Sigma. \sigma \preceq \sigma' \Rightarrow \neg\hat{P}(\sigma'))$
2. $\neg D(\sigma) \Rightarrow (\forall \sigma' \in \Sigma. \sigma \preceq \sigma' \Rightarrow \exists \tau \in \mathcal{A}^*. (\sigma; \tau \preceq \sigma' \wedge \hat{P}(\sigma; \tau)))$

Proof See Appendix A. ■

Intuitively, when encountering an illegal portion of an input sequence, a truncation automaton has two options. It may halt if there is no way to correct this sequence, or it may accept and continue if it can decide that this illegal input will *always* be corrected by future actions.

3.3 Effectively Enforceable Properties

In addition to harnessing the power of constraints placed on the set of executions (Σ) to precisely enforce non-safety properties, truncation automata can take advantage of a system's equivalence relation to effectively enforce non-safety properties, even on uniform systems.

Consider a uniform system whose only actions are to open and close a login window. We wish to enforce the policy that the first action in any sequence must close the login window. The policy also allows the login window to be closed and then opened and closed any number of times, as this is considered equivalent to closing the window once. All sequences (and only those sequences) of the form $\text{close}; (\text{open}; \text{close})^*$ satisfy the property, and $\text{close}; (\text{open}; \text{close})^* \cong \text{close}$. This is not a safety property because there is an illegal sequence $(\text{close}; \text{open})$ that can be extended into a legal sequence $(\text{close}; \text{open}; \text{close})$.

A truncation automaton can effectively enforce this property by checking that the first action closes the login window. If it does then the automaton accepts the action and halts; otherwise, it does not accept but still halts. This effectively enforces the property because any legal sequence will be equivalent to the accepted close action and all sequences emitted by the automaton (either \cdot or close) obey the property.

As with our example of precise enforcement on nonuniform systems, this illustrates both that truncation automata can effectively enforce more than just safety properties and that these non-safety properties are very limited. Before we can formally characterize the properties effectively enforceable by truncation automata, we must generalize the transformations truncation automata induce on instruction sequences into functions that act over sequences of actions. Given a set of actions \mathcal{A} , a function $\alpha^* : \mathcal{A}^* \rightarrow \mathcal{A}^*$ is a *truncation-rewrite function* if it satisfies the following conditions.

1. $\alpha^*(\cdot) = \cdot$
2. $\forall \sigma \in \mathcal{A}^* \forall a \in \mathcal{A}$.
 $\alpha^*(\sigma; a) = \alpha^*(\sigma); a$, or
 $\alpha^*(\sigma; a) = \alpha^*(\sigma) \wedge H(\alpha^*, \sigma; a)$

where $H(\rho^*, \sigma)$ is a predicate defined on a general rewrite function ρ^* and a sequence σ as follows.

$$H(\rho^*, \sigma) \Leftrightarrow (\forall \sigma' \in \mathcal{A}^*. \sigma \preceq \sigma' \Rightarrow \rho^*(\sigma') = \rho^*(\sigma))$$

Informally, $H(\rho^*, \sigma)$ is true if and only if an automaton whose output matches that of the rewrite function ρ^* halts on input σ , ceasing to output any additional actions. That is, whenever an automaton stops processing its input σ (for example, by applying rule T-STOP), it cannot examine any extensions to σ , so its output on any such extension must match its output on σ alone.

The restrictions placed on truncation-rewrite functions capture the operational restrictions of truncation automata. Condition (1) above ensures that truncation-rewrite functions do not output actions without receiving any input, and condition (2) stipulates that when examining the current action a , the rewrite function must either accept and output a (after it has finished outputting actions corresponding to earlier input) or halt and output nothing more. Note that when deciding how to transform an action a , the rewrite function can base its decision on the history of the entire execution up to a . Hence the input to the function in condition (2) is $\sigma; a$ and not just a .

The following theorem specifies the properties effectively enforceable by truncation automata on both uniform (where $\Sigma = \mathcal{A}^*$) and nonuniform systems.

Theorem 3 (Effective T-Enforcement)

A property \hat{P} on the system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be effectively enforced by some truncation automaton if and only if there exists a computable truncation-rewrite function α^ such that for all executions $\sigma \in \Sigma$,*

1. *If $\neg \hat{P}(\sigma)$ then $\hat{P}(\alpha^*(\sigma))$*
2. *If $\hat{P}(\sigma)$ then $\sigma \cong \alpha^*(\sigma)$*

Proof See Appendix A. ■

4 Suppression Automata

Given our novel view of security automata as sequence transformers, it is a short step to define new sorts of automata that have greater transformational capabilities than truncation automata. In this section, we describe *suppression automata* and characterize the properties they enforce precisely and effectively.

4.1 Definition

A *suppression automaton* S is a finite or infinite state machine (Q, q_0, δ, ω) that is defined with respect to some system (\mathcal{A}, Σ) . As before, Q is the set of all possible machine states, q_0 is a distinguished starting state for the machine, and the partial function δ specifies the transition function. The partial function $\omega : \mathcal{A} \times Q \rightarrow \{-, +\}$ has the same domain as δ and indicates whether or not the action in question is to be suppressed ($-$) or emitted ($+$).

$$\boxed{(\sigma, q) \xrightarrow{\tau}_S (\sigma', q')}$$

$$(\sigma, q) \xrightarrow{a}_S (\sigma', q') \quad (\text{S-STEP A})$$

if $\sigma = a; \sigma'$
and $\delta(a, q) = q'$
and $\omega(a, q) = +$

$$(\sigma, q) \xrightarrow{\cdot}_S (\sigma', q') \quad (\text{S-STEP S})$$

if $\sigma = a; \sigma'$
and $\delta(a, q) = q'$
and $\omega(a, q) = -$

$$(\sigma, q) \xrightarrow{\cdot}_S (\cdot, q) \quad (\text{S-STOP})$$

otherwise

Note that although rule S-STOP is not strictly necessary (rather than halting, the automaton could suppress all further actions), using S-STOP simplifies the automaton's specification and decreases its running time.

As before, we extend the single-step semantics to a multi-step semantics using the reflexivity and transitivity rules from Section 2.3.

4.2 Precisely Enforceable Properties

Similarly to truncation automata, in uniform systems suppression automata can precisely enforce any safety property and no other properties.

Theorem 4 (Uniform Precise S-Enforcement)

A property \hat{P} on the uniform system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be precisely enforced by a suppression automaton if and only if $\forall \sigma \in \text{pre}(\Sigma). \neg \hat{P}(\sigma) \Rightarrow \forall \sigma' \in \Sigma. (\sigma \leq \sigma' \Rightarrow \neg \hat{P}(\sigma'))$.

Proof See Appendix A. ■

In a nonuniform system, suppression automata can precisely enforce non-safety properties not enforceable by truncation automata. For example, consider the following system \mathcal{S} .

$$\begin{aligned}
\mathcal{A} &= \{\text{aq}, \text{use}, \text{rel}\} \\
\Sigma &= \{\text{aq}; \text{rel}, \\
&\quad \text{aq}; \text{use}; \text{rel}, \\
&\quad \text{aq}; \text{use}; \text{use}; \text{rel}\}
\end{aligned}$$

The symbols `aq`, `use`, and `rel` denote the acquisition, use, and release of a resource. The set of executions includes zero, one, or two uses of the resource. Such a scenario might arise were we to publish a policy that programs must release the resource after using it at most two times. After publishing such a policy, we might find a bug in our implementation that makes it impossible for us to handle the load we were predicting. Naturally we would want to tighten the security policy as soon as possible, but we might not be able to change the policy we have published. Fortunately, we can use a suppression automaton to suppress extra uses and dynamically change the policy from a two-use policy to a one-use policy. Notice that an ordinary truncation automaton is not sufficient to make this change because it can only terminate execution, in which case the `rel` symbol would be missing from the end of the sequence. After terminating a two-use application, a truncation automation would be unable to insert the `rel` necessary to satisfy the policy.

To formally characterize the properties that can be precisely enforced by suppression automata, we again generalize our automata to functions that act over sequences of symbols. Given a set of actions \mathcal{A} , a function $\omega^* : \mathcal{A}^* \rightarrow \mathcal{A}^*$ is a *suppression-rewrite function* if it satisfies the following conditions.

1. $\omega^*(\cdot) = \cdot$
2. $\forall \sigma \in \mathcal{A}^* \forall a \in \mathcal{A}$.
 $\omega^*(\sigma; a) = \omega^*(\sigma); a$, or
 $\omega^*(\sigma; a) = \omega^*(\sigma)$

These conditions encode the restrictions placed on suppression automata that (1) nothing can be output when no actions are input and (2) an input action can only be accepted or suppressed. Halting is equivalent to suppressing all additional input actions.

A suppression automaton can precisely enforce the following properties.

Theorem 5 (Nonuniform Precise S-Enforcement)

A property \hat{P} on the system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be precisely enforced by some suppression automaton if and only if there exists a computable suppression-rewrite function ω^* such that for all sequences $\sigma \in \Sigma$,

1. If $\neg \hat{P}(\sigma)$ then $\hat{P}(\omega^*(\sigma))$.
2. If $\hat{P}(\sigma)$ then $\omega^*(\sigma) = \sigma$.

Proof See Appendix A. ■

4.3 Effectively Enforceable Properties

Even on uniform systems, suppression automata can effectively enforce a proper superset of the properties enforceable by truncation automata. As in the preceding discussion of nonuniform precise enforcement, we illustrate this with an example.

Consider an authenticated-login policy that requires users who wish to login to use an authenticated login, on a uniform system with actions for both an unauthenticated login (`ulogin`) and an authenticated login (`alogin`). On this system, an unauthenticated login followed by an authenticated login is semantically equivalent to a single authenticated login (i.e., $\text{ulogin}; \text{alogin} \cong \text{alogin}$); however, logging in more than once with an authenticated login is not equivalent to a single authenticated login (i.e., $\text{alogin}; \text{alogin} \not\cong \text{alogin}$).

A truncation automaton cannot effectively enforce this property because upon seeing a `ulogin` it can take no valid action. It cannot accept because an authenticated login may never occur, which violates the property. It also cannot halt because the next action may be `alogin`, making the automaton's input (`ulogin; alogin`) obey the property but not be equivalent to its output (\cdot). In contrast, a suppression automaton can effectively enforce this property simply by suppressing the `ulogin`. If the next action is `alogin`, then the suppression automaton accepts this and will have correctly effectively enforced the property since its output (`alogin`) is equivalent to its input (`ulogin; alogin`). If the `alogin` does not appear (and hence the input sequence violates the property), then the suppression automaton will not have emitted anything, thereby satisfying the property.

