

Runtime Enforcement with Partial Control

Raphaël Khoury^(✉) and Sylvain Hallé

Laboratoire d'informatique formelle, Département d'informatique et de mathématique, Université du Québec à Chicoutimi, Chicoutimi, Canada
raphael.khoury@uqac.ca, shalle@acm.org

Abstract. This study carries forward the line of enquiry that seeks to characterize precisely which security policies are enforceable by runtime monitors. In this regard, Basin et al. recently refined the structure that helps distinguish between those actions that the monitor can potentially suppress or insert in the execution, from those that the monitor can only observe. In this paper, we generalize this model by organizing the universe of possible actions in a lattice that naturally corresponds to the levels of monitor control. We then delineate the set of properties that are enforceable under this paradigm and relate our results to previous work in the field. Finally, we explore the set of security policies that are enforceable if the monitor is given greater latitude to alter the execution of its target, which allows us to reflect on the capabilities of different types of monitors.

1 Introduction

Runtime monitoring is an approach to enforcing security policies that seeks to allow untrusted code to run safely by observing its execution and reacting as needed to prevent a violation of a user-supplied security policy. This method of ensuring the security of code is rapidly gaining acceptance in practice and several implementations exist [18]. One question seems to recur frequently in multiple studies: exactly which set of properties are *monitorable*, in the sense that they are enforceable by monitors. Previous research has identified several factors that can affect the set of security policies enforceable by monitors. These include the means at the disposal of monitors to react to a potential violation of the security policy [3], the availability of statically gathered data about the target program's possible executions [3, 9], memory and computability constraints [13, 14] etc.

One specific aspect that can have a considerable impact on the monitor's expressiveness is its ability to either suppress certain actions performed by the target program from occurring during the execution (while allowing the remainder of the execution to continue unaffected) or to insert additional events in an ongoing execution. These abilities, when available, extend the monitor's enforcement power considerably. Indeed, a lower bound on the enforcement power of

S. Hallé—The authors gratefully acknowledge the financial support of the Natural Sciences and Engineering Research Council of Canada (NSERC).

monitors is given by Schneider [26] who shows that the set of properties enforceable by a monitor whose only possible reaction to a potential violation of the desired security policy is to abort the execution coincides with the set of safety properties. Conversely, Ligatti et al. [23] consider the case of a monitor with an unlimited ability to delay any event performed by the target program until it has ascertained that its occurrence in the execution would not violate the security policy. In effect, the monitor is simulating the execution of the program until it is certain that the behaviour it has so far witnessed is correct. When behaving in this manner, the monitor can enforce a vast range of security properties, termed the set of infinite renewal properties, which includes all safety policies, some liveness policies and some policies that are neither safety nor liveness.

Yet, it may not be realistic to assume that the monitor has an unlimited ability to simulate the execution of the target program. To this end, Basin et al. propose a middle ground [2]. They partition the set of possible program actions in two disjoint subsets: a set of controllable actions, which the monitor may freely suppress from the execution, and a set of observable actions, whose occurrence the monitor can only observe. This allows for a more precise characterization to the set of monitorable properties. Section 2 will discuss these concepts in more detail.

In Sect. 3, we further generalize this analysis by organizing the set of possible actions along a lattice that distinguishes between four types of atomic actions, namely *controllable* actions (which a monitor can insert or block from an execution), *insertable* actions (which a monitor can add to the execution but not suppress), *suppressible* actions (the converse) and *observable* actions (which the monitor can only observe). We then delineate the set of properties that are enforceable under this paradigm and relate our results to previous work in the field.

Finally, we explore in Sect. 4 the set of security policies that are enforceable if the monitor is given greater latitude to alter the execution of its target, rather than be bounded to return a syntactically identical execution sequence if the original execution is valid. In particular, we consider a monitor which can *add* any action into the execution, but cannot prevent any action from occurring if the target program requests it. We also consider a monitor can *remove* potentially malicious actions performed by the target program but cannot add any action to the execution. We show how both can be handled by our model by simply considering a different equivalence relation between traces.

2 Preliminaries

2.1 Executions

Executions are modelled as sequences of atomic actions taken from a finite or countably infinite set of actions Σ . The empty sequence is noted ϵ , the set of all finite length sequences is noted Σ^* , that of all infinite length sequences is noted Σ^ω , and the set of all possible sequences is noted $\Sigma^\infty = \Sigma^\omega \cup \Sigma^*$. Let $\tau \in \Sigma^*$ and $\sigma \in \Sigma^\infty$ be two sequences of actions. We write $\tau;\sigma$ for the concatenation of τ and σ . We say that τ is a prefix of σ noted $\tau \preceq \sigma$, or equivalently $\sigma \succeq \tau$ iff there exists a sequence σ' such that $\tau;\sigma' = \sigma$. Let $\tau, \sigma \in \Sigma^\infty$ be a sequence, we write

acts (σ) for the set of actions present in σ . We write $res_{\hat{P}}(\sigma)$ for the residual of \hat{P} with regard to σ , i.e. the set of sequence $S \subseteq \Sigma^\infty$ s.t. $\forall \tau \in S : \sigma; \tau \in \hat{P}$. Finally, let $\tau, \sigma \in \Sigma^\infty$, τ is said to be a suffix of σ iff there exists a $\sigma' \in \Sigma^*$ such that $\sigma = \sigma'; \tau$.

Following [23], if σ has already been quantified, we freely write $\forall \tau \preceq \sigma$ (resp. $\exists \tau \preceq \sigma$) as an abbreviation to $\forall \tau \in \Sigma^* : \tau \preceq \sigma$ (resp. $\exists \tau \in \Sigma^* : \tau \preceq \sigma$). Likewise, if τ has already been quantified, $\forall \sigma \in \Sigma^\infty : \sigma \succeq \tau$ (resp. $\exists \sigma \in \Sigma^\infty : \sigma \succeq \tau$) can be abbreviated as $\forall \sigma \succeq \tau$ (resp. $\exists \sigma \succeq \tau$).

