



HAL
open science

Runtime Enforcement Using Büchi Games

Matthieu Renard, Antoine Rollet, Yliès Falcone

► **To cite this version:**

Matthieu Renard, Antoine Rollet, Yliès Falcone. Runtime Enforcement Using Büchi Games. Model Checking Software - 24rd International Symposium, SPIN 2017, Co-located with ISSSTA 2017, Jul 2017, Santa Barbara, United States. hal-01509827v3

HAL Id: hal-01509827

<https://hal.science/hal-01509827v3>

Submitted on 17 May 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Runtime Enforcement Using Büchi Games

Matthieu Renard
LaBRI
Bordeaux INP
University of Bordeaux
Bordeaux, France
matthieu.renard@labri.fr

Antoine Rollet
LaBRI
Bordeaux INP
University of Bordeaux
Bordeaux, France
antoine.rollet@labri.fr

Yliès Falcone
Univ. Grenoble Alpes, Inria
Laboratoire d'Informatique de
Grenoble
F-38000, Grenoble, France
yliès.falcone@univ-grenoble-alpes.fr

ABSTRACT

We leverage Büchi games for the runtime enforcement of regular properties with uncontrollable events. Runtime enforcement consists in modifying the execution of a running system to have it satisfy a given regular property, modelled by an automaton. We revisit runtime enforcement with uncontrollable events and propose a framework where we model the runtime enforcement problem as a Büchi game and synthesise sound, compliant, and optimal enforcement mechanisms as strategies. We present algorithms and a tool implementing enforcement mechanisms. We reduce the complexity of the computations performed by enforcement mechanisms at runtime by pre-computing the decisions of enforcement mechanisms ahead of time.

KEYWORDS

Runtime Verification, Runtime Enforcement, Games, Automata

1 INTRODUCTION

Runtime Verification (RV) consists in checking if the execution of a running system satisfies some given specification. Unlike static verification, RV studies a real execution of a system, possibly after deployment. This paper deals with *runtime enforcement*, an extension of runtime verification where executions are corrected when they violate the property [1, 10, 11, 15]; see [7] for a tutorial on runtime enforcement. The considered properties are regular properties, that are represented by deterministic and complete automata. All such properties are monitorable since they are defined with a semantics over finite words [8]. An enforcement mechanism modifies an execution: it takes an execution as input and outputs a possibly-different execution. Enforcement mechanisms may operate *online*, meaning that they modify the executions of a system while it is running, or *offline*, by reading a log of system events. While working online, enforcement mechanisms can add a time overhead due to their need to compute a correct output. We distinguish two categories of events: *controllable* events that can be modified by an enforcement mechanism, and *uncontrollable* events that can only be observed by the enforcement mechanism. Enforcement mechanisms should be *sound* and *compliant*, meaning that the output should satisfy the specification when it is possible, and that the output should be as close to the input as possible, respectively. The general scheme is given in Fig. 1.

Motivations. In this paper, we improve the modelling of the enforcement mechanisms proposed in [13, 14], as well as the computation of their output. Such mechanisms should impact the system as little as possible, thus reducing the time spent by enforcement mechanisms to compute their behaviour allows us to use them in a more realistic way. For example, in interactive systems, where the system interacts with a human user, if an event takes too long to be output, the user may think that the system failed. It could also be useful for embedded systems, where computing power may be reduced. Computing the behaviour of the enforcement mechanism ahead of the execution and storing it ensures that the computation does not depend on the size of the automaton, thus allowing the time spent on online computations by enforcement mechanisms to be reduced and more predictable. Indeed, not exploring the whole execution tree for all possible outputs at runtime, as a naive approach would do, allows us to have computation times that vary less (with a naive approach, the computation time can become very important with an increasing number of stored controllable events).

Challenges. Storing the behaviour of enforcement mechanisms to improve their online computation time induces some changes compared to previous work (as [13, 14]). The main difficulty resides in the fact that the number of states of the enforcement mechanism is infinite. Indeed, the mechanism has the possibility to store controllable events that it may choose to release or not. The number of events that can be stored at the same time is not bounded, thus the number of states of the enforcement mechanism is infinite. Therefore, computing and storing its behaviour for all possible input traces entails defining appropriate abstractions.