The following theorem formally specifies the properties effectively enforceable by suppression automata.

Theorem 6 (Effective S-Enforcement)

A property \hat{P} on the system $S = (\mathcal{A}, \Sigma)$ can be effectively enforced by some suppression automaton if and only if there exists a computable suppression-rewrite function ω^* such that for all executions $\sigma \in \Sigma$,

1. If $\neg \hat{P}(\sigma)$ then $\hat{P}(\omega^*(\sigma))$
2. If $\hat{P}(\sigma)$ then $\sigma \cong \omega^*(\sigma)$

Proof See Appendix A. ■

The essence of suppression automata as effective enforcers is that they can suppress an action if it is potentially bad and at the same time unnecessary. If the input sequence turns out to be in the property, then because the suppressed action was unnecessary, no harm was done.

5 Insertion Automata

This section introduces *insertion automata*, which have the power to insert actions into the action stream or halt the target program but cannot suppress

actions. As with truncation and suppression automata, we define them formally and then consider what properties they can enforce precisely and effectively.

5.1 Definition

An *insertion automaton* I is a finite or infinite state machine (Q, q_0, δ, γ) that is defined with respect to some system (\mathcal{A}, Σ) . The partial function $\delta : \mathcal{A} \times Q \rightarrow Q$ specifies the transition function as before. The new element is a partial function γ that specifies the insertion of a finite sequence of actions into the program's action sequence. We call this the *insertion function* and it has type $\mathcal{A} \times Q \rightarrow \vec{\mathcal{A}} \times Q$, where the first component in the returned pair indicates the finite and non-empty sequence of actions to be inserted. In order to maintain the determinacy of the automaton, we require the domain of the insertion function to be disjoint from the domain of the transition function.

We specify the execution of an insertion automaton as before. The single-step relation is defined below.

$$\boxed{(\sigma, q) \xrightarrow{\tau}_I (\sigma', q')}$$

$$\begin{array}{ll}
 (\sigma, q) \xrightarrow{a}_I (\sigma', q') & \text{(I-STEP)} \\
 \text{if } \sigma = a; \sigma' & \\
 \text{and } \delta(a, q) = q' & \\
 \\
 (\sigma, q) \xrightarrow{\tau}_I (\sigma, q') & \text{(I-INS)} \\
 \text{if } \sigma = a; \sigma' & \\
 \text{and } \gamma(a, q) = \tau, q' & \\
 \\
 (\sigma, q) \xrightarrow{\cdot}_I (\cdot, q) & \text{(I-STOP)} \\
 \text{otherwise} &
 \end{array}$$

We also extend this single-step semantics to a multi-step semantics as before.

5.2 Precisely Enforceable Properties

Identically to truncation and suppression automata, on uniform systems insertion automata precisely enforce exactly the set of safety properties. Perhaps surprisingly, the power to suppress or insert actions does not enable security automata to precisely enforce more properties on uniform systems. This is because every illegal sequence must be transformed into a legal sequence, but doing so prevents any legal extension from being precisely enforced (because the sequence has already been modified and cannot be accepted in an entirely lockstep manner). Therefore, regardless of the transformational capabilities a security automaton has on execution sequences, on uniform systems it can only precisely enforce safety properties.

Theorem 7 (Uniform Precise I-Enforcement)

A property \hat{P} on the uniform system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be precisely enforced by an insertion automaton if and only if $\forall \sigma \in \text{pre}(\Sigma). \neg \hat{P}(\sigma) \Rightarrow \forall \sigma' \in \Sigma. (\sigma \preceq \sigma' \Rightarrow \neg \hat{P}(\sigma'))$.

Proof See Appendix A. ■

Once again, considering nonuniform systems expands the set of precisely enforceable properties. Let us consider the nonuniform **aq; use; rel** system of Section 4.2. This time, we want to enforce the property that the resource will be used exactly once or twice, rather than zero or one times. The only sequences which satisfy our property are **aq; use; rel** and **aq; use; use; rel**. A truncation automaton cannot precisely enforce this property because upon encountering an **aq** it can take no action. It cannot accept because the full input sequence may be **aq; rel** and the automaton has no way of inserting the required **use**; and it cannot truncate because then it would give up on precisely enforcing any input sequences that start with **aq** and actually do obey the property, such as **aq; use; rel**.

An insertion automaton precisely enforces this property by accepting **aq** and **use** actions. When it encounters a **rel**, it will insert a **use** if one was not already present in the input.

We formally characterize the properties that can be precisely enforced by an insertion automaton as follows.

Theorem 8 (Nonuniform Precise I-Enforcement)

A property \hat{P} on the system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be precisely enforced by some insertion automaton if and only if there exists a computable function γ_p such that for all executions $\sigma \in \text{pre}(\Sigma)$, if $\neg \hat{P}(\sigma)$ then

1. $\forall \sigma' \in \Sigma. \sigma \preceq \sigma' \Rightarrow \neg \hat{P}(\sigma')$, or
2. $\sigma \notin \Sigma$ and $\hat{P}(\sigma; \gamma_p(\sigma))$

Proof See Appendix A. ■

We can also compare the power of insertion automata with suppression automata. A suppression automaton cannot precisely enforce the one-or-two-use policy described in this section because it cannot insert **use** events that are necessary if the program releases the resource without using it, and it cannot simply suppress resource acquisitions because this would modify sequences that do satisfy the policy. Hence, insertion automata can precisely enforce some properties that suppression automata cannot. On the other hand, an insertion automaton can precisely enforce the zero-or-one-use **aq; use; rel** policy of Section 4.2 simply by inserting a **rel** event and then halting whenever a bad sequence is found (i.e., **aq; use; use** would be transformed to **aq; use; rel**). Generalizing from this, we see that insertion automata are strictly more powerful precise enforcers than suppression automata.

Theorem 9 (Precise S-Enforceable Subset of Precise I-Enforceable)

If there exists a suppression automaton $S = (Q_s, q_{0_s}, \delta_s, \omega)$ that precisely enforces \hat{P} on some system (\mathcal{A}, Σ) then there exists an insertion automaton $I = (Q_i, q_{0_i}, \delta_i, \gamma)$ that precisely enforces \hat{P} on the same system.

Proof See Appendix A. ■

5.3 Effectively Enforceable Properties

Let us now consider insertion automata operating as effective, rather than precise, enforcers. As with the other automata seen so far, insertion automata viewed in this light can enforce a greater range of properties, even on uniform systems. We illustrate this with a very application-specific example: the San Francisco cable car.

The property we want to enforce is that if a person boards the car (`board`), she must show her ticket either to the conductor (`show_conductor`) or to the driver (`show_driver`). She may show them her ticket either before or after getting onto the car, and it is OK, though redundant, to show the ticket twice. A person may board the car at most once. The only sequences not in the property are getting onto the car without ever showing a ticket and boarding more than once.

An insertion automaton can effectively enforce this property because upon seeing a bad sequence where no ticket has been shown but a passenger boards (`board`), it can insert the action `show_driver`. If the stream turns out to already contain either `show_driver` or `show_conductor`, no harm has been done (because `show_driver; board; show_driver` \cong `board; show_driver` and `show_driver; board; show_conductor` \cong `board; show_conductor`).

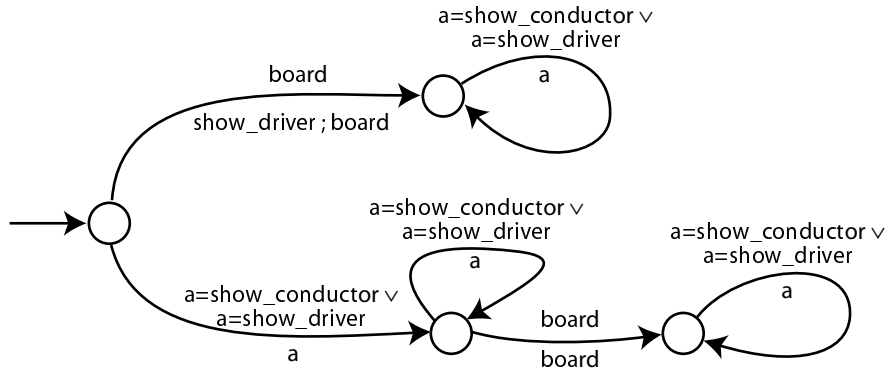


Figure 1: An insertion automaton that effectively enforces the cable-car policy on a uniform system.

Figure 1 shows an insertion automaton that effectively enforces the cable-

car policy. The nodes in the picture represent the automaton states and the arcs represent the transitions. The action above an arc triggers the transition, and the sequence below an arc represents the actions that are emitted. An arc with multiple symbols below it is an insertion transition. All other transitions are accepting transitions, and if there is no arc for the current action then the automaton halts.

A truncation automaton cannot effectively enforce this property because it has no way to handle an initial `board` action. If it accepts and there are no further actions, then it has allowed a passenger to board without showing a ticket; if it halts and the sequence does obey the property, then it has not preserved the semantics of the input sequence. Either way, it cannot effectively enforce the property.

As with precise enforcement, we can compare the powers of insertion and suppression automata as effective enforcers. The authenticated-login policy of Section 4.3 cannot be effectively enforced by an insertion automaton. When confronted with the sequence `ulogin`, an insertion automaton cannot accept because if there is no further input, the output would not obey the property; it cannot halt because the sequence could potentially be in the property (and `ulogin; alogin` $\not\equiv \cdot$); and it cannot insert an `alogin` because possible future `alogin` actions would make the stream obey the property yet not be equivalent to the stream with the inserted `alogin`.

This shows that there are properties that are effectively enforceable by suppression automata that are not effectively enforceable by insertion automata. Conversely, not all properties effectively enforceable by insertion automata can be effectively enforced by suppression automata. The cable-car example, for instance, cannot be effectively enforced by a suppression automaton. When confronted with a `board` action, a suppression automaton cannot accept because $\neg \hat{P}(\text{board})$, and cannot halt or suppress because that could change a legal sequence into one that is not equivalent to it (e.g., `board; show_driver` $\not\equiv \text{show_driver}$ and `board; show_driver` $\not\equiv \cdot$).