Let τ be a sequence and a be an action, we write $\tau \setminus a$ for the left cancellation of a from τ , which is defined as the removal from τ of the first occurrence of a . Formally:

$$a; \tau \setminus a' = \begin{cases} \tau & \text{if } a = a'; \\ a; (\tau \setminus a') & \text{otherwise} \end{cases}$$

Observe that $\epsilon \setminus a = \epsilon$. Abusing the notation, we write $\tau \setminus \tau'$ to denote the sequence obtained by left cancellation of each action of τ' from τ . Formally, $\tau \setminus a; \tau' = (\tau \setminus a) \setminus \tau'$. For example, $abcada \setminus daa = bca$.

A finite word $\tau \in \Sigma^*$ is said to be a subword of a word ω , noted $\tau \triangleleft_\Sigma \omega$, iff $\tau = a_0 a_1 a_2 a_3 \dots a_k$ and $\omega = \sigma_0 a_0 \sigma_1 a_1 \sigma_2 a_2 \sigma_3 a_3 \dots \sigma_k a_k v$ with $\sigma_0, \sigma_1, \sigma_2 \dots \in \Sigma^*$ and $v \in \Sigma^\infty$. Let τ, σ be sequences form Σ^* . We write $cs_\tau(\sigma)$ to denote the longest subword of τ which is also a subword of σ . For any $\tau \neq \epsilon$, $\tau.last$ denotes the last action of sequence τ .

2.2 Security Policies and Security Properties

A security policy P is a property iff it can be characterized as a set of sequences for which there exists a decidable predicate \hat{P} over the executions of $\Sigma^\infty : \hat{P}(\sigma)$ iff σ is in the policy [26]. In other words, a property is a policy for which the membership of any sequence can be determined by examining only the sequence itself¹. Such a sequence is said to be *valid* or to *respect* the property. Since, by definition, all policies enforceable by monitors are properties, P and \hat{P} are used interchangeably in our context. Additionally, since the properties of interest represent subsets of Σ^∞ , we follow the common usage in the literature and freely use \hat{P} to refer to these sets.

A number of classes of properties have been defined in the literature and are of special interest in the study of monitoring. First are *safety* properties [21], which proscribe the occurrence of a certain “bad thing” during the execution. Formally, let Σ be a set of actions and \hat{P} be a property. \hat{P} is a *safety* property iff

$$\forall \sigma \in \Sigma^\infty : \neg \hat{P}(\sigma) \Rightarrow \exists \sigma' \preceq \sigma : \forall \tau \succeq \sigma' : \neg \hat{P}(\tau) \quad (\text{safety})$$

¹ Security policies whose enforcement necessitates the examination of multiples execution sequences, such as noninterference policies, are not generally enforceable by monitors.

Informally, this states that any sequence does not respect the security property if there exists a prefix of that sequence from which any possible extension does not respect the security policy. This implies that a violation of a safety property is irremediable: once a violation occurs, nothing can be done to correct the situation.

Alternatively, a *liveness* property [1] is a property prescribing that a certain “good thing” must occur in any valid execution. Formally, for an action set Σ and a property \hat{P} , \hat{P} is a liveness property iff

$$\forall \sigma \in \Sigma^* : \exists \tau \in \Sigma^\infty : \tau \succeq \sigma \wedge \hat{P}(\tau) \quad (\text{liveness})$$

Informally, the definition states that a property is a liveness property if any finite sequence can be extended into a valid sequence.

Another class of security properties that are of interest is that of *renewal* properties [23]. A property is in renewal if every infinite valid sequence has infinitely many valid prefixes, while every infinite invalid sequence has only finitely many such prefixes. Observe that every property over finite sequences is in infinite renewal. The set of renewal properties is equivalent to the set of response properties in the safety-progress classification [10].

$$\forall \hat{P} \subseteq \Sigma^\omega : \hat{P}(\sigma) \Leftrightarrow \exists \sigma' \preceq \sigma : \exists \tau \preceq \omega : \tau \succeq \sigma' \wedge \hat{P}(\tau). \quad (\text{renewal})$$

It is often useful to restrict our analysis to properties for which the empty sequence ϵ is valid. Such properties are said to be *reasonable* [23]. Formally,

$$\forall \hat{P} \subseteq \Sigma^\infty : \hat{P}(\epsilon) \Leftrightarrow \hat{P} \text{ is reasonable} \quad (\text{reasonable})$$

In the remainder of this paper, we will only consider reasonable properties. Furthermore, in order to avoid having the main topic of this paper be sidestepped by decidability issues, will consider that $\hat{P}(\sigma)$ is decidable for all properties and all execution sequences. Likewise, we also consider that other predicates or functions over sequences are decidable.

2.3 Security Property Enforcement

Finally, we need to provide a definition of what it means to “enforce” a security property \hat{P} . A number of possible definitions have been suggested. The most widely used is $\text{effective}_{\preceq}$ enforcement [3]. Under this definition, a property is $\text{effective}_{\preceq}$ enforced iff the following two criterion are respected.

1. Soundness: All observable behaviours of the target program respect the desired property, i.e. every output sequence is present in the set of executions defined by \hat{P} .
2. Transparency: The semantics of valid executions is preserved, i.e. if the execution of the unmonitored program already respects the security property, the monitor must output an equivalent sequence, with respect to an equivalence relation $\cong \subseteq \Sigma^\infty \times \Sigma^\infty$.

Syntactic equality is the most straightforward equivalence relation, and the one that has been the most studied in the literature. It models the behaviour of a monitor that enforces the desired property by suppressing (or simulating) part of the execution, only allowing it to be output when it has ascertained that the execution up to that point is valid. In Sect. 4, we consider two alternative notions of equivalences, which can be used to characterize alternative behaviours on the part of the enforcement mechanism.

2.4 Related Work

Initial work on the question of delineating which security policies are or are not enforceable by monitor was performed by Schneider [26]. He considered the capabilities of a monitor that observes the execution of its target, with no knowledge of its possible future behaviour and no means to affect the target except by aborting the execution. Each time the target program attempts to perform an action, the monitor has to either accept it immediately, or abort the execution. Under these constraints, the set of properties enforceable by monitors coincides with the set of *safety* properties.