Contributions. A first approach of enforcement with uncontrollable events has been presented in [14] providing sound and compliant enforcement mechanisms. An optimal version of this approach may be found in [13]. In this paper, we propose to extend this work by computing the behaviour of the enforcement mechanism using Büchi games. Using Büchi games allows us to compute the behaviour of the enforcement mechanism before the execution, thus allowing to trade time complexity with space complexity. When




Figure 1: Schematic description of an enforcement mechanism E , modifying the execution σ of the system S to $E(\sigma)$, so that it satisfies the property φ .

using an *online* enforcement mechanism (i.e. an enforcement mechanism on a running system), it allows us to produce the output faster in the worst case than in [14] and [13] where the behaviour of the enforcement mechanism was computed every time an event was received. Leveraging games, when an event is received, the enforcement mechanism only follows a path in a graph, and the destination vertex is sufficient to indicate if a stored event should be output or not. Moreover, the generated graph can be visualised to understand the behaviour of the enforcement mechanism, and it also shows clearly when an enforcement mechanism can effectively ensure soundness. We redefine soundness, compliance and optimality using a set-theoretic view of the system, thus providing a global vision of the system based on inputs and outputs at any instant. We give the algorithms (and their complexity analysis) implementing the behaviour of enforcement mechanisms. We finally present a tool implementing the proposed approach.

2 PRELIMINARIES AND NOTATION

General notions. An *alphabet* is a finite set of symbols. A *word* over an alphabet Σ is a sequence over Σ . The set of finite words over Σ is denoted Σ^* . The *length* of a finite word w is noted $|w|$, and the *empty word* is noted ϵ . Σ^+ stands for $\Sigma^* \setminus \{\epsilon\}$. A *language* over Σ is any subset $L \subseteq \Sigma^*$. The concatenation of two words w and w' is noted $w.w'$ (or ww' when clear from the context). A word w' is a *prefix* of a word w , noted $w' \preceq w$, if there exists a word w'' such that $w = w'.w''$. Word w'' is called the *residual* of w after reading the prefix w' , noted $w'' = w'^{-1}.w$. Note that $w'.w'' = w'.w'^{-1}.w = w$. These definitions are extended to languages in the natural way. A language $L \subseteq \Sigma^*$ is *extension-closed* if for any words $w \in L$ and $w' \in \Sigma^*$, $w.w' \in L$. Given a word w and an integer i such that $1 \leq i \leq |w|$, we note $w(i)$ the i -th element of w . Given a tuple $e = (e_1, e_2, \dots, e_n)$ of size n , for an integer i such that $1 \leq i \leq n$, we note Π_i the projection on the i -th coordinate, i.e. $\Pi_i(e) = e_i$. The tuple (e_1, e_2, \dots, e_n) is sometimes noted $\langle e_1, e_2, \dots, e_n \rangle$ in order to help reading. It can be used, for example, if a tuple contains a tuple. Given a word $w \in \Sigma^*$ and $\Sigma' \subseteq \Sigma$, we define the *restriction* of w to Σ' , noted $w|_{\Sigma'}$, as the word $w' \in \Sigma'^*$ whose letters are the letters of w belonging to Σ' in the same order. Formally, $\epsilon|_{\Sigma'} = \epsilon$ and $\forall \sigma \in \Sigma^*, \forall a \in \Sigma, (w.a)|_{\Sigma'} = w|_{\Sigma'}.a$ if $a \in \Sigma'$, and $(w.a)|_{\Sigma'} = w|_{\Sigma'}$ otherwise. We also note $=_{\Sigma'}$ the equality of the restrictions of two words to Σ' : for σ and σ' in Σ^* , $\sigma =_{\Sigma'} \sigma'$ if $\sigma|_{\Sigma'} = \sigma'|_{\Sigma'}$. We define in the same way $\preceq_{\Sigma'}$: $\sigma \preceq_{\Sigma'} \sigma'$ if $\sigma|_{\Sigma'} \preceq \sigma'|_{\Sigma'}$.