Before we characterize the properties that can be effectively enforced by insertion automata, we again generalize our automata into functions that act over sequences of symbols. Given a set of actions \mathcal{A} , a function $\gamma^* : \mathcal{A}^* \rightarrow \mathcal{A}^*$ is an *insertion-rewrite function* if it satisfies the following conditions.

1. $\gamma^*(\cdot) = \cdot$
2. $\forall \sigma \in \mathcal{A}^* \forall a \in \mathcal{A} \exists \tau \in \mathcal{A}^* .$
 $\gamma^*(\sigma; a) = \gamma^*(\sigma); \tau; a$, or
 $\gamma^*(\sigma; a) = \gamma^*(\sigma); \tau \wedge (\forall \tau' \in \mathcal{A}^* . \tau \neq \tau'; a) \wedge H(\gamma^*, \sigma; a)$

where $H(\gamma^*, \sigma)$ is the computable predicate defined in Section 3.3.

The intuition for insertion-rewrite functions is similar to that for the other rewrite functions, except that there are more possible behaviors of the corresponding automata to consider. When processing an action a , an insertion automaton can react in any of the following ways: it may accept a , insert some

number of actions (the number of inserted actions is finite in this article because we limit our analysis to finite-length sequences) before accepting a , halt the target, or insert some number of actions before halting the target. The first clause of condition (2) above captures the first two of these possibilities by allowing the rewrite function to insert some number (possibly zero) of actions before accepting. The second part of condition (2) captures the other possibilities where some number (possibly zero) of actions are inserted before halting. For the sake of determinacy—that is, to make sure that only one of the two clauses in condition (2) applies—the case in which the automaton inserts some sequence that ends with a and then halts is handled as if a were accepted before halting.

The following theorem formally specifies the properties effectively enforceable by insertion automata.

Theorem 10 (Effective I-Enforcement)

A property \hat{P} on the system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be effectively enforced by some insertion automaton if and only if there exists a computable insertion-rewrite function γ^* such that for all executions $\sigma \in \Sigma$,

1. If $\neg\hat{P}(\sigma)$ then $\hat{P}(\gamma^*(\sigma))$
2. If $\hat{P}(\sigma)$ then $\sigma \cong \gamma^*(\sigma)$

Proof See Appendix A. ■

The defining feature of insertion automata as effective enforcers is that they allow required actions to be inserted into a sequence whenever it is OK for the inserted actions to be repeated. If the automaton’s input actually does obey the property, the automaton’s output remains semantically equivalent to the input.

6 Edit Automata

We now turn to a far more powerful security automaton, the *edit automaton*. After defining edit automata formally and considering the properties they can enforce precisely and effectively, we present an extended example of an edit automaton that effectively enforces a market transaction policy.

6.1 Definition

We form an *edit automaton* E by combining the insertion automaton with the suppression automaton. Our machine is now described by a 5-tuple of the form $(Q, q_0, \delta, \gamma, \omega)$. The operational semantics is derived from the composition of the operational rules of the two previous automata. Again, the partial functions δ and ω have the same domain while δ and γ have disjoint domains.

$$\boxed{(\sigma, q) \xrightarrow{\tau}_E (\sigma', q')}$$

	$(\sigma, q) \xrightarrow{a}_E (\sigma', q')$	(E-STEPA)
if $\sigma = a; \sigma'$ and $\delta(a, q) = q'$ and $\omega(a, q) = +$		
	$(\sigma, q) \xrightarrow{\cdot}_E (\sigma', q')$	(E-STEPS)
if $\sigma = a; \sigma'$ and $\delta(a, q) = q'$ and $\omega(a, q) = -$		
	$(\sigma, q) \xrightarrow{\tau}_E (\sigma, q')$	(E-INS)
if $\sigma = a; \sigma'$ and $\gamma(a, q) = \tau, q'$		
	$(\sigma, q) \xrightarrow{\cdot}_E (\cdot, q)$	(E-STOP)
otherwise		

Although rule E-STEPA is not strictly necessary (rather than accepting an action a , the automaton could insert a and then suppress the original a), its presence allows acceptance of an action in only one step. This simplifies the specification of the automaton and in general decreases its running time. Similarly, and as with suppression automata, the effect of rule E-STOP can be accomplished by suppressing any further input, again at the cost of increased running time and specification complexity.

As with the other security automata, we extend the single-step semantics of edit automata to a multi-step semantics with the rules for reflexivity and transitivity.

6.2 Precisely Enforceable Properties

On uniform systems, edit automata precisely enforce exactly the set of safety properties. This follows immediately from the discussion in Section 5.2—as precise enforcers on uniform systems, edit automata have the same power as truncation, suppression, and insertion automata.

Theorem 11 (Uniform Precise E-Enforcement)

A property \hat{P} on the uniform system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be precisely enforced by an edit automaton if and only if $\forall \sigma \in \text{pre}(\Sigma). \neg \hat{P}(\sigma) \Rightarrow \forall \sigma' \in \Sigma. (\sigma \preceq \sigma' \Rightarrow \neg \hat{P}(\sigma'))$.

Proof See Appendix A. ■

When we consider nonuniform systems, edit automata precisely enforce exactly those properties precisely enforceable by insertion automata. This result is not particularly surprising because edit automata are just a composition of insertion and suppression automata, and insertion automata precisely enforce a superset of the properties precisely enforceable by suppression automata.

Theorem 12 (Nonuniform Precise E-Enforcement)

A property \hat{P} on the system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be precisely enforced by some edit automaton if and only if there exists a computable function γ_p such that for all executions $\sigma \in pre(\Sigma)$, if $\neg\hat{P}(\sigma)$ then

1. $\forall\sigma' \in \Sigma. \sigma \preceq \sigma' \Rightarrow \neg\hat{P}(\sigma')$, or
2. $\sigma \notin \Sigma$ and $\hat{P}(\sigma; \gamma_p(\sigma))$

Proof See Appendix A. ■

6.3 Effectively Enforceable Properties

Similarly to truncation, suppression, and insertion automata, considering edit automata as effective enforcers enables them to enforce a wider range of properties. In addition, edit automata can effectively enforce properties effectively enforceable by neither suppression nor insertion automata. Section 6.4 illustrates this with a detailed example.

Edit automata are very powerful effective enforcers because they can insert actions that are required and later suppress those same actions if they appear in the original input sequence. Conversely, an edit automaton can suppress a sequence of potentially illegal actions, and if the sequence is later determined to be legal, just re-insert it. Interestingly, with this technique edit automata can effectively enforce any property—the automaton simply suppresses all actions until it can confirm that the current prefix obeys the property, at which point it inserts all the suppressed actions. Any legal input will thus be output without modification. If an input is illegal, an edit automaton will output its longest valid prefix.

The following theorem formalizes this idea, and its proof shows how to construct an edit automaton to effectively enforce any property. This theorem holds on all systems \mathcal{S} because it only uses strict equality as the equivalence relation and places no constraints on the set of possible executions (so \mathcal{S} may be uniform).

Theorem 13 (Effective E-Enforcement)

Any property \hat{P} on the system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be effectively enforced by some edit automaton.

Proof See Appendix A. ■

6.4 An Example: Transactions

To demonstrate the power of our edit automata, we show how to implement the monitoring of transactions. The desired properties of atomic transactions [EN94], commonly referred to as the ACID properties, are atomicity (either the entire transaction is executed or no part of it is executed), consistency preservation (upon completion of the transaction the system must be in a consistent state), isolation (the effects of a transaction should not be visible to other concurrently executing transactions until the first transaction is committed), and durability or permanence (the effects of a committed transaction cannot be undone by a future failed transaction).

The first property, atomicity, can be modeled using an edit automaton by creating an intentions log. That is, the automaton suppresses input actions from the start of the transaction and, if the transaction completes successfully, the entire sequence of actions is emitted atomically to the output stream; otherwise it is discarded. Consistency preservation can be enforced by simply verifying that the sequence to be emitted leaves the system in a consistent state. The durability or permanence of a committed transaction is ensured by the fact that committing a transaction is modeled by emitting the corresponding sequence of actions to the output stream. Once an action has been written to the output stream it can no longer be touched by the automaton; furthermore, failed transactions output nothing. We only model the actions of a single agent in this example and therefore ignore issues of isolation.




Figure 2: An edit automaton that effectively enforces the market policy.

To make our example more concrete, we will model a simple market system with two main actions, $\text{take}(n)$ and $\text{pay}(n)$, which represent acquisition of n apples and the corresponding payment. We let \mathbf{a} range over all the actions that might occur in the system (such as take , pay , window-shop , browse , etc.). Our policy is that every time an agent takes n apples it must pay for those apples. Payments may come before acquisition or vice versa, and $\text{take}(n); \text{pay}(n)$

is semantically equivalent to $\text{pay}(n); \text{take}(n)$. The automaton enforces the atomicity of this transaction by emitting $\text{take}(n); \text{pay}(n)$ only when the transaction completes. If payment is made first, the automaton allows clients to perform other actions such as `browse` before paying (the take-pay transaction appears atomically after all such intermediary actions). On the other hand, if apples are taken and not paid for immediately, we issue a warning and abort the transaction. Consistency is ensured by remembering the number of apples taken or the size of the prepayment in the state of the machine. Once acquisition and payment occur, the sale is final and there are no refunds (durability). Figure 2 displays the infinite-state edit automaton that effectively enforces our market policy. Arcs with no symbols beneath them represent suppression transitions.

Neither an insertion automaton nor a suppression automaton can effectively enforce the market policy. An insertion automaton, if it encounters $\text{pay}(n)$, can take no action: it cannot halt because the next action may be $\text{take}(n)$, it cannot accept because there may be no more actions, and it cannot insert $\text{take}(n)$ because that may disrupt a valid sequence of $\text{take}(n)$ and $\text{pay}(n)$ in an input that obeys the property. A suppression automaton, upon encountering $\text{take}(n)$, can take no action: it cannot suppress or halt because the next action may be $\text{pay}(n)$, and it cannot accept because there may be no more actions.

7 Related Work

Schneider [Sch00] was the first to analyze the power of security mechanisms. He defined the set of security properties precisely enforceable by truncation automata and observed that it was a subset of the safety properties. Schneider also briefly mentions mechanisms more powerful than truncation automata, but he does not give a detailed analysis of their power. Schneider’s definition provides an upper bound on the set of properties that can be precisely enforced by security automata on a uniform input set. However, it is a loose upper bound. Viswanathan, Kim, and others [Vis00, KKL⁺02] demonstrate that to obtain a tighter bound on the power of run-time monitors, one must add computability constraints to the definition of truncation automata. Viswanathan [Vis00] has also demonstrated that the set of properties enforceable through run-time monitoring is equivalent to the CoRE properties.