Ligatti et al. [22] extend Schneider’s modelling of monitors along three axes:

1. According to the means at the disposal of the monitor to react to a potential violation of the security policy. These include truncating the execution, inserting new actions into the execution, suppressing some part of the executions or both inserting and suppressing actions
2. According to the availability of statically gathered data describing the target program possible execution paths
3. According to how much latitude the monitor is given to alter executions that already respect the security policy.

By combining these three criteria, they build a rich taxonomy of enforceable properties, and contrast the enforcement power of different types of monitors.

Basin et al. [2] generalize Schneider’s model by distinguishing between observable actions, whose occurrence the monitor cannot prevent, and controllable actions, which the monitor can prevent from occurring by aborting the execution.

The enforcement power of monitor operating with memory constraints is studied in [4, 13, 27–29].

The computability constraints that can further restrict a monitor’s enforcement power are discussed in [14, 20]; that of monitors relying upon an *a priori* model of the program’s possible behaviour is discussed in [9, 22].

Falcone et al. [11, 12] show that the set of infinite renewal properties coincides with the union of four of the 6 classes of the safety-progress classification of security properties [10]. Khoury and Tawbi [16, 17] and Bielova et al. [5–7] further refine the notion of enforcement by suggesting alternative definitions of enforcement. In [24] Ligatti and Reddy introduced an alternative model, the mandatory-result automaton. This model distinguishes between the action set of the target and that of the system with which it interacts. This distinction

makes it easier to study the interaction between the target program, the monitor and the system. A thorough survey of the question of enforceable properties by monitors is provided in [18].

3 Monitoring with Partial Control

The previous works each consider monitors where actions belong to particular sets. For example, Schneider’s model assumes that all actions can be suppressed by the monitor; conversely, Ligatti et al. assume that all actions can be indefinitely delayed. Basin et al. propose a middle ground where every action can either be freely suppressed, or can only be observed. In this section, we define a generalized model of actions where each of these works becomes a particular case. We then study what properties are enforceable in this generalized model.

3.1 A Lattice of Actions

We organize the set of possible actions along a lattice that distinguishes between four types of atomic actions: namely controllable actions (\mathcal{C}), insertable actions (\mathcal{I}), suppressible actions (\mathcal{D} —for delete) and observable actions (\mathcal{O}), as is shown in Fig. 1.

- Controllable actions (\mathcal{C}) are the basic actions such as opening a file or sending data on the network, which the monitor can either insert into the execution or prevent from occurring if they violate the security policy. In Ligatti’s model, all actions are controllable.
- Insertable actions (\mathcal{I}) can be added into the execution but not suppressed if they are present. An example of an insertable action is an additional delay before processing a request to bring it into compliance with a resource usage policy. The monitor may add such actions to the execution sequence but cannot remove them if the target program executes them.
- Suppressible actions (\mathcal{D} —for delete) are those actions that the monitor can prevent from occurring, but cannot insert into the execution if the target program does not request them. Sending an email, decrypting a file or receiving a user input are all examples of suppressible actions. In Schneider’s model, all actions are suppressible.
- Observable actions \mathcal{O} can only be observed by the monitor, which can neither insert them into the execution when they are not present nor suppress them if they occur. In Basin et al.’s model, all actions are either suppressible or observable.

As the examples above illustrate, we believe that the lattice model we propose is a more realistic description of the reality encountered by the monitor, and will thus allow a more precise characterization of the set of enforceable properties. Observe that the monitor may only abort the execution if the next action is in $\mathcal{C} \cup \mathcal{D}$. Moreover, the sets $\mathcal{C}, \mathcal{I}, \mathcal{D}, \mathcal{O}$ are disjoint and that the universe of possible program actions is $\Sigma = \mathcal{C} \cup \mathcal{I} \cup \mathcal{D} \cup \mathcal{O}$.

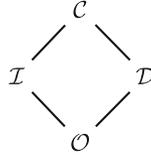


Fig. 1. The lattice of possible actions

The following notation is useful for comparing different enforcement mechanisms. Let Σ be a universe of actions and let \mathcal{L} be a lattice over the set Σ as described in Sect. 2. Let $\mathcal{S} \subseteq \Sigma^\infty$ stand for a subset of possible execution sequences and let $\cong \subseteq \Sigma^\infty \times \Sigma^\infty$ be an equivalence relation. We write $\mathcal{L}^{\mathcal{S}}$ -enforceable $_{\cong}$ to denote the set of proprieties that are enforceable $_{\cong}$ by a monitor when the set of possible sequences is \mathcal{S} and the set of possible actions is organized alongside lattice \mathcal{L} .

Let \mathcal{L} be a lattice as described above. We write $\mathcal{L}_{\mathcal{O}}$ (resp. $\mathcal{L}_{\mathcal{I}}$, $\mathcal{L}_{\mathcal{D}}$, $\mathcal{L}_{\mathcal{C}}$) for the set \mathcal{O} (resp. \mathcal{I} , \mathcal{D} , \mathcal{C}) in \mathcal{L} . We write $\mathcal{L} = \langle \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4 \rangle$ for the lattice where $\mathcal{L}_{\mathcal{O}} = \Sigma_1$, $\mathcal{L}_{\mathcal{I}} = \Sigma_2$, $\mathcal{L}_{\mathcal{D}} = \Sigma_3$ and $\mathcal{L}_{\mathcal{C}} = \Sigma_4$. Let $A, B \in \{\mathcal{O}, \mathcal{I}, \mathcal{D}, \mathcal{C}\}$ and let $a \in \Sigma$, we write $\mathcal{L}_{A \xrightarrow{a} B}$ to indicate the lattice \mathcal{L}' defined such that $\mathcal{L}'_A = \mathcal{L}_A \setminus \{a\}$, $\mathcal{L}'_B = \mathcal{L}_B \cup \{a\}$ and $\forall C \in \{\mathcal{O}, \mathcal{I}, \mathcal{D}, \mathcal{C}\} : C \notin \{A, B\} \Rightarrow \mathcal{L}'_C = \mathcal{L}_C$. In other words, $\mathcal{L}_{A \xrightarrow{a} B}$ is the lattice built by moving only element a from set A to another B , leaving all other sets unchanged.