Automata. An *automaton* is a tuple $\langle Q, q_0, \Sigma, \rightarrow, F \rangle$, where Q is the set of *states*, $q_0 \in Q$ is the initial state, Σ is the alphabet, $\rightarrow \subseteq Q \times \Sigma \times Q$ is the transition relation and $F \subseteq Q$ is the set of accepting states. Whenever there exists $(q, a, q') \in \rightarrow$, we note it $q \xrightarrow{a} q'$. Relation \rightarrow is extended to its reflexive and transitive closure in the usual way. Moreover, for any $q \in Q$, $q \xrightarrow{\epsilon} q$ always holds. An automaton $\mathcal{A} = \langle Q, q_0, \Sigma, \rightarrow, F \rangle$ is *deterministic* if $\forall q \in Q, \forall a \in \Sigma, (q \xrightarrow{a} q' \wedge q \xrightarrow{a} q'') \implies q' = q''$. \mathcal{A} is *complete* if $\forall q \in Q, \forall a \in \Sigma, \exists q' \in Q, q \xrightarrow{a} q'$. A word w is *accepted* by \mathcal{A} if there exists $q \in F$ such that $q_0 \xrightarrow{w} q$. The language (i.e. set of all words) accepted by \mathcal{A} is noted $\mathcal{L}(\mathcal{A})$. A *property* is a language accepted over an alphabet Σ . A regular property is a language accepted by an automaton. In the sequel, we assume that a property φ is




Figure 2: Property φ_{ex} modelling writes on a shared storage device

represented by a deterministic and complete automaton \mathcal{A}_φ . Given a complete and deterministic automaton $\mathcal{A} = \langle Q, q_0, \Sigma, \rightarrow, F \rangle$ and a word $\sigma \in \Sigma^*$, for $q \in Q$, we note q after σ the only state such that $q \xrightarrow{\sigma} (q \text{ after } \sigma)$. The completeness of \mathcal{A} ensures that q after σ exists, and its determinism ensures that it is unique. We also note $\text{Reach}(\sigma) = q_0 \text{ after } \sigma$. We extend these definitions to languages: if L is a language, q after $L = \bigcup_{\sigma \in L} q$ after σ and $\text{Reach}(L) = q_0$ after L .

Graphs and Büchi games. A *graph* is a couple $\langle V, E \rangle$ such that V is a set of elements called *vertices*, $E \subseteq V \times V$ is a relation defining *edges* between the vertices. Given a graph $G = \langle V, E \rangle$ and a partition of V into two subsets V_0 and V_1 , it is possible to play a two-player game in the *arena* $A = (V_0, V_1, E)$. A *play* over A is a path in G , i.e. a sequence of vertices such that there exists an edge in G between any two consecutive vertices in the sequence. A *strategy* for player P_0 is a mapping $\sigma : V^*V_0 \rightarrow V$ such that for all $\pi \in V^*$, for all $v_0 \in V_0$, $(v_0, \sigma(\pi.v_0)) \in E$, i.e. the strategy gives a vertex that can be reached from v_0 . Note that V_0 is thus the set of vertices from which P_0 can play, whereas the other player, P_1 , plays from the vertices in V_1 . Strategies for P_1 are defined in a similar way, replacing V_0 by V_1 . A play $\pi = v_0, v_1, \dots$ is *consistent* with the strategy σ if for any $v_i \in V_0, v_{i+1} = \sigma(v_0.v_1 \dots v_i)$, meaning that the strategy was followed for any vertex in V_0 . The goal of a game can be, for example, to reach a state in a given subset of V (reachability game), or to ensure that a given subset of V is visited an infinite number of times (Büchi games). Thus, given a subset $F_G \subseteq V$ of vertices, the Büchi game (A, F_G) for P_0 consists in finding a *winning strategy* σ such that all plays π over A consistent with σ visit an infinite number of times the set F_G (i.e. if π is consistent with σ , $\pi \in (V^*F_G)^\omega$). It is known that it is possible to compute the set W_0 of winning vertices for P_0 (i.e. the set of vertices from where there exists a winning strategy for P_0), and the associated winning strategy from all these vertices. From all the other vertices (in $V \setminus W_0$), there exists a winning strategy for P_1 , i.e. $W_1 = V \setminus W_0$, thus P_0 can not win the game if P_1 plays perfectly from one of these vertices.