Concurrently with our work, Hamlin, Morrisett, and Schneider [HMS02] have begun to investigate the power of a broader set of enforcement mechanisms. They consider security mechanisms based upon static analysis, run-time monitoring, and program rewriting and compare and contrast the power of these mechanisms. They observe that the statically enforceable properties correspond to the recursively decidable properties of programs. Taken with Viswanathan’s result, this fact implies that run-time monitors (when given access to the source program text) can enforce strictly more properties than can be enforced through static analysis. Hamlin et al. also prove that program rewriters do not correspond to any complexity class in the arithmetic hierarchy.

In contrast with these other theoretical research efforts, our work provides

a detailed analysis of the power of run-time monitors. We introduce several kinds of monitors, each with different run-time capabilities (i.e., truncation, suppression, and insertion). We clearly specify what it means to enforce a property and we demonstrate that the power of these monitors varies depending upon the context in which they are used.

Implementation efforts [ET99, UES99, UES00, KVBA⁺99] are considerably more advanced than the corresponding theoretical investigations. In general, these systems allow arbitrary code to be executed in response to a potential security violation, so edit automata provide a reasonable model for attempting to understand their behavior. In most cases, these languages can be considered domain-specific aspect-oriented programming languages [KHH⁺01], which have been designed to facilitate the enforcement of security properties.

Other researchers have investigated optimization techniques for run-time monitors [CF00, Thi01], and certification of programs instrumented with security checks [Wal00]. Kim et al. [KKL⁺02] analyze the performance cost of run-time monitors and show that the language one uses to define monitors has a significant impact on performance. They also define optimizations to decrease the overhead of run-time monitoring.

Run-time monitoring and checking can also be used in settings where security is not necessarily the primary focus. Lee et al. [KVBA⁺99], for example, have developed a monitoring system specifically for improving the reliability of real-time systems. Sandholm and Schwartzbach have used run-time monitoring to ensure that concurrently executing Web scripts obey important safety properties [SS98].

8 Conclusions

In this article, we have introduced a detailed framework for reasoning about the power of security mechanisms that are able to intercept and modify untrusted program actions at run time.

Before we begin the analysis proper, we carefully define what it means to enforce a property at run time. Our definitions of policy enforcement is motivated by two main considerations:

(Soundness) The final output of a monitored system must obey the policy. Consequently, bad programs that would otherwise violate the policy must have their executions modified by the enforcement mechanism.

(Transparency) Whenever the untrusted program obeys the policy in question, a run-time enforcement mechanism should preserve the semantics of the untrusted program. In other words, the actions of the enforcement mechanism should not be observable when monitoring good programs that do not violate the policy.

One obtains slightly different definitions of enforcement depending upon the interpretation of what it means for a mechanism to be “semantics-preserving.”

Consequently, it is necessary to ensure that the definition of enforcement is appropriate for the application and context at hand.

Once the definition of enforcement is set, we investigate the power of a hierarchy of monitors, each with varying run-time capabilities, including truncation, suppression, and insertion of security-relevant program actions. We demonstrate that the power of these different kinds of monitors varies, and we characterize their enforceable properties precisely. We also observe that when the monitor’s input space is constrained (possibly by a prior program rewriting pass or other mechanisms), a monitor may be able to enforce a larger set of properties. A summary of our main findings is presented in Figures 3 and 4.

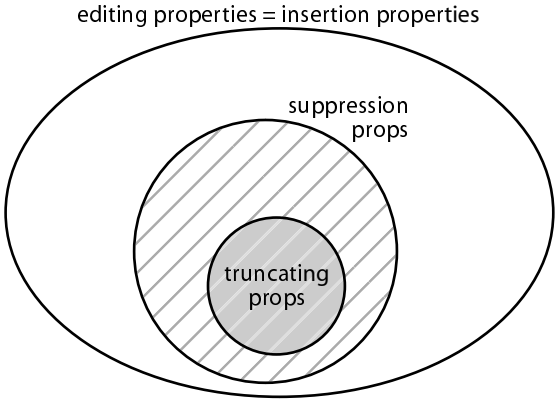


Figure 3: A taxonomy of *precisely* enforceable security properties on nonuniform systems.

This article sets some of the fundamental limits of security mechanisms that operate by monitoring and modifying program behavior at run time. However, there are many open questions that continue to spark our interest in this area. In particular, we desire to understand further the impact of constraining the resources available either to the running program or the run-time monitor. For example, are there practical properties that can be enforced by exponential-time but not polynomial-time monitors? What effect does bounding the space available to edit automata have on the set of enforceable properties? What if we limit the program’s access to random bits and therefore its ability to use strong cryptography? Can we generalize our analyses to infinite sequences? These unanswered questions and many similar ones suggest a variety of new research directions in this nascent research field.




Figure 4: A taxonomy of *effectively* enforceable security properties on both uniform and nonuniform systems.

Acknowledgments

Enlightening discussions with Kevin Hamlin, Greg Morrisett, and Fred Schneider helped to stimulate our research and improve this article. We would also like to thank the anonymous reviewers for their helpful and thorough comments.

References

- [AS85] Bowen Alpern and Fred B. Schneider. Defining liveness. *Information Processing Letters*, 21(4):181–185, October 1985. 6
- [AS87] Bowen Alpern and Fred Schneider. Recognizing safety and liveness. *Distributed Computing*, 2:117–126, 1987. 5, 6
- [BLW02] Lujo Bauer, Jarred Ligatti, and David Walker. More enforceable security policies. In *Foundations of Computer Security*, Copenhagen, Denmark, July 2002. 1
- [CF00] Thomas Colcombet and Pascal Fradet. Enforcing trace properties by program transformation. In *Twenty-Seventh ACM Symposium on Principles of Programming Languages*, pages 54–66, Boston, January 2000. ACM Press. 26
- [EN94] Ramez Elmasri and Shamkant B. Navathe. *Fundamentals of database systems*. The Benjamin/Cummings Publishing Company, Inc., 1994. 24

- [ET99] David Evans and Andrew Twyman. Flexible policy-directed code safety. In *IEEE Security and Privacy*, Oakland, CA, May 1999. 26
- [HMS02] Kevin Hamlin, Greg Morrisett, and Fred Schneider. In pursuit of punctuality: Computability classes for enforcement mechanisms. Unpublished Manuscript, November 2002. 25
- [KHH⁺01] Gregor Kiczales, Erik Hilsdale, Jim Hugunin, Mik Kersten, Jeffrey Palm, and William Griswold. An overview of AspectJ. In *European Conference on Object-oriented Programming*. Springer-Verlag, 2001. 26
- [KKL⁺02] Moonjoo Kim, Sampath Kannan, Insup Lee, Oleg Sokolsky, and Mahesh Viswanathan. Computational analysis of run-time monitoring—fundamentals of Java-MaC. In *Run-time Verification*, June 2002. 25, 26
- [KVBA⁺99] Moonjoo Kim, Mahesh Viswanathan, Hanene Ben-Abdallah, Sampath Kannan, Insup Lee, and Oleg Sokolsky. Formally specified monitoring of temporal properties. In *European Conference on Real-time Systems*, York, UK, June 1999. 26
- [Lam77] Leslie Lamport. Proving the correctness of multiprocess programs. *IEEE Transactions of Software Engineering*, 3(2):125–143, 1977. 5
- [Sch00] Fred B. Schneider. Enforceable security policies. *ACM Transactions on Information and Systems Security*, 3(1):30–50, February 2000. 2, 4, 6, 10, 25
- [SS98] Anders Sandholm and Michael Schwartzbach. Distributed safety controllers for web services. In *Fundamental Approaches to Software Engineering*, volume 1382 of *Lecture Notes in Computer Science*, pages 270–284. Springer-Verlag, 1998. 26
- [Thi01] Peter Thiemann. Enforcing security properties by type specialization. In *European Symposium on Programming*, Genova, Italy, April 2001. 26
- [UES99] Úlfar Erlingsson and Fred B. Schneider. SASI enforcement of security policies: A retrospective. In *Proceedings of the New Security Paradigms Workshop*, pages 87–95, Caledon Hills, Canada, September 1999. 11, 26
- [UES00] Úlfar Erlingsson and Fred B. Schneider. IRM enforcement of Java stack inspection. In *IEEE Symposium on Security and Privacy*, pages 246–255, Oakland, California, May 2000. 11, 26

- [Vis00] Mahesh Viswanathan. *Foundations for the Run-time Analysis of Software Systems*. PhD thesis, University of Pennsylvania, 2000. 25
- [Wal00] David Walker. A type system for expressive security policies. In *Twenty-Seventh ACM Symposium on Principles of Programming Languages*, pages 254–267, Boston, January 2000. 26

A Proofs of Theorems

Theorem 1 (Uniform Precise T-Enforcement)

A property \hat{P} on the uniform system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be precisely enforced by a truncation automaton if and only if $\forall \sigma \in \text{pre}(\Sigma). \neg \hat{P}(\sigma) \Rightarrow \forall \sigma' \in \Sigma. (\sigma \preceq \sigma' \Rightarrow \neg \hat{P}(\sigma'))$.

Proof (If Direction) We construct a truncation automaton that precisely enforces any such \hat{P} as follows.

- States: $q \in \text{pre}(\Sigma)$ (the sequence of actions seen so far)
- Start state: $q_0 = \cdot$ (the empty sequence)
- Transition function (δ):
Consider processing the action a in state σ .
 - (A) If $\hat{P}(\sigma; a)$ then emit the action a and continue in state $\sigma; a$.
 - (B) If $\neg \hat{P}(\sigma; a)$ then simply stop (i.e., we leave δ is undefined in this case).

The automaton maintains the invariant $\text{INV}_p(q)$ that $q = \sigma$ and σ has been output and $\hat{P}(\sigma)$. The automaton can initially establish $\text{INV}_p(q_0)$ since our definition of a property assumes $\hat{P}(\cdot)$ for all properties. A simple inductive argument on the length of the input σ suffices to show that the invariant is maintained for all inputs.

We now show that this automaton precisely enforces \hat{P} on any $\sigma \in \Sigma$ (where $\Sigma = \mathcal{A}^*$ because the system is uniform). There are two cases to consider.