3.2 Enforceable Properties

We begin by reflecting on the set of properties that are $\mathcal{L}^{\Sigma^\infty}$ -enforceable $_{=}$, i.e. properties that are enforceable if the monitor is bounded to output any valid sequence exactly as it occurs (with syntactic equality as the equivalence relation between valid inputs and the monitor’s output). This is the enforcement paradigm that has been the most studied in the literature. A monitor that seeks to enforce a property in this manner may take any one of three strategies, depending on the desired property and the ongoing execution, and the set of $\mathcal{L}^{\Sigma^\infty}$ -enforceable $_{=}$ properties can be derived by combining the three.

First, using a model in which every action is controllable, Ligatti et al. argued that a monitor can enforce any reasonable renewal property by suppressing the execution until a valid prefix is reached at which point the monitor can output the suffix of the execution it has previously suppressed. We generalized the definition of renewal as follows:

$$\forall \sigma \in \Sigma^\infty : \hat{P}(\sigma) \Leftrightarrow \forall \sigma' \preceq \sigma : (\exists \tau \preceq \sigma : \sigma' \preceq \tau \wedge \hat{P}(\tau) \wedge \forall \tau' \preceq \tau : \tau' \succeq \sigma' \Rightarrow \tau'.last \in \mathcal{C}). \quad (\mathcal{L}\text{-Renewal})$$

Observe that the definition now applies to all sequences in Σ^∞ , rather than just to infinite sequences. The second half of the equation always evaluates to true if all the actions are controllable (Ligatti’s model) and always evaluates

to false if all the actions are suppressible or observable (Schneider or Basin’s models).

Second, Ligatti et al. observe that a property is enforceable if there exists a prefix beyond which there is only one valid extension. In that case, the monitor can abort the execution and output that sequence. This allows some nonsafety properties to be monitored. Ligatti et al. refer to this case as the “corner case” of effective enforcement. We generalize this case as follows:

$$\forall \sigma \in \Sigma^\infty : \hat{P}(\sigma) \vee \exists \sigma' \preceq \sigma : \forall \sigma'; \tau \succeq \sigma' : \hat{P}(\sigma'; \tau) \Rightarrow \sigma'; \tau = \sigma \wedge \sigma'.last \in \mathcal{D} \cup \mathcal{C} \wedge acts(\tau) \subseteq \mathcal{I} \cup \mathcal{C}. \quad (\mathcal{L}\text{-Corner case})$$

Once again, observe that the equation restricts all sequences, rather than only infinite ones. This equation always evaluates to false in the Schneider-Basin model and the second conjunct always evaluates to true in Ligatti’s model.

Finally, the monitor can simply abort the execution if it is irremediably invalid. This is a generalization of the set of safety properties to our framework.

$$\forall \sigma \in \Sigma^\infty : \neg \hat{P}(\sigma) \Rightarrow \exists \tau; a \preceq \sigma : \hat{P}(\tau) \wedge a \in \mathcal{D} \cup \mathcal{C} \wedge \neg \exists \tau' \succeq \tau : \hat{P}(\tau'). \quad (\mathcal{L}\text{-Safety})$$

The set of $\mathcal{L}^{\Sigma^\infty}$ -enforceable₌ properties can now be stated. The definition is not simply the conjunction of the tree preceding set as it must take into account the possibility that enforcement might begin by suppressing and reinserting part of the execution, and then abort the execution using one of the other methods.

Theorem 1

$$\begin{aligned} \hat{P} \in \mathcal{L}^{\Sigma^\infty} - \text{enforceable}_= &\Leftrightarrow \forall \sigma \in \Sigma^\infty : \\ (\hat{P}(\sigma) &\Leftrightarrow \forall \sigma' \preceq \sigma : \exists \tau \preceq \sigma : \sigma' \preceq \tau \wedge \hat{P}(\tau) \wedge (\forall \tau' \preceq \tau : \neg \hat{P}(\tau') \Rightarrow \tau'.last \in \mathcal{C}) \vee \\ (\exists \tau; a \preceq \sigma : a &\in \mathcal{D} \cup \mathcal{C} \wedge \forall \tau' \preceq \tau; a : \neg \hat{P}(\tau') \Rightarrow \tau'.last \in \mathcal{C} \wedge (\forall \tau' \succeq \tau; a : \hat{P}(\tau') \Rightarrow \\ &\tau' = \sigma \vee (\hat{P}(\tau) \wedge \neg \exists \tau' \succeq \tau : \hat{P}(\tau') \wedge acts(\sigma \setminus \tau'; a) \in \mathcal{C} \cup \mathcal{I}))) \end{aligned}$$

Proof. Proofs of the theorems and corollaries in this work are contained in our companion Technical Report [15]. We omit them here due to space constraints.

This definition narrows the set of monitorable properties somewhat, compared with previous work. For example, bounded availability is given in [23] as an example of a monitorable policy. In fact, it is monitorable only if every action that occurs between the acquisition and release of a protected resource is controllable. Likewise, the “no send after read” property described in [26] is only enforceable if every action which might violate the property is deletable or controllable.

While the set-theoretic characterization of enforceable property is somewhat involved, an LTL property can characterize such properties.