3 ENFORCEMENT MONITORING OF PROPERTIES USING BÜCHI GAMES

This paper revisits and extends the approach described in [14] by writing the definitions in a set-theoretic view instead of a functional

way and proposing a new synthesis technique of *Enforcement Mechanisms* (EM) using Büchi games.

In this section, φ is a regular property defined by a complete and deterministic automaton $\mathcal{A}_\varphi = \langle Q, q_0, \Sigma, \rightarrow, F \rangle$. Recall that the general scheme of an EM is given in Fig. 1, where S represents the running system, σ its execution, E the enforcement mechanism, φ the property to enforce, and $E(\sigma)$ the output of the enforcement mechanism, which should satisfy φ .

We consider uncontrollable events¹ in the set $\Sigma_u \subseteq \Sigma$. These events cannot be modified by an EM, so they must be output by the EM whenever they are received. Let us note $\Sigma_c = \Sigma \setminus \Sigma_u$ the set of controllable events, which can be modified by the EM. An EM can decide to buffer them to delay their emission, but it cannot suppress them (nevertheless, it can delay them endlessly, keeping their order unchanged).² Thus, an EM may interleave controllable and uncontrollable events.

3.1 Enforcement Functions and their Requirements

We consider an alphabet of actions Σ . We consider functions as sets: a *function* from a set A to a set B is a set $f \subseteq A \times B$ such that for any element a in A , there is a unique b in B such that $(a, b) \in f$. We note $\mathcal{F}(A, B)$ the set of all functions from A to B . An enforcement function is a description of the input/output behaviour of an EM. It is a function from Σ^* to Σ^* , increasing on Σ^* (with respect to \preceq):

Definition 3.1 (Enforcement function). A function $f \in \mathcal{F}(\Sigma^*, \Sigma^*)$ is an *enforcement function* if $\forall i_1 \in \Sigma^*, \forall i_2 \in \Sigma^*, (i_1 \preceq i_2 \wedge (i_1, o_1) \in f \wedge (i_2, o_2) \in f) \implies o_1 \preceq o_2$. We note \mathcal{F}_{enf} the set of all enforcement functions.

An enforcement function is a function that modifies an execution, and that cannot remove events it has already output.

In the sequel, we define the requirements on an EM and express them on enforcement functions. As stated previously, the usual purpose of an EM is to ensure that the executions of a running system satisfy a property, thus its enforcement function has to be *sound*, meaning that its output always satisfies the property:

Definition 3.2 (Soundness). An enforcement function $E \in \mathcal{F}_{\text{enf}}$ is *sound* with respect to φ in an extension-closed set $S \subseteq \Sigma^*$ if $\forall i \in S, (i, o) \in E \implies o \models \varphi$. We note $\mathcal{F}_{\text{snd}}(S)$ the set of all enforcement functions that are sound in S .

The reception of uncontrollable events could lead to the property not being satisfied by the output of the enforcement mechanism. Moreover, some uncontrollable sequences could lead to a state of the property that would be a non-accepting sink state. Thus, the enforcement mechanism would not be able to make the property satisfied. Consequently, in Definition 3.2, soundness is not defined for all words in Σ^* , but in a subset S , since the property could not be enforceable from the initial state. In practice, S needs to be extension-closed to ensure that once the enforcement mechanism becomes sound, its output will always satisfy the property afterwards.

¹This notion of uncontrollable event should not be confused with the notion of uncontrollable transition used in some game theory (e.g. [4])

²This choice appeared to us as the most realistic one. Extending the notions presented in this section in order to handle enforcement mechanisms with suppression is rather simple.

The usual notion of *transparency* in enforcement monitoring (cf. [11, 15]) states that the output of an enforcement function is the longest prefix of the input satisfying the property, implying that correct executions are left unchanged. However, because of uncontrollable events, events may be released in a different order from the one they are received. Therefore, transparency can not be ensured, and we define the weaker notion of *compliance*.