- **Case $\hat{P}(\sigma)$:**
Consider any prefix τ of σ . If $\neg \hat{P}(\tau)$ then by the safety constraint stated in the theorem and the fact that $\tau \preceq \sigma$, we would have $\neg \hat{P}(\sigma)$, contrary to the assumption that $\hat{P}(\sigma)$. Therefore, $\hat{P}(\tau)$ must be true for all prefixes τ of σ , so by examination of the transition function given above, the automaton must accept every prefix of σ without ever halting.
- **Case $\neg \hat{P}(\sigma)$:**
 INV_p maintains that regardless of the state of the automaton, it has always emitted some σ such that $\hat{P}(\sigma)$.

The automaton correctly precisely enforces \hat{P} in both cases.

(Only-If Direction) Consider any $\sigma \in pre(\Sigma)$ such that $\neg\hat{P}(\sigma)$. Because $\Sigma = \mathcal{A}^*$, we have $\sigma \in \Sigma$, so the truncation automaton T must edit σ when σ is supplied as input. That is, since T precisely enforces \hat{P} and $\neg\hat{P}(\sigma)$ and $\sigma \in \Sigma$, it cannot be the case that $(\sigma, q_0) \xrightarrow{\sigma}_T (\cdot, q')$ for some q' . Suppose for the sake of obtaining a contradiction that $\exists\sigma' \in \Sigma.\sigma \preceq \sigma' \wedge \hat{P}(\sigma')$. Then by the definition of precise enforcement, all actions of σ' must be accepted without any editing, implying (since σ is later extended to σ') that $(\sigma, q_0) \xrightarrow{\sigma}_T (\cdot, q')$ for some q' , which we just showed cannot be true. Therefore, $\exists\sigma' \in \Sigma.\sigma \preceq \sigma' \wedge \hat{P}(\sigma')$ is not true, so $\forall\sigma' \in \Sigma.\sigma \preceq \sigma' \Rightarrow \neg\hat{P}(\sigma')$. ■

Theorem 2 (Nonuniform Precise T-Enforcement)

A property \hat{P} on the system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be precisely enforced by some truncation automaton if and only if there exists a decidable predicate D such that for all executions $\sigma \in pre(\Sigma)$, if $\neg\hat{P}(\sigma)$ then

1. $D(\sigma) \Rightarrow (\forall\sigma' \in \Sigma.\sigma \preceq \sigma' \Rightarrow \neg\hat{P}(\sigma'))$
2. $\neg D(\sigma) \Rightarrow (\forall\sigma' \in \Sigma.\sigma \preceq \sigma' \Rightarrow \exists\tau \in \mathcal{A}^*.\sigma \preceq \tau \wedge \hat{P}(\sigma; \tau))$

Proof (If Direction) We construct a truncation automaton that precisely enforces any such \hat{P} as follows.

- States: $q \in pre(\Sigma) \times \{+, -\}$ (the sequence of actions seen so far paired with + (-) to indicate that this sequence is (is not) in the property)
- Start state: $q_0 = \langle \cdot, + \rangle$ (the empty sequence)
- Transition function (δ):
Consider processing the action a in state q .
 - (A) If $q = \langle \sigma, + \rangle$ and $\hat{P}(\sigma; a)$ then emit the action a and continue in state $\langle \sigma; a, + \rangle$.
 - (B) If $q = \langle \sigma, + \rangle$ and $\neg\hat{P}(\sigma; a)$ and $D(\sigma; a)$ then simply stop (i.e., δ is undefined in this case).
 - (C) If $q = \langle \sigma, + \rangle$ and $\neg\hat{P}(\sigma; a)$ and $\neg D(\sigma; a)$ then emit a and continue in state $\langle \sigma; a, - \rangle$
 - (D) If $q = \langle \sigma, - \rangle$ and $\neg\hat{P}(\sigma; a)$ then emit the action a and continue in state $\langle \sigma; a, - \rangle$.
 - (E) If $q = \langle \sigma, - \rangle$ and $\hat{P}(\sigma; a)$ then emit the action a and continue in state $\langle \sigma; a, + \rangle$.

The automaton maintains the following invariant $INV_p(q)$ when processing some input $\tau \in \Sigma$.

- If $q = \langle \sigma, + \rangle$ then σ is the prefix of τ input so far and σ has been emitted and $\hat{P}(\sigma)$.

- If $q = \langle \sigma, - \rangle$ then σ is the prefix of τ input so far and σ has been emitted and $\neg \hat{P}(\sigma)$ and $\forall \sigma' \in \Sigma. \sigma \preceq \sigma' \Rightarrow \exists \tau \in \mathcal{A}^*. (\sigma; \tau \preceq \sigma' \wedge \hat{P}(\sigma; \tau))$.

The automaton can initially establish $\text{INV}_p(q_0)$ since our definition of a property assumes $\hat{P}(\cdot)$ for all properties. A simple inductive argument on the length of the input σ suffices to show that the invariant is maintained for all inputs.

The automaton never halts in some state $\langle \sigma, - \rangle$ because the precondition for entering this state given in automaton transition (C) ($\neg D(\sigma)$, implying that $\forall \sigma' \in \Sigma. \sigma \preceq \sigma' \Rightarrow \exists \tau \in \mathcal{A}^*. \sigma; \tau \preceq \sigma' \wedge \hat{P}(\sigma; \tau)$) guarantees that we will eventually reach some prefix of the input that is in the property. When this occurs after having input some $\sigma; \tau$, the automaton will follow transition (E) to return to state $\langle \sigma; \tau, + \rangle$.

In addition, we note that all transitions of the truncation automaton have decidable boolean guards, so the automaton must indeed halt on all inputs. Thus, when some $\sigma' \in \Sigma$ is input into this automaton, the machine either halts in state $\langle \sigma', + \rangle$ or halts by following transition (B). In the latter case, by examination of the guard for transition (B) and INV_p , the final state must have been $\langle \sigma, + \rangle$ for some $\sigma \preceq \sigma'$. In any case, INV_p ensures that whenever the automaton halts, it has emitted some σ such that $\hat{P}(\sigma)$.

Finally, we must show that this automaton accepts every action in any input sequence σ' such that $\hat{P}(\sigma')$. In such a case, the predicate $\forall \sigma' \in \Sigma. \sigma \preceq \sigma' \Rightarrow \neg \hat{P}(\sigma')$ can clearly never be true for any σ such that $\sigma \preceq \sigma'$, so by the contrapositive of clause (1) in the theorem, $\neg D(\sigma)$ for all σ such that $\sigma \preceq \sigma'$. The only transition that could prevent the automaton from accepting all actions in σ' is (B), but this transition cannot be followed on input σ' because it requires $D(\sigma)$ to be true for some $\sigma \preceq \sigma'$. The automaton can therefore never follow transition (B) on input σ' , so it must accept every action in σ' , as required.

(Only-If Direction) Given some truncation automaton T , we must construct a decision procedure D such that for all $\sigma \in \text{pre}(\Sigma)$ where $\neg \hat{P}(\sigma)$,

1. $D(\sigma) \Rightarrow (\forall \sigma' \in \Sigma. \sigma \preceq \sigma' \Rightarrow \neg \hat{P}(\sigma'))$
2. $\neg D(\sigma) \Rightarrow (\forall \sigma' \in \Sigma. \sigma \preceq \sigma' \Rightarrow \exists \tau \in \mathcal{A}^*. (\sigma; \tau \preceq \sigma' \wedge \hat{P}(\sigma; \tau)))$

The decision procedure simply runs T on input σ and evaluates to true if and only if T prematurely halts. That is, let $D(\sigma)$ be true if and only if T at some point steps via rule $T\text{-STOP}$ on input σ . T is a precise enforcer and as such must halt within finite time on all inputs (since $\forall \sigma \in \Sigma. \exists q'. \exists \sigma' \in \mathcal{A}^*. (\sigma, q_0) \xrightarrow{\sigma'}_T (\cdot, q')$), so D is a decidable predicate. Also, $D(\sigma)$ implies that $\forall \sigma' \in \Sigma. \sigma \preceq \sigma' \Rightarrow \neg \hat{P}(\sigma')$ because halting on σ prevents the automaton from precisely enforcing \hat{P} on any σ' such that $\sigma \preceq \sigma'$ and $\hat{P}(\sigma')$. Finally, $\neg D(\sigma)$ implies that $\forall \sigma' \in \Sigma. \sigma \preceq \sigma' \Rightarrow \exists \tau \in \mathcal{A}^*. (\sigma; \tau \preceq \sigma' \wedge \hat{P}(\sigma; \tau))$. In this case, we have $(\sigma, q_0) \xrightarrow{\sigma}_T (\cdot, q')$ for some q' , so because T is a truncation automaton precisely enforcing \hat{P} and must therefore emit some prefix $\sigma; \tau$ of its input σ' such that $\hat{P}(\sigma; \tau)$ (when T has already emitted σ on input σ and $\sigma \preceq \sigma'$), we

conclude that $\forall \sigma' \in \Sigma. \sigma \preceq \sigma' \Rightarrow \exists \tau \in \mathcal{A}^*. (\sigma; \tau \preceq \sigma' \wedge \hat{P}(\sigma; \tau))$. ■

Theorem 3 (Effective T-Enforcement)

A property \hat{P} on the system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be effectively enforced by some truncation automaton if and only if there exists a computable truncation-rewrite function α^* such that for all executions $\sigma \in \Sigma$,

1. If $\neg \hat{P}(\sigma)$ then $\hat{P}(\alpha^*(\sigma))$
2. If $\hat{P}(\sigma)$ then $\sigma \cong \alpha^*(\sigma)$

Proof (If Direction) We construct a truncation automaton that effectively enforces any such \hat{P} as follows.

- States: $q \in \text{pre}(\Sigma)$ (the sequence of actions seen so far)
- Start state: $q_0 = \cdot$ (for the empty sequence)
- Transition function (δ):

Consider processing the action a in state σ .

(A) If $\alpha^*(\sigma; a) = \alpha^*(\sigma); a$ then emit a and continue in state $\sigma; a$.

(B) If $\alpha^*(\sigma; a) = \alpha^*(\sigma)$ and $H(\alpha^*, \sigma; a)$ then halt (i.e., we leave δ undefined in this case).

The automaton maintains the invariant $\text{INV}_p(q)$ that $q = \sigma$ and σ is the input so far and the automaton has emitted $\alpha^*(\sigma)$. Initially, $\text{INV}_p(q_0)$ because $\alpha^*(\cdot) = \cdot$, and a simple inductive argument on the length of the input σ suffices to show that the invariant is maintained for all inputs.