Theorem 2. *Let valid be a predicate identifying a valid sequence:*

$$\text{valid}(\sigma) \Leftrightarrow \hat{P}(\sigma)$$

Let cc be a predicate that identifies a sequence in the \mathcal{L} -corner case:

$$\begin{aligned} \text{cc}(\sigma) \Leftrightarrow \hat{P}(\sigma) \wedge \exists \sigma' \preceq \sigma : \forall \sigma' ; \tau \succeq \sigma' : \hat{P}(\sigma' ; \tau) \Rightarrow \\ \sigma' ; \tau = \sigma \wedge \sigma'.\text{last} \in \mathcal{D} \cup \mathcal{C} \wedge \text{acts}(\tau) \subseteq \mathcal{I} \cup \mathcal{C} \end{aligned}$$

Let \mathcal{C} be a predicate identifying a sequence ending on a controllable action, and \mathcal{D} a predicate that identifies a sequence ending on a deletable action:

$$\begin{aligned} \mathcal{C}(\sigma) \Leftrightarrow \sigma.\text{last} \in \mathcal{C} \\ \mathcal{D}(\sigma) \Leftrightarrow \sigma.\text{last} \in \mathcal{D} \end{aligned}$$

Then we have:

$$\begin{aligned} \hat{P} \in \mathcal{L}^{\Sigma^\infty}\text{-enforceable}_= \Leftrightarrow \forall \sigma \in \Sigma^\infty : \\ \mathbf{G}(C \mathbf{W} \text{valid}) \vee (C \mathbf{W} \text{valid} \vee \mathbf{X}((D \vee C) \wedge (\mathbf{G} \neg \text{valid} \vee \text{cc}))) \end{aligned}$$

In this theorem, we assume that any execution sequence from Σ^∞ is possible. In other words, at each step of the execution, the monitor must assume that any action from Σ is a possible next action, or that the execution of the target program could stop. This is called the uniform enforcement context. The monitor often operates in a context where it knows that certain executions are impossible (the nonuniform context). This situation occurs when the monitor benefits from a static analysis of its target, that provides it with a model of the target's possible behaviour. We can adapt the above theorem to take into account the fact that the monitor could operate in a nonuniform context.

Theorem 3

$$\begin{aligned} \hat{P} \in \mathcal{L}^{\mathcal{S}}\text{-enforceable}_= \Leftrightarrow \forall \sigma \in \mathcal{S} : \\ (\hat{P}(\sigma) \Leftrightarrow \forall \sigma' \preceq \sigma : \exists \tau \preceq \sigma : \sigma' \preceq \tau \wedge \hat{P}(\tau) \wedge (\forall \tau' \preceq \tau : \neg \hat{P}(\tau') \Rightarrow \\ (\tau'.\text{last} \in \mathcal{C} \wedge \tau \in \mathcal{S})) \vee \\ (\exists \tau ; a \preceq \sigma : a \in \mathcal{D} \cup \mathcal{C} \wedge \forall \tau' \preceq \tau ; a : (\neg \hat{P}(\tau') \wedge \tau' \in \mathcal{S}) \Rightarrow \tau'.\text{last} \in \mathcal{C} \wedge \\ (\forall \tau' \succeq \tau ; a : \hat{P}(\tau') \Rightarrow \tau' = \sigma \vee (\hat{P}(\tau) \wedge \neg \exists \tau' \succeq \tau : \\ \hat{P}(\tau') \wedge \tau' \in \mathcal{S} \wedge \text{acts}(\sigma \setminus \tau' ; a) \in \mathcal{C} \cup \mathcal{I})))) \end{aligned}$$

We can now restate the results of previous research in our new formalism.

Theorem 4 (from [26]). *If $\mathcal{L} = \langle \emptyset, \emptyset, \Sigma, \emptyset \rangle$ then $\mathcal{L}^{\Sigma^\infty}\text{enforceable}_=$ is Safety.*

Theorem 5 (from [23]). *If $\mathcal{L} = \langle \emptyset, \emptyset, \emptyset, \Sigma \rangle$ then $\mathcal{L}^{\Sigma^\infty}\text{enforceable}_=$ is the union of Infinite Renewal and the corner case.*

3.3 Additional Results

Corollary 1. *Let \mathcal{L} be a lattice over a set of actions Σ as described above and let $\Sigma_1, \Sigma_2 \in \{\mathcal{O}, \mathcal{D}, \mathcal{I}, \mathcal{C}\}$ and $\Sigma_1 \sqsubset \Sigma_2$. $\forall a \in \Sigma : \mathcal{L}_{\Sigma_1}$ -enforceable $_= \subset \mathcal{L}_{\Sigma_1 \xrightarrow{a} \Sigma_2}$ -enforceable.*

Corollary 1 indicates that the set of enforceable properties increases monotonically with the capabilities of the monitor. Thus, any effort made to improve the capabilities of the monitor to control its target is rewarded by a augmented set of enforceable security properties. Conversely, if every action from the set Σ is in \mathcal{O} , only the inviolable property is enforceable.

Corollary 2. *Let $\mathcal{L} = \langle \Sigma, \emptyset, \emptyset, \emptyset \rangle$, $\hat{P} \in \mathcal{L}_{\Sigma}$ -enforceable $\Leftrightarrow \hat{P} = \Sigma^\infty$.*

Corollary 2 follows immediately from Theorem 1.

4 Alternative Equivalence Relations

The above results apply only to the case of effective $_=$ enforcement. This corresponds to the enforcement power of a monitor that sometimes delays the occurrence of actions in its target, or abort its execution, but does not add additional actions or permanently suppress part of the execution sequence, allowing the remainder of the execution to continue. Syntactic equality (and subclasses of that relation) is the only equivalence relation that has been extensively studied in the literature. This naturally does not exhaust the enforcement capabilities of monitors. In this section, we explore alternative equivalence relations (that characterize alternative enforcement mechanisms) and determine the set of enforceable properties for each.

4.1 Subword Equivalence and Insertion Enforcement

The first alternative equivalence relation that we examine is subword equivalence, noted \cong_{\triangleleft} . This corresponds to the enforcement power of a monitor which can *add* any action into the execution, but cannot prevent any action from occurring if the target program requests it. For example, the monitor can enforce a property stating that any opened file is eventually closed by adding the missing close file action before the end of the program's execution. This model is interesting to understand the capabilities of certain types of inline monitors, that are injected into the program in the form of guards or run in parallel with their target, such as those based upon the aspect-oriented programming paradigm [8, 19, 25].

Let $\sigma, \sigma', \tau \in \Sigma^*$, we write $\sigma_\tau \cong_{\triangleleft} \sigma' \Leftrightarrow (\sigma \triangleleft_\Sigma \tau \wedge \sigma' \triangleleft_\Sigma \tau)$, and designate by enforceable $_{\cong_{\triangleleft}}$ the set of properties that are enforceable when $\tau \cong_{\triangleleft} \sigma'$ is the equivalence relation and τ is the original input. This is a very permissive equivalence relation, which in effect allows the monitor to insert any action or actions into the execution, and consider the transformed execution equivalent to the original execution (τ) if all actions performed by the target program are still

present. While this equivalence relation may be too permissive to be realistic, it does allow us to deduce an upper bound to the set of policies enforceable under this paradigm.