Definition 3.3 (Compliance). $E \in \mathcal{F}_{\text{enf}}$ is *compliant* with respect to Σ_u and Σ_c , noted $\text{compliant}(E, \Sigma_u, \Sigma_c)$, if $\forall i \in \Sigma^*, (i, o) \in E \implies (o \preceq_{\Sigma_c} i \wedge o =_{\Sigma_u} i \wedge \forall u \in \Sigma_u, ((i.u, o') \in E \implies o.u \preceq o'))$. We note $\mathcal{F}_{\text{cpl}}(\Sigma_u, \Sigma_c)$ the set of all enforcement functions that are compliant with respect to Σ_u and Σ_c .

Intuitively, compliance states that the EM does not change the order of the controllable events and emits uncontrollable events simultaneously with their reception, possibly followed by stored controllable events. When clear from the context, the partition is not mentioned: E is said to be compliant, we note it $\text{compliant}(E)$, and the set of all compliant functions is then denoted \mathcal{F}_{cpl} .

We say that a property φ is *enforceable* whenever there exists a compliant function that is sound with respect to φ .

In addition, an enforcement mechanism should be optimal in the sense that its output sequences should be maximal while preserving soundness and compliance. We define the optimality of sound and compliant enforcement functions as follows:

Definition 3.4 (Optimality). An enforcement function $E \in \mathcal{F}_{\text{snd}}(S) \cap \mathcal{F}_{\text{cpl}}(\Sigma_u, \Sigma_c)$ is *optimal* in S if:

$$\forall E' \in \mathcal{F}_{\text{snd}}(S) \cap \mathcal{F}_{\text{cpl}}(\Sigma_u, \Sigma_c), \forall i \in S, \forall a \in \Sigma, ((i, o) \in E \cap E' \wedge (i.a, o') \in E \wedge (i.a, p') \in E') \implies p' \preceq o'.$$

Intuitively, optimality states that outputting a longer word than an optimal enforcement function breaks soundness or compliance. Since it is not always possible to satisfy the property from the beginning, this condition is restrained to an extension-closed subset of Σ^* , as in the definition of soundness (see Definition 3.2).

Example 3.5. We consider a simple shared storage device. After Authentication, a user can write a value only if the storage is unlocked. (Un)locking the device is decided by another entity, meaning that it is not controllable by the user. Property φ_{ex} (see Fig. 2) formalises the above requirement. φ_{ex} is not enforceable if the uncontrollable alphabet is $\{\text{LockOn}, \text{LockOff}, \text{Auth}\}$ ³ since reading the word *LockOn* from q_0 leads to q_3 , which is not an accepting state. However, the existence of such a word does not prevent φ_{ex} from being enforced for some other input words. If word *Auth* is read, then state q_1 is reached, and from this state, it is possible to enforce φ_{ex} by emitting *Write* only when in state q_1 .

3.2 Synthesising Enforcement Functions

Example 3.5 shows that some input words cannot be corrected by the EM because of uncontrollable events. Nevertheless, since the received events may lead to a state from which it is possible to ensure that φ will be satisfied (meaning that for any events received as input, the enforcement mechanism can output a sequence that satisfies φ), it is then possible to define a subset of Σ^* in which an enforcement function is sound.

³Uncontrollable events are emphasised in italics.