We now show that the automaton emits $\alpha^*(\sigma)$ when the input is σ . There are two cases, derived by inspection of the truncation automaton. In the first case, the automaton halts in state σ , so by $\text{INV}_p(\sigma)$, $\alpha^*(\sigma)$ has been emitted. Otherwise, the automaton must halt in some state σ' such that $\sigma' \preceq \sigma$ and $H(\alpha^*, \sigma')$. $\text{INV}_p(\sigma')$ then implies that $\alpha^*(\sigma')$ has been emitted when the automaton halts. However, $H(\alpha^*, \sigma')$ and $\sigma' \preceq \sigma$ imply that $\alpha^*(\sigma') = \alpha^*(\sigma)$, so the automaton has in fact emitted $\alpha^*(\sigma)$ when it halts in this case as well.

Consider any execution $\sigma \in \Sigma$. By clause (1) in the theorem statement, if $\neg \hat{P}(\sigma)$ then $\hat{P}(\alpha^*(\sigma))$. Because $\alpha^*(\sigma)$ is actually what the automaton outputs on input σ , it correctly effectively enforces any $\sigma \in \Sigma$ such that $\neg \hat{P}(\sigma)$. Similarly, clause (2) of the theorem states that if $\hat{P}(\sigma)$ then $\sigma \cong \alpha^*(\sigma)$, so the automaton's output is equivalent to its input when $\hat{P}(\sigma)$. The automaton therefore effectively enforces \hat{P} .

(Only-If Direction) Define $\alpha^*(\sigma)$ to be whatever sequence is emitted by the truncation automaton on input σ . By this definition, $\alpha^*(\sigma)$ is a computable function because the automaton is an effective enforcer and as such acts deterministically and halts on all inputs in the set of executions (Σ). We first show that this is a truncation-rewrite function. Clearly, $\alpha^*(\cdot) = \cdot$. When processing some action a after having already processed any sequence σ , the automaton may step via T-STEP or T-STOP.

- **Case T-Step:** The automaton emits whatever has been emitted in processing σ (by definition, this is $\alpha^*(\sigma)$), followed by a .
- **Case T-Stop:** The automaton emits only what has already been emitted, so $\alpha^*(\sigma; a) = \alpha^*(\sigma)$. Because the automaton halts, it may not examine any remaining input, implying that $H(\alpha^*, \sigma; a)$ is true.

In any case, the truncation automaton's output adheres to the definition of a truncation-rewrite function.

Now consider any arbitrary $\sigma \in \Sigma$. By the definition of effective enforcement and that $\alpha^*(\sigma)$ is defined to be whatever sequence is emitted by the truncation automaton on input σ , $\hat{P}(\alpha^*(\sigma))$ and $\hat{P}(\sigma) \Rightarrow \sigma \cong \alpha^*(\sigma)$. ■

Theorem 4 (Uniform Precise S-Enforcement)

A property \hat{P} on the uniform system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be precisely enforced by a suppression automaton if and only if $\forall \sigma \in pre(\Sigma). \neg \hat{P}(\sigma) \Rightarrow \forall \sigma' \in \Sigma. (\sigma \preceq \sigma' \Rightarrow \neg \hat{P}(\sigma'))$.

Proof Both directions proceed completely analogously to the proof of Uniform Precise T-Enforcement. ■

Theorem 5 (Nonuniform Precise S-Enforcement)

A property \hat{P} on the system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be precisely enforced by some suppression automaton if and only if there exists a computable suppression-rewrite function ω^* such that for all sequences $\sigma \in \Sigma$,

1. If $\neg \hat{P}(\sigma)$ then $\hat{P}(\omega^*(\sigma))$.
2. If $\hat{P}(\sigma)$ then $\omega^*(\sigma) = \sigma$.

Proof (If Direction) We construct a suppression automaton that precisely enforces any such \hat{P} as follows.

- States: $q \in pre(\Sigma)$ (the sequence of actions seen so far)
- Start state: $q_0 = \cdot$ (the empty sequence)
- Transition and suppression functions (combined for simplicity):
Consider processing the action a in state $q = \sigma$.

- (A) If $\omega^*(\sigma; a) = \omega^*(\sigma); a$ then emit a and continue in state $\sigma; a$.
- (B) If $\omega^*(\sigma; a) = \omega^*(\sigma)$ then suppress a and continue in state $\sigma; a$.

This suppression automaton maintains the invariant $\text{INV}_p(q)$ that $q = \sigma$ and σ is the input seen so far and $\omega^*(\sigma)$ has been output. Initially, $\text{INV}_p(q_0)$ because $\omega^*(\cdot) = \cdot$, and a simple inductive argument on the length of the input σ suffices to show that the invariant is maintained for all inputs.

By inspection of the automaton, we note that when processing some input $\sigma \in \Sigma$, the automaton always halts in state σ because there is no more input to process. Thus, $\text{INV}_p(\sigma)$ ensures that the automaton always emits $\omega^*(\sigma)$ whenever the input to the automaton is σ .

Consider any execution $\sigma \in \Sigma$. By clause (1) in the theorem statement, if $\neg \hat{P}(\sigma)$ then $\hat{P}(\omega^*(\sigma))$. Because $\omega^*(\sigma)$ is actually what the automaton outputs on input σ , it correctly precisely enforces any $\sigma \in \Sigma$ such that $\neg \hat{P}(\sigma)$. Similarly, by clause (2), the automaton never suppresses any action (i.e., it accepts every action) in its input σ when $\hat{P}(\sigma)$. The automaton therefore precisely enforces \hat{P} .

(Only-If Direction) Define $\omega^*(\sigma)$ to be whatever sequence is emitted by the suppression automaton on input σ . By this definition, $\omega^*(\sigma)$ is a computable function because the automaton is a precise enforcer and as such acts deterministically and halts on all inputs in the set of executions (Σ). We first show that this is indeed a suppression-rewrite function. Clearly, $\omega^*(\cdot) = \cdot$. When processing some action a after having already processed any sequence σ , the automaton may step via S-STEP_A, S-STEP_S, or S-STOP. In the S-STEP_A case, the automaton emits whatever has been emitted in processing σ (by definition, this is $\omega^*(\sigma)$), followed by a . In the other cases, the automaton emits only $\omega^*(\sigma)$. Hence, ω^* is a valid suppression function.

Now consider any arbitrary $\sigma \in \Sigma$. By the definition of precise enforcement and that $\omega^*(\sigma)$ is defined to be whatever sequence is emitted by the suppression automaton on input σ , $\hat{P}(\omega^*(\sigma))$ (because all final outputs must be in the property) and $\hat{P}(\sigma) \Rightarrow \sigma = \omega^*(\sigma)$ (because no suppressions or halting may occur on inputs that are in the property). ■

Theorem 6 (Effective S-Enforcement)

A property \hat{P} on the system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be effectively enforced by some suppression automaton if and only if there exists a computable suppression-rewrite function ω^* such that for all executions $\sigma \in \Sigma$,

1. If $\neg \hat{P}(\sigma)$ then $\hat{P}(\omega^*(\sigma))$
2. If $\hat{P}(\sigma)$ then $\sigma \cong \omega^*(\sigma)$

Proof (If Direction) We construct a suppression automaton that effectively enforces any such \hat{P} as follows.

- States: $q \in pre(\Sigma)$ (the sequence of actions seen so far)
- Start state: $q_0 = \cdot$ (the empty sequence)
- Transition and suppression functions (combined for simplicity):
Consider processing the action a in state $q = \sigma$.
 - (A) If $\omega^*(\sigma; a) = \omega^*(\sigma); a$ then emit a and continue in state $\sigma; a$.
 - (B) If $\omega^*(\sigma; a) = \omega^*(\sigma)$ then suppress a and continue in state $\sigma; a$.

This is the same automaton constructed in the proof of Nonuniform Precise S-Enforcement, so by the same reasoning given there, this automaton always emits $\omega^*(\sigma)$ whenever the input to the automaton is σ .

Analogously to the argument given in the Effective T-Enforcement theorem, because the automaton always outputs $\omega^*(\sigma)$ on input σ , clauses (1) and (2) in the theorem statement ensure that the automaton effectively enforces \hat{P} .

(Only-If Direction) Define $\omega^*(\sigma)$ to be whatever sequence is emitted by the suppression automaton on input σ . This is a computable suppression-rewrite function by the same argument given in the Only-If Direction of the proof of Nonuniform Precise S-Enforcement. Moreover, given that the automaton is an effective enforcer of \hat{P} and $\omega^*(\sigma)$ always matches the output of the automaton on input σ , both requirements of \hat{P} stated in the theorem are clearly met. ■

Theorem 7 (Uniform Precise I-Enforcement)

A property \hat{P} on the uniform system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be precisely enforced by an insertion automaton if and only if $\forall \sigma \in pre(\Sigma). \neg \hat{P}(\sigma) \Rightarrow \forall \sigma' \in \Sigma. (\sigma \preceq \sigma' \Rightarrow \neg \hat{P}(\sigma'))$.

Proof Both directions proceed completely analogously to the proof of Uniform Precise T-Enforcement. ■

Theorem 8 (Nonuniform Precise I-Enforcement)

A property \hat{P} on the system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be precisely enforced by some insertion automaton if and only if there exists a computable function γ_p such that for all executions $\sigma \in pre(\Sigma)$, if $\neg \hat{P}(\sigma)$ then

1. $\forall \sigma' \in \Sigma. \sigma \preceq \sigma' \Rightarrow \neg \hat{P}(\sigma')$, or
2. $\sigma \notin \Sigma$ and $\hat{P}(\sigma; \gamma_p(\sigma))$

Proof (If Direction) We construct an insertion automaton that precisely enforces any such \hat{P} as follows.

- States: $q \in pre(\Sigma) \cup \{\text{end}\}$ (the sequence of actions seen so far, or end if the automaton will stop on the next step)

- Start state: $q_0 = \cdot$ (the empty sequence)
- Transition and insertion functions (for simplicity, we combine δ and γ): Consider processing the action a in state q .
 - (A) If $q = \text{end}$ then stop (i.e., δ and γ are undefined, so rule I-STOP applies).
 - (B) Else if $q = \sigma$ and $\hat{P}(\sigma; a)$ then emit a and continue in state $\sigma; a$.
 - (C) Else if $q = \sigma$ and $\neg\hat{P}(\sigma; a)$ and $\sigma; a \notin \Sigma$ and $\hat{P}(\sigma; a; \gamma_p(\sigma; a))$ then emit a and continue in state $\sigma; a$
 - (D) Else if $q = \sigma$ and $\neg\hat{P}(\sigma; a)$ and $\forall\sigma' \in \Sigma.\sigma; a \preceq \sigma' \Rightarrow \neg\hat{P}(\sigma')$ and $\hat{P}(\sigma)$ then stop.
 - (E) Else if $q = \sigma$ and $\neg\hat{P}(\sigma; a)$ and $\forall\sigma' \in \Sigma.\sigma; a \preceq \sigma' \Rightarrow \neg\hat{P}(\sigma')$ and $\neg\hat{P}(\sigma)$ then insert $\gamma_p(\sigma)$ and continue in state end .