We begin by considering the cases that occur when the monitor has only limited control over the action set. When the monitor can only suppress actions or abort the execution, the set of enforceable $_{\cong_{\triangleleft}}$ properties coincides with that of safety properties, since the monitor is unable to take advantage of the permissiveness of the equivalence relation.

Theorem 6. *Let $\mathcal{L} = \langle \emptyset, \emptyset, \Sigma, \emptyset \rangle$. $\mathcal{L}^{\Sigma^\infty}$ -enforceable $_{\cong_{\triangleleft}}$ is the set of safety properties.*

Another interesting case occurs when the monitor cannot delay the occurrence of the actions present in the execution sequence, but can only react to them by inserting additional actions afterwards. In other words, any action output by the target program must be immediately accepted, but can be followed by other actions inserted by the monitor. An intuitive lower bound to the set of enforceable properties in this context is the intersection of renewal properties and liveness properties. That is, properties for which any invalid sequence can be corrected into a valid sequence with a finite number of corrective steps. Additionally, some safety properties and some persistence properties are also enforceable. For example, the property \hat{P}_{-aa} imposing that no two ‘a’ actions occur consecutively is a safety property since an invalid sequence cannot be corrected by the insertion of any subsequent actions. However, the policy is $\mathcal{L}^{\Sigma^\infty}$ -enforceable $_{\cong_{\triangleleft}}$ since the monitor can insert any action other than ‘a’ after the occurrence of each ‘a’ action to ensure the respect of the security property. What characterizes properties outside the intersection of renewal and liveness that are nonetheless $\mathcal{L}^{\Sigma^\infty}$ -enforceable $_{\cong_{\triangleleft}}$ is that there exists a property \hat{P}' in renewal \cap liveness such that the property of interest includes \hat{P}' . In the case of the example above, the property “any a action is immediately followed by some action different from a (or the end of sequence token)” is a renewal property included in \hat{P}_{-aa} . The monitor can enforce the property \hat{P}_{-aa} by enforcing \hat{P}' .

Theorem 7. *Let $\hat{P} \subseteq \mathcal{P}(\Sigma^\infty)$ and let $\mathcal{L} = \langle \emptyset, \Sigma, \emptyset, \emptyset \rangle$. If the monitor cannot delay the occurrence of actions performed by the target program, then $\hat{P} \in \mathcal{L}^{\Sigma^\infty}$ -enforceable $_{\cong_{\triangleleft}} \Leftrightarrow \exists \hat{P}' : \hat{P}' \subseteq \hat{P}$ and \hat{P}' is infinite renewal \cap liveness.*

Allowing the monitor to insert a finite number of actions before or after an action taken by the target program increases the set of enforceable properties further. Consider for example the property \hat{P}_{os} stating that a *write* action only be performed on a previously opened file. The property is a safety property, and falls outside the set of $\hat{P} \in \mathcal{L}^{\Sigma^\infty}$ -enforceable $_{\cong_{\triangleleft}}$ properties defined in Theorem 7 if the number of files is infinite. The property can be $\mathcal{L}^{\Sigma^\infty}$ -enforced $_{\cong_{\triangleleft}}$ by a monitor that inserts the corresponding *open* file action anytime the target program attempts to write to a file that has not yet been opened. More generally, if all actions are in the set \mathcal{C} , then a property \hat{P} is enforceable iff there exists a property $\hat{P}' \subseteq \hat{P}$, s.t. \hat{P}' is in the intersection of Renewal and Liveness or if it is a

safety property and for any action $a \in \Sigma$, and for any sequence σ in \hat{P}' , there is a sequence τ in the residual of \hat{P}' with regard to σ s.t. $a \in \text{acts}(\sigma)$. In other words, if for any sequence in the subproperty \hat{P}' and for any possible action a , there is a continuation τ s.t. τ contains a . The monitor can $\mathcal{L}^{\Sigma^\infty}$ -enforce $_{\cong_q}$ such a property by appending a sequence from the residual containing any action requested by the target program.

Theorem 8. *Let $\mathcal{L} = \langle \emptyset, \Sigma, \emptyset, \emptyset \rangle$. $\mathcal{L}^{\Sigma^\infty}$ -enforceable $_{\cong_q}$ iff there exists a property $\hat{P}' \subseteq \hat{P} \subseteq \text{Liveness} \cap \text{Renewal} \cup \text{Safety} : \forall \sigma \in \hat{P}' : \bigcup_{\sigma' \in \text{res}_{\hat{P}'}(\sigma)} \text{acts}(\sigma') = \Sigma$.*

The upper bound naturally occurs when all actions are controllable. We found it harder to give a specific upper bound to the set of enforceable $_{\cong_q}$ properties. Indeed, this set seems to include almost every properties, with the exception of a number of hard to define special cases. Observe that because the subword relation is reflexive, a monitor that enforce $_{=}$ the property also enforce $_{\cong_q}$ it. This equivalence relation is so permissive that only a few very particular cases seem unenforceable.