A compliant enforcement mechanism may store the received controllable events to emit them after having received another event, possibly uncontrollable. To ensure soundness, the enforcement mechanism must know if it is possible to emit some of its stored controllable events in order to reach an accepting state, from which it will be able to reach an accepting state even if some uncontrollable events are received later on. Thus, it should compute the set of words it can emit to reach such an accepting state. This set will be called G , and to define it, we solve a Büchi game over a graph representing the possible actions of an enforcement monitor. Solving a Büchi game over a graph with a set of goal nodes (referred to as Büchi nodes) consists in finding the nodes of the graph from where it is possible to reach a Büchi node infinitely often, no matter what the opponent's strategy is. Note that a winning node of a Büchi game is also a winning node for the reachability game on the same graph, with the set of Büchi nodes used as the set of goal nodes, since the Büchi criterion requires that a node is reachable infinitely often (thus, in particular, once). Besides, Büchi games are more permissive than safety games, where one wants to remain in the given set of goal nodes. Since the definition of soundness requires that the output of the enforcement mechanism always eventually satisfies the property, the natural choice of game is a Büchi game. Solving a Büchi game is made by computing a set of nodes of the graph from which there exists a winning strategy. Then, from any of these winning nodes, the player can always come back to a Büchi state, whatever the strategy of the adversary is. Here, we construct a graph such that the enforcement mechanism is a player, and we compute its winning nodes, with the Büchi nodes representing a valid execution. The vertices of the graph are composed of a state in Q and the stored controllable events of the enforcement mechanism. There exists two of each of these vertices, one that belongs to player P_0 and one that belongs to player P_1 . Player P_0 represents the enforcement mechanism, and P_1 the environment.

Definition 3.6 (Game graph). The game graph \mathcal{G} is defined as $\mathcal{G} = \langle V, E \rangle$, where

- $V = Q \times \Sigma_c^* \times \{0, 1\}$,
- $E_1 = \{ \langle (q, w, 0), (q, w, 1) \rangle \in V \times V \}$,
- $E_2 = \{ \langle (q, c.w, 0), (q \text{ after } c, w, 0) \rangle \in V \times V \mid c \in \Sigma_c \}$,
- $E_3 = \{ \langle (q, w, 1), (q \text{ after } u, w, 0) \rangle \in V \times V \mid u \in \Sigma_u \}$,
- $E_4 = \{ \langle (q, w, 1), (q, w.c, 0) \rangle \in V \times V \mid c \in \Sigma_c \}$,
- $E_5 = \{ \langle (q, w, 1), (q, w, 0) \rangle \in V \times V \}$,
- $E = E_1 \cup E_2 \cup E_3 \cup E_4 \cup E_5$.

A vertex $\langle q, w, l \rangle \in V$ represents the state of the enforcement mechanism: $q \in Q$ is the state of \mathcal{A}_φ that has been reached so far by the output of the enforcement mechanism, $w \in \Sigma_c^*$ is the stored controllable events of the enforcement mechanism, and $l \in \{0, 1\}$ indicates that the vertex belongs to the player P_l . In the definition of E , each set of edges represents an action of the enforcement mechanism or the environment. The enforcement mechanism can only take two decisions: doing nothing, i.e. letting the environment play (set E_1), or emitting the first stored controllable event (set E_2), in which case it continues to play (since the destinations of the edges in E_2 belong to P_0). The sets E_3 and E_4 represent the reception of an uncontrollable and a controllable event, respectively. Receiving an event lets the enforcement mechanism (P_0) play. Since games are infinite, and we only consider finite executions, the environment

can also decide to let the enforcement mechanism play without any new event (set E_5). This allows us to consider finite executions that produce an infinite path in the game by looping on an edge in E_1 and then one in E_5 .

Unfortunately, this graph has an infinite number of vertices, it is thus impossible to compute a winning strategy in a Büchi game played over this graph. To overcome this, the graph is reduced to a graph with a finite number of vertices. To do this, first note that the number of vertices is infinite because the set Σ_c^* is not bounded. Thus, Σ_c^* must be abstracted to a finite set. Since the goal is to reach a state in F , the stored controllable events are used to reach some states in Q . Since Q is finite, having more controllable events than $|Q|$ means that (following the Pumping lemma) there is a loop, i.e. some state in Q is reached twice when emitting all the controllable events. Thus, the enforcement mechanism can emit all the events until it reaches this state for the second time, and then its decision will only depend on the remaining controllable events. Thus, the number of controllable events can be reduced to at most $|Q|$. More precisely, we can reduce Σ_c^* to the set of words that allow to reach a new state (i.e. a state that is not reached by one of its prefixes) from at least one state in Q . Let us call this set Σ_c^n , and define it as follows:

$$\Sigma_c^n = \{ w \in \Sigma_c^* \mid \exists q \in Q, \exists c \in \Sigma_c, \forall w' \preceq w, q \text{ after } w.c \neq q \text{ after } w' \}$$

As explained previously, since Q is finite, Σ_c^n is finite as well. Now, let us redefine \mathcal{G} to an abstraction of the game graph:

Definition 3.7 (Abstracted game graph). $\mathcal{G} = \langle V', E' \rangle$, where $V' = Q \times \Sigma_c^n \times \{0, 1\}$, and E' is the same set as E , but considering vertices in V' instead of V .