This automaton halts on all inputs $\sigma \in \Sigma$ because only decidable boolean predicates need to be evaluated during execution. Note that if transitions (A), (B), and (C) are not followed, then by clauses (1) and (2) given in the statement of this theorem, it must be the case that $\forall\sigma' \in \Sigma.\sigma; a \preceq \sigma' \Rightarrow \neg\hat{P}(\sigma')$. Therefore, after eliminating transitions (A), (B), and (C), the only test that needs to be performed in order to determine which of transitions (D) and (E) to apply is whether $\hat{P}(\sigma)$, which is assumed to be a decidable predicate (possibly in contrast to the predicate $\forall\sigma' \in \Sigma.\sigma; a \preceq \sigma' \Rightarrow \neg\hat{P}(\sigma')$).

If σ is the input so far, the automaton maintains the following invariant $\text{INV}_p(q)$.

- If $q = \sigma$ then σ has been emitted and either $\hat{P}(\sigma)$ or $(\neg\hat{P}(\sigma)$ and $\sigma \notin \Sigma$ and $\hat{P}(\sigma; \gamma_p(\sigma))$).
- Otherwise, $q = \text{end}$ and the automaton has emitted $\sigma; \gamma_p(\sigma)$ and $\hat{P}(\sigma; \gamma_p(\sigma))$.

Initially, $\text{INV}_p(q_0)$ since our definition of a property assumes $\hat{P}(\cdot)$ for all properties. A simple inductive argument on the length of the input σ suffices to show that the invariant is maintained for all inputs.

Given this invariant, it is straightforward to show that the automaton processes every input $\sigma \in \Sigma$ properly and precisely enforces \hat{P} . There are two cases.

- **Case $\hat{P}(\sigma)$:**
 Consider any prefix $\sigma[..i]$. By induction on i , we show that the automaton accepts $\sigma[..i]$ without stopping, inserting any actions, or moving to the state end .
 - If $\hat{P}(\sigma[..i])$ then the automaton accepts this prefix using transition (B) and continues.

- If $\neg\hat{P}(\sigma[..i])$ then since $\sigma[..i] \preceq \sigma$ and $\hat{P}(\sigma)$, clause (1) in the theorem statement cannot be true of $\sigma[..i]$, so $\sigma[..i] \notin \Sigma$ and $\hat{P}(\sigma[..i]; \gamma_p(\sigma[..i]))$. Hence, the automaton accepts this prefix using transition (C) and continues.

- **Case $\neg\hat{P}(\sigma)$:**

The automaton may halt by applying transitions (A) or (D), or by running out of input actions to process. Thus, the final state prior to halting must be one of the following.

- end, if transition (A) is applied
- σ' such that $\sigma' \preceq \sigma$ and $\hat{P}(\sigma')$, if transition (D) is applied
- σ , if the automaton ran out of input actions

In the first case, $\text{INV}_p(\text{end})$ ensures that the automaton's output is in the property. The second case dictates that $\hat{P}(\sigma')$, and $\text{INV}_p(\sigma')$ implies that σ' is the automaton's output. The third case cannot occur because $\neg\hat{P}(\sigma)$ and $\sigma \in \Sigma$ contradict $\text{INV}_p(\sigma)$. Therefore, the automaton always outputs some sequence that is in the property.

(Only-If Direction) To show that any property \hat{P} precisely enforced by an insertion automaton I conforms to the constraints outlined in the theorem, we define a decidable algorithm for computing $\gamma_p(\sigma)$ for all $\sigma \in \text{pre}(\Sigma)$. This definition follows.

1. Run insertion automaton I on input σ to completion.
2. If I does not output exactly σ on input σ then let $\gamma_p(\sigma) = \cdot$.
3. Otherwise, begin searching for a $\tau \in \mathcal{A}^*$ such that $\hat{P}(\sigma; \tau)$. Once the first such τ is found, let $\gamma_p(\sigma) = \tau$.

Steps (1) and (2) require only finite time because I is a precise enforcer and as such must halt within finite time on all inputs (since $\forall \sigma \in \Sigma. \exists q'. \exists \sigma' \in \mathcal{A}^*. (\sigma, q_0) \xrightarrow{\sigma'}_I (\cdot, q')$). Now consider step (3), which occurs when I outputs exactly σ on input σ . Because $\sigma \in \text{pre}(\Sigma)$, there must be some $\tau' \in \Sigma$ such that $\sigma \preceq \tau'$. Moreover, since I enforces \hat{P} , I must output some sequence σ' on input τ' such that $\hat{P}(\sigma')$, but because $(\sigma, q_0) \xrightarrow{\sigma}_I (\cdot, q')$, $(\tau', q_0) \xrightarrow{\sigma'}_I (\cdot, q'')$, and $\sigma \preceq \tau'$, we have $\sigma' = \sigma; \tau$ for some $\tau \in \mathcal{A}^*$. Thus guaranteed that at least one τ exists such that $\hat{P}(\sigma; \tau)$, we can search \mathcal{A}^* for this τ in finite time because \mathcal{A}^* is countable (though possibly countably infinite). This algorithm therefore halts on all inputs $\sigma \in \text{pre}(\Sigma)$, meeting the constraint that γ_p be a computable function.

Next, we consider any arbitrary $\sigma \in \text{pre}(\Sigma)$ such that $\neg\hat{P}(\sigma)$ and show that one of the following must hold.

1. $\forall \sigma' \in \Sigma. \sigma \preceq \sigma' \Rightarrow \neg\hat{P}(\sigma')$, or

2. $\sigma \notin \Sigma$ and $\hat{P}(\sigma; \gamma_p(\sigma))$

Consider any case in which (1) does not hold. Then, $\exists \sigma' \in \Sigma. \sigma \preceq \sigma' \wedge \hat{P}(\sigma')$, so I must output exactly σ on input σ (because σ may later be extended to σ' , which must be precisely enforced by accepting every action in σ'). By the definition of $\gamma_p(\sigma)$, then, $\hat{P}(\sigma; \gamma_p(\sigma))$. Also, since $(\sigma, q_0) \xrightarrow{\sigma}_I (\cdot, q')$, it must be that case that $\sigma \notin \Sigma$ (if $\sigma \in \Sigma$ then $\hat{P}(\sigma)$ because I is a precise enforcer, contrary to the assumption that $\neg \hat{P}(\sigma)$). Thus, (2) must hold whenever (1) does not. ■

Theorem 9 (Precise S-Enforceable Subset of Precise I-Enforceable)

If there exists a suppression automaton $S = (Q_s, q_{0_s}, \delta_s, \omega)$ that precisely enforces \hat{P} on some system (\mathcal{A}, Σ) then there exists an insertion automaton $I = (Q_i, q_{0_i}, \delta_i, \gamma)$ that precisely enforces \hat{P} on the same system.

Proof We construct I as follows and then show that it precisely enforces \hat{P} .

- States: $Q_i = pre(\Sigma) \cup \{\text{end}\}$ (a state $q \in Q_i$ represents the sequence of actions seen so far, or **end** if the automaton will stop on the next step)
- Start state: $q_{0_i} = \cdot$ (the empty sequence)
- Transition and insertion functions (we combine δ_i and γ_i for simplicity): Consider processing the action a in state q .
 - (A) If $q = \text{end}$ then halt.
 - (B) If $q = \sigma$ and $(\sigma; a, q_{0_s}) \xrightarrow{\sigma; a}_S (\cdot, q')$ for some q' then accept a and continue in state $\sigma; a$.
 - (C) If $q = \sigma$ and $(\sigma; a, q_{0_s}) \xrightarrow{\sigma}_S (\cdot, q')$ for some q' then insert the first $\tau \in \mathcal{A}^*$ found such that $\hat{P}(\sigma; \tau)$ and continue in state **end**.

We show that I precisely enforces \hat{P} by considering two cases of $\sigma \in \Sigma$. If $\hat{P}(\sigma)$ then because S precisely enforces \hat{P} , S must accept every action in σ . Therefore, the boolean guard on transition (B) in insertion automaton I will always evaluate to true (in finite time because S is a precise enforcer and as such must halt on all inputs in finite time), so I will also accept every action in σ , thereby correctly precisely enforcing \hat{P} on any $\sigma \in \Sigma$ such that $\hat{P}(\sigma)$.

In the other case, $\neg \hat{P}(\sigma)$, implying that S must somehow edit σ so that its output is in the property. Let $\sigma[j]$ be the first action in σ that is not accepted by S , implying that S (and therefore I) accepts and emits $\sigma[..j-1]$. Since S precisely enforces \hat{P} , there must exist some $\tau \in \mathcal{A}^*$ such that $\hat{P}(\sigma[..j-1]; \tau)$, so that S can emit $\sigma[..j-1]; \tau$ on input σ (in the case of S halting at $\sigma[j]$, we have $\tau = \cdot$). Because S 's first suppression or halting occurs at $\sigma[j]$, insertion automaton I must follow transition (C) when it reaches $\sigma[j]$, and because there must exist some $\tau \in \mathcal{A}^*$ such that $\hat{P}(\sigma[..j-1]; \tau)$, automaton I will find and insert the requisite τ in finite time (since \mathcal{A}^* must be countable, though possibly countably infinite). Therefore, the final output of I on input σ (i.e., $\sigma[..j-1]; \tau$)

is in the property, as required to precisely enforce \hat{P} . ■

Theorem 10 (Effective I-Enforcement)

A property \hat{P} on the system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be effectively enforced by some insertion automaton if and only if there exists a computable insertion-rewrite function γ^* such that for all executions $\sigma \in \Sigma$,

1. If $\neg\hat{P}(\sigma)$ then $\hat{P}(\gamma^*(\sigma))$
2. If $\hat{P}(\sigma)$ then $\sigma \cong \gamma^*(\sigma)$

Proof (If Direction) We construct an insertion automaton that effectively enforces any such \hat{P} as follows.