Observe that while this result indicates that the security properties that are $\mathcal{L}^{\Sigma^\infty}$ -enforceable $_{\cong_q}$ are largely the same as those that are $\mathcal{L}^{\Sigma^\infty}$ -enforceable $_{=}$, the manner of enforcement is quite different. $\mathcal{L}^{\Sigma^\infty}$ -enforce $_{=}$ guarantees that any valid sequence is output *as is*, without any modification by the monitor. For invalid sequences, $\mathcal{L}^{\Sigma^\infty}$ -enforceable $_{=}$ ensures that the longest valid prefix is always output [5]. $\mathcal{L}^{\Sigma^\infty}$ -enforceable $_{\cong_q}$ does not provide these guarantees. Instead, as seen above, for non-safety properties $\mathcal{L}^{\Sigma^\infty}$ -enforceable $_{\cong_q}$ can ensure that any action present in the original program is eventually output, whether the execution sequence is valid or not. $\mathcal{L}^{\Sigma^\infty}$ -enforceable $_{\cong_q}$ also evidently has a much reduced memory overhead, since this enforcement paradigm does not impose on the monitor that it keep in memory an indefinitely long segment of the execution trace, as $\mathcal{L}^{\Sigma^\infty}$ -enforceable $_{=}$ does. Since memory constraints were showed in [13] to significantly affect the set of enforceable properties it is certain that once such constraints are taken into account, some properties will be found to be $\mathcal{L}^{\Sigma^\infty}$ -enforceable $_{\cong_q}$ but not $\mathcal{L}^{\Sigma^\infty}$ -enforceable $_{=}$.

4.2 Suppression Enforcement

Another interesting enforcement paradigm is suppression enforcement [3], which occurs when the monitor can remove potentially malicious actions performed by the target program but cannot add any action to the execution. In that case, the monitor's enforcement is bounded by the reverse subword equivalence, meaning that the monitor's output is a subword of the original sequence. This enforcement paradigm is similar to the one described in [24], where the monitor is interposed between the target program and the system. Any action requested by the target program is intercepted by the monitor which must accept or reject it, and allows us to pose an upper bound to this enforcement paradigm.

Let $\sigma, \sigma', \tau \in \Sigma^*$, we write $\sigma_\tau \cong_b \sigma' \Leftrightarrow (\tau \triangleleft_\Sigma \sigma \wedge \tau \triangleleft_\Sigma \sigma')$. As was the case in Sect. 4.1, this is a very permissive equivalence relation, which characterizes the

behaviour of a monitor that can potentially *suppress* any action or actions from the execution. We write \cong_{\triangleright} for this equivalence relation where τ is the original input.

We begin by determining an upper bound to the set, which occurs when every action is suppressible. In that case, every reasonable property is enforceable, simply by always outputting the empty sequence (or possibly the longest valid prefix). While this may not be a particularly useful enforcement, it does however serve as a useful upper bound to begin reflecting about the capabilities of different types of monitors. It also argues for a stronger notion of transparency, as discussed in [17], which allows us to reason about the capabilities of monitor in a context that is more similar to that of a real-life monitor. Such monitors would normally be bounded with respect to the alterations that they are allowed to performed on valid and invalid executions alike.

Theorem 9. *If $\mathcal{L} = \langle \emptyset, \emptyset, \Sigma, \emptyset \rangle$ then $\mathcal{L}^{\Sigma^\infty}$ -enforceable $_{\cong_{\triangleright}} = \mathcal{P}(\Sigma^\infty)$.*

Corollary 3. *Let $\mathcal{L} = \langle \emptyset, \emptyset, \Sigma, \emptyset \rangle$. $\forall a \in \mathcal{L}_{\mathcal{D}} : \mathcal{L}^{\Sigma^\infty}$ -enforceable $_{\cong_{\triangleright}} = \mathcal{L}^{\Sigma^\infty}_{\mathcal{D} \xrightarrow{a} \mathcal{C}}$ -enforceable $_{\cong_{\triangleright}}$.*

Corollary 4. *Let $\mathcal{L} = \langle \emptyset, \emptyset, \Sigma, \emptyset \rangle$. $\forall \mathcal{S} \subseteq \Sigma^\infty : \mathcal{L}^{\Sigma^\infty}$ -enforceable $_{\cong_{\triangleright}} = \mathcal{L}^{\mathcal{S}}$ -enforceable $_{\cong_{\triangleright}}$.*

A more interesting characterization occurs when we consider that some subset of Σ is unsuppressible. These may be actions that the monitor lacks the ability to suppress (such as internal system computations) or actions that cannot be deleted without affecting the functionality of the target program. When that is the case, a property is enforceable iff every invalid sequence has a valid prefix ending on an action in \mathcal{D} . In that case, the property of interest includes a safety property that can be enforced by truncation.

Theorem 10. *If $\mathcal{L} = \langle \mathcal{O}, \emptyset, \mathcal{D}, \emptyset \rangle$ then $\forall \sigma \in \Sigma^\infty : \hat{P}(\sigma) \in \mathcal{L}^{\Sigma^\infty}$ -enforceable $_{\cong_{\triangleright}} \Leftrightarrow \exists \hat{P}' \subseteq \hat{P} : \hat{P}' \in \mathcal{L}$ -safety.*

5 Conclusion

In this paper, we reexamined the delimitation of enforceable properties by monitors, and proposed a finer characterization that distinguishes between actions that are only observable, actions which the monitor can delete but not insert into the execution, actions which the monitor can insert in the execution but not suppress if they are already present and completely controllable actions. Our study is a generalization of previous work on the same topic, and provides a finer characterization of the set of security properties that are enforceable by monitors in different contexts.

Additionally, we explored the set of properties that are enforceable by the monitor is given broader latitude to transform valid sequences, rather than be bounded to return a syntactically identical execution sequence if the original execution is valid. We argue that our results point to the need for an alternative definition of enforcement.

Part of the reason the sets of $\mathcal{L}^{\Sigma^\infty}$ -enforceable $_{\cong_b}$ and $\mathcal{L}^{\Sigma^\infty}$ -enforceable $_{\cong_q}$ properties are so large is that, in the definition of effective $_{\cong}$ enforcement, the transparency requirement is so weak. This leads to monitors with unusually broad licence to alter invalid sequences in order to correct them. For monitors with broad capabilities to add and remove actions from the execution sequence, the desired behaviour of real-life security policy enforcement mechanism would be more accurately characterized by a more constraining definition of *enforcement*. For example, a practical suppression monitor should be bounded to remove from an invalid execution sequence only those actions that violate the security policy. Any valid behaviour present in an otherwise invalid sequence should be preserved.

In the future, we would like to consider that the cost of inserting or deleting an action may not be the same for all actions. This would allow us to contrast different enforcement strategies for the same security property. A lattice-based framework is well-suited to model such a restriction.