\mathcal{G} then has a finite number of vertices. Let us now consider $W_0 \subseteq V$ the set of vertices that are winning for P_0 in the Büchi game over \mathcal{G} , with the set of Büchi (accepting) vertices being $F \times \Sigma_c^n \times \{0, 1\}$.

Example 3.8. The graph in Fig. 3 is computed from property φ_{ex} , with `Write` abbreviated `w` in the second member of the nodes. The Büchi nodes are double circled, and the winning nodes for player 0 (i.e. nodes in W_0) are in blue and rounded rectangles (in our example, all the Büchi nodes are winning). Each edge has a different colour and a different head depending on the set it belongs to. Blue edges (empty triangular head) belong to E_1 , green edges (filled triangular head) belong to E_2 , orange edges (empty diamond head) belong to E_4 , and red edges (filled diamond head) belong to $E_3 \cup E_5$. Each edge is represented only once, even if there are multiple edges in the set (for example, because multiple uncontrollable events lead to the same state from one state). The squared vertex is the initial vertex, and “-” stands for “ ϵ ” (empty buffer).

Since the initial vertex is black (not rounded), this means that it is impossible to ensure that the property will be satisfied from the beginning. The only way to reach a winning state is to follow a red edge from a vertex in $\{q_0\} \times \{\epsilon, w, w.w\} \times \{1\}$, that corresponds to receiving the uncontrollable event `Auth` (since it leads to a state in $\{q_2\} \times \{\epsilon, w, w.w\} \times \{0\}$). Then, `Write` events can only be emitted when in state q_1 . This behaviour is the one expected, since in φ_{ex} , the only way to reach a state in F from q_0 is to follow a path labelled by `Auth`, reaching q_1 . Then, from q_1 , it is possible to emit `Write` events, and if some uncontrollable events are received that lead to

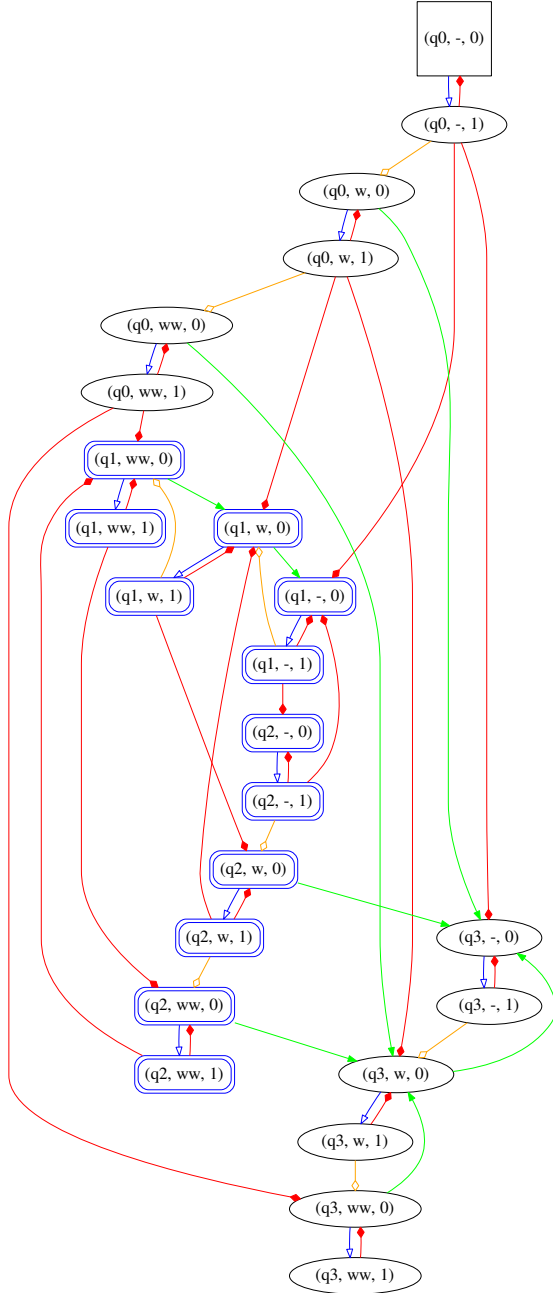


Figure 3: Graph of the game associated to φ_{ex}

q_2 , a *LockOff* event must occur to get back to q_1 and be able to emit another *Write* event.