- States: $q \in (\text{pre}(\Sigma) \times \{+, -\}) \cup \{\text{end}\}$ (the sequence of actions seen so far paired with + (-) to indicate that the automaton did not (did) most recently step via rule I-Ins, or end if the automaton will stop on the next step)
- Start state: $q_0 = \langle \cdot, + \rangle$ (for the empty sequence)
- Transition and insertion functions (for simplicity, we combine δ and γ): Consider processing the action a in state q .
 - (A) If $q = \text{end}$ then stop.
 - (B) If $q = \langle \sigma; a, - \rangle$ then emit the action a and continue in state $\langle \sigma; a, + \rangle$.
 - (C) If $q = \langle \sigma, + \rangle$ and $\gamma^*(\sigma; a) = \gamma^*(\sigma); \tau; a$ then insert τ and continue in state $\langle \sigma; a, - \rangle$.
 - (D) If $q = \langle \sigma, + \rangle$ and $\gamma^*(\sigma; a) = \gamma^*(\sigma); \tau$ and $\forall \tau' \in \mathcal{A}^*. \tau \neq \tau'; a$ and $H(\gamma^*, \sigma; a)$ then insert τ and continue in state **end**.

The automaton maintains the following invariant $\text{INV}_p(q)$.

- If $q = \langle \sigma, + \rangle$ then σ is the input so far and the automaton has emitted $\gamma^*(\sigma)$.
- If $q = \text{end}$ then σ is the input so far and the automaton has emitted $\gamma^*(\sigma)$ and $H(\gamma^*, \sigma)$.
- Otherwise, $q = \langle \sigma; a, - \rangle$ and $\sigma; a$ is the input so far and $\gamma^*(\sigma; a) = \gamma^*(\sigma); \tau; a$ and the automaton has emitted $\gamma^*(\sigma); \tau$.

Initially, $\text{INV}_p(q_0)$ because $\gamma^*(\cdot) = \cdot$, and a simple inductive argument on the length of the input σ suffices to show that the invariant is maintained for all inputs.

We now show that the automaton emits $\gamma^*(\sigma)$ when the input is σ . There are two cases, derived by inspection of the insertion automaton. In the first case, the automaton halts in state $\langle \sigma, + \rangle$, so by $\text{INV}_p(\sigma)$, $\gamma^*(\sigma)$ has been emitted.

Otherwise, the automaton must halt in state `end`, having only seen input σ' such that $\sigma' \preceq \sigma$. In this case, $\text{INV}_p(\text{end})$ implies that $\gamma^*(\sigma')$ has been emitted when the automaton halts and $H(\gamma^*, \sigma')$. However, $H(\gamma^*, \sigma')$ and $\sigma' \preceq \sigma$ imply that $\gamma^*(\sigma') = \gamma^*(\sigma)$, so the automaton has in fact emitted $\gamma^*(\sigma)$ when it halts in this case as well. Note that the automaton cannot halt in state $\langle \sigma; a, - \rangle$ because the next step will always be to emit a and continue in state $\langle \sigma; a, + \rangle$ (using transition (B)).

Analogously to the argument given in the Effective T-Enforcement theorem, because the automaton always outputs $\gamma^*(\sigma)$ on input σ , clauses (1) and (2) in the theorem statement ensure that the automaton effectively enforces \hat{P} .

(Only-If Direction) Define $\gamma^*(\sigma)$ to be whatever sequence is emitted by the insertion automaton on input σ . By this definition, $\gamma^*(\sigma)$ is a computable function because the automaton is an effective enforcer and as such acts deterministically and halts on all inputs in the set of executions (Σ) . We first show that this is an insertion-rewrite function. Clearly, $\gamma^*(\cdot) = \cdot$. When processing some action a after having already processed any sequence σ , the automaton may step via I-STEP, I-INS, or I-STOP.

- **Case I-Step:** Here the automaton emits whatever has been emitted in processing σ (by definition, this is $\gamma^*(\sigma)$), followed by a . Thus, $\gamma^*(\sigma; a) = \gamma^*(\sigma); \tau; a$, where $\tau = \cdot$.
- **Case I-Stop:** Here the automaton emits only what has already been emitted, so $\gamma^*(\sigma; a) = \gamma^*(\sigma); \tau$, where $\tau = \cdot$ and $\forall \tau' \in \mathcal{A}^*. \tau \neq \tau'; a$. Because the automaton halts, it may not examine any remaining input, implying that $H(\gamma^*, \sigma; a)$.
- **Case I-Ins:** Because the automaton effectively enforces \hat{P} , it may only insert a finite sequence of symbols before either accepting or halting on the current action. That is, although the automaton may apply rule I-INS multiple times in succession, the end result must be the insertion of only a finite sequence of actions, followed by acceptance or termination. In the case of ultimate acceptance, the net output of the automaton, $\gamma^*(\sigma; a)$, is $\gamma^*(\sigma); \tau; a$. In the case of ultimate termination, no more input may be examined, so $H(\gamma^*, \sigma; a)$ and either $\gamma^*(\sigma; a) = \gamma^*(\sigma); \tau; a$ (if a was the last inserted action), or $\gamma^*(\sigma; a) = \gamma^*(\sigma); \tau$ and $\forall \tau' \in \mathcal{A}^*. \tau \neq \tau'; a$ (if a was not the last inserted action).

In any case, the insertion automaton's output adheres to the definition of an insertion-rewrite function.

Now consider an arbitrary $\sigma \in \Sigma$. By the definition of effective enforcement and that $\gamma^*(\sigma)$ is defined to be whatever sequence is emitted by the insertion automaton on input σ , $\hat{P}(\gamma^*(\sigma))$ and $\hat{P}(\sigma) \Rightarrow \sigma \cong \gamma^*(\sigma)$. ■

Theorem 11 (Uniform Precise E-Enforcement)

A property \hat{P} on the uniform system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be precisely enforced by an edit automaton if and only if $\forall \sigma \in \text{pre}(\Sigma). \neg \hat{P}(\sigma) \Rightarrow \forall \sigma' \in \Sigma. (\sigma \preceq \sigma' \Rightarrow \neg \hat{P}(\sigma'))$.

Proof Both directions proceed completely analogously to the proof of Uniform Precise T-Enforcement. ■

Theorem 12 (Nonuniform Precise E-Enforcement)

A property \hat{P} on the system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be precisely enforced by some edit automaton if and only if there exists a computable function γ_p such that for all executions $\sigma \in \text{pre}(\Sigma)$, if $\neg \hat{P}(\sigma)$ then

1. $\forall \sigma' \in \Sigma. \sigma \preceq \sigma' \Rightarrow \neg \hat{P}(\sigma')$, or
2. $\sigma \notin \Sigma$ and $\hat{P}(\sigma; \gamma_p(\sigma))$

Proof (If Direction) By the Nonuniform Precise I-Enforcement theorem, given a function γ_p and a property \hat{P} satisfying the requirements stated above, we may build an insertion automaton $I = (Q, q_0, \delta, \gamma)$ to enforce \hat{P} . Then, we can construct an edit automaton $E = (Q, q_0, \delta, \gamma, \omega)$, where ω is defined to be + over all of its domain (the same domain as δ). Clearly, E and I enforce the same property because whenever I steps via I-STEP, I-INS, or I-STOP, E respectively steps via E-STEP, E-INS, or E-STOP, while emitting exactly the same actions as I . Because I precisely enforces \hat{P} , so too must E .

(Only-If Direction) This direction proceeds completely analogously to the Only-If Direction of the proof of Nonuniform Precise I-Enforcement. ■

Theorem 13 (Effective E-Enforcement)

Any property \hat{P} on the system $\mathcal{S} = (\mathcal{A}, \Sigma)$ can be effectively enforced by some edit automaton.

Proof We construct an edit automaton that effectively enforces any \hat{P} as follows.

- States: $q \in \text{pre}(\Sigma) \times \mathcal{A}^* \times \{+, -\}$ (the sequence of actions seen so far, the actions seen but not emitted, and + (-) to indicate that the automaton must not (must) suppress the current action)
- Start state: $q_0 = \langle \cdot, \cdot, + \rangle$ (for the empty sequence)
- Transition, insertion, and suppression functions (for simplicity, we combine δ , γ , and ω):
Consider processing the action a in state q .
 - (A) If $q = \langle \sigma, \tau, + \rangle$ and $\neg \hat{P}(\sigma; a)$ then suppress a and continue in state $\langle \sigma; a, \tau; a, + \rangle$.

- (B) If $q = \langle \sigma, \tau, + \rangle$ and $\hat{P}(\sigma; a)$ then insert $\tau; a$ and continue in state $\langle \sigma; a, \cdot, - \rangle$.
- (C) If $q = \langle \sigma; a, \cdot, - \rangle$ then suppress a and continue in state $\langle \sigma; a, \cdot, + \rangle$.

The automaton maintains the following invariant $\text{INV}_p(q)$.

- If $q = \langle \sigma, \tau, + \rangle$ then σ is the input so far and $\sigma = \sigma'; \tau$ (for some $\sigma' \in \mathcal{A}^*$) and the automaton has emitted σ' and σ' is the longest prefix of σ such that $\hat{P}(\sigma')$.
- Otherwise, $q = \langle \sigma; a, \cdot, - \rangle$ and $\sigma; a$ is the input so far and $\sigma; a$ has been emitted and $\hat{P}(\sigma; a)$.

Initially, $\text{INV}_p(q_0)$ because $\hat{P}(\cdot)$, and a simple inductive argument on the length of the input σ suffices to show that the invariant is maintained for all inputs.

By inspection, we note that the edit automaton cannot halt in state $\langle \sigma; a, \cdot, - \rangle$ because the next step will always be to suppress a and continue in state $\langle \sigma; a, \cdot, + \rangle$ (using transition (C)). Therefore, the automaton always halts in some state $\langle \sigma, \tau, + \rangle$ on input σ , so $\text{INV}_p(\langle \sigma, \tau, + \rangle)$ ensures that the automaton emits some σ' such that $\hat{P}(\sigma')$ on all inputs. In addition, $\text{INV}_p(\langle \sigma, \tau, + \rangle)$ implies that if $\hat{P}(\sigma)$ then $\sigma = \sigma'$, so the automaton must emit exactly σ on input σ whenever $\hat{P}(\sigma)$. This edit automaton thus effectively enforces \hat{P} . ■