References

1. Alpern, B., Schneider, F.: Defining liveness. *Inf. Process. Lett.* **21**(4), 181–185 (1985)
2. Basin, D., Jugé, V., Klaedtke, F., Zalinescu, E.: Enforceable security policies revisited. *ACM Trans. Inf. Syst. Secur.* **16**(1), 3 (2013)
3. Bauer, L., Ligatti, J., Walker, D.: More enforceable security policies. In: *Foundations of Computer Security*, Copenhagen, Denmark (2002)
4. Beauquier, D., Pin, J.-E.: Languages and scanners. *Theoret. Comput. Sci.* **84**(1), 3–21 (1991)
5. Bielova, N., Massacci, F.: Do you really mean what you actually enforced? - edit automata revisited. *Int. J. Inf. Secur.* **10**(4), 239–254 (2011)
6. Bielova, N., Massacci, F.: Predictability of enforcement. In: Erlingsson, U., Wieringa, R., Zannone, N. (eds.) *ESSoS 2011*. LNCS, vol. 6542, pp. 73–86. Springer, Heidelberg (2011)
7. Bielova, N., Massacci, F., Micheletti, A.: Towards practical enforcement theories. In: Jøsang, A., Maseng, T., Knapskog, S.J. (eds.) *NordSec 2009*. LNCS, vol. 5838, pp. 239–254. Springer, Heidelberg (2009)
8. Bodden, E., Lam, P., Hendren, L.J.: Partially evaluating finite-state runtime monitors ahead of time. *ACM Trans. Program. Lang. Syst.* **34**(2), 7 (2012)
9. Chabot, H., Khoury, R., Tawbi, N.: Extending the enforcement power of truncation monitors using static analysis. *Comput. & Secur.* **30**(4), 194–207 (2011)
10. Chang, E., Manna, Z., Pnueli, A.: The safety-progress classification. In: Bauer, F., Brauer, W., Schwichtenberg, H. (eds.) *Logic and Algebra of Specifications*. NATO ASI Series, vol. 94, pp. 143–202. Springer, Heidelberg (1991)
11. Falcone, Y., Fernandez, J.-C., Mounier, L.: Synthesizing enforcement monitors wrt. the safety-progress classification of properties. In: Sekar, R., Pujari, A.K. (eds.) *ICISS 2008*. LNCS, vol. 5352, pp. 41–55. Springer, Heidelberg (2008)
12. Falcone, Y., Fernandez, J.-C., Mounier, L.: Runtime verification of safety-progress properties. In: Bensalem, S., Peled, D.A. (eds.) *RV 2009*. LNCS, vol. 5779, pp. 40–59. Springer, Heidelberg (2009)

13. Fong, P.: Access control by tracking shallow execution history. In: Proceedings of the 2004 IEEE Symposium on Security and Privacy, California, USA, May, Oakland (2004)
14. Hamlen, K.W., Morrisett, G., Schneider, F.B.: Computability classes for enforcement mechanisms. *ACM Trans. Program. Lang. Syst.* **28**, 175–205 (2006)
15. Khoury, R., Hallé, S.: Runtime enforcement with partial control. Technical report, Université du Québec à Chicoutimi (2015)
16. Khoury, R., Tawbi, N.: Using equivalence relations for corrective enforcement of security policies. In: Kotenko, I., Skormin, V. (eds.) *MMM-ACNS 2010*. LNCS, vol. 6258, pp. 139–154. Springer, Heidelberg (2010)
17. Khoury, R., Tawbi, N.: Corrective enforcement: A new paradigm of security policy enforcement by monitors. *ACM Trans. Inf. Syst. Secur.* **15**(2), 10:1–10:27 (2012)
18. Khoury, R., Tawbi, N.: Which security policies are enforceable by runtime monitors? a survey. *Comput. Sci. Rev.* **6**(1), 27–45 (2012)
19. Kiczales, G., Hilsdale, E.: Aspect-oriented programming. *SIGSOFT Softw. Eng. Notes* **26**(5), 313 (2001)
20. Kim, M., Kannan, S., Lee, I., Sokolsky, O., Viswanathan, M.: Computational analysis of run-time monitoring - fundamentals of java-mac. *Electr. Notes Theor. Comput. Sci.* **70**(4), 80–94 (2002)
21. Lammport, L.: Proving the correctness of multiprocess programs. *IEEE Trans. Softw. Eng.* **3**(2), 125–143 (1977)
22. Ligatti, J., Bauer, L., Walker, D.: Edit automata: Enforcement mechanisms for run-time security policies. *Int. J. Inf. Secur.* **4**(1–2), 2–16 (2004)
23. Ligatti, J., Bauer, L., Walker, D.: Run-time enforcement of nonsafety policies. *ACM Trans. Inf. Syst. Secur.* **12**(3), 1–41 (2009)
24. Ligatti, J., Reddy, S.: A theory of runtime enforcement, with results. In: Gritzalis, D., Preneel, B., Theoharidou, M. (eds.) *ESORICS 2010*. LNCS, vol. 6345, pp. 87–100. Springer, Heidelberg (2010)
25. Meredith, P., Roşu, G.: Runtime verification with the RV system. In: Barringer, H., et al. (eds.) *RV 2010*. LNCS, vol. 6418, pp. 136–152. Springer, Heidelberg (2010)
26. Schneider, F.: Enforceable security policies. *Inf. Syst. Secur.* **3**(1), 30–50 (2000)
27. Talhi, C., Tawbi, N., Debbabi, M.: Execution monitoring enforcement for limited-memory systems. In: Proceedings of the PST 2006 Conference (Privacy, Security, Trust), October 2006
28. Talhi, C., Tawbi, N., Debbabi, M.: Execution monitoring enforcement under memory-limitation constraints. In: Proceedings of FCS-ARSPA 2006 (Joint Workshop on Foundations of Computer Security and Automated Reasoning for Security Protocol Analysis) associated with FLOC 2006 (Federated Logic Conference), August 2006
29. Talhi, C., Tawbi, N., Debbabi, M.: Execution monitoring enforcement under memory-limitations constraints. *Inf. Comput.* **206**(1), 158–184 (2008)