Now, we can use W_0 to define G , the set of words that can be emitted from a state $q \in Q$ by an enforcement mechanism with a buffer $\sigma \in \Sigma_c^*$.

Definition 3.9 (G). For a state $q \in Q$ and a word of controllable events $\sigma \in \Sigma_c^*$, we define the set $G(q, \sigma)$ as follows:

$$G(q, \sigma) = \{w \in \Sigma_c^* \mid w \preceq \sigma \wedge q \text{ after } w \in F \wedge \langle q \text{ after } w, \max_{\preceq}(\{w' \preceq w^{-1} \cdot \sigma \mid w' \in \Sigma_c^n\}), 1 \rangle \in W_0\}.$$

Intuitively, G is the set of words that can be output by a compliant enforcement mechanism to ensure soundness.

Now, we use G to define the functional behaviour of the enforcement mechanism.

Definition 3.10 (Functions $\text{store}_\varphi, E_\varphi$).⁴ Function $\text{store}_\varphi \in \Sigma^* \times (\Sigma^* \times \Sigma_c^*)$ is defined by induction on its first member as follows:

- $(\epsilon, \langle \epsilon, \epsilon \rangle) \in \text{store}_\varphi$;
- for $\sigma \in \Sigma^*$ and $a \in \Sigma$, let $(\sigma, \langle \sigma_s, \sigma_c \rangle) \in \text{store}_\varphi$, then:

$$\begin{cases} (\sigma.a, \langle \sigma_s.a.\sigma'_s, \sigma'_c \rangle) \in \text{store}_\varphi & \text{if } a \in \Sigma_u \\ (\sigma.a, \langle \sigma_s.\sigma''_s, \sigma''_c \rangle) \in \text{store}_\varphi & \text{if } a \in \Sigma_c \end{cases}, \text{ where:}$$

$$\begin{aligned} \kappa_\varphi(q, w) &= \max_{\preceq}(G(q, w) \cup \{\epsilon\}), \text{ for } q \in Q \text{ and } w \in \Sigma_c^*, \\ \sigma'_s &= \kappa_\varphi(\text{Reach}(\sigma_s.a), \sigma_c), & \sigma'_c &= \sigma_s'^{-1} \cdot \sigma_c, \\ \sigma''_s &= \kappa_\varphi(\text{Reach}(\sigma_s), \sigma_c.a), & \sigma''_c &= \sigma_s''^{-1} \cdot (\sigma_c.a). \end{aligned}$$

The enforcement function $E_\varphi \in \mathcal{F}_{\text{enf}}$ is defined as:

$$E_\varphi = \{(\sigma, \sigma') \mid \exists w \in \Sigma_c^*, (\sigma, \langle \sigma', w \rangle) \in \text{store}_\varphi\}.$$




Figure 4: Enforcement function

Figure 4 gives a scheme of the behaviour of the enforcement function. Intuitively, σ_s is the word that can be released as output, whereas σ_c is the buffer containing the events that are already read/received, but cannot be released as output yet because they lead to an unsafe state from which it would be possible to violate the property reading only uncontrollable events (i.e. they lead to a vertex in $W_1 = V \setminus W_0$). Upon receiving a new event a , the enforcement mechanism distinguishes two cases:

- If a belongs to Σ_u , then it is output, as required by compliance. Then, the longest prefix of σ_c that satisfies φ and leads to a vertex in W_0 is also output.
- If a is in Σ_c , then it is added to σ_c , and the longest prefix of this new buffer that satisfies φ and leads to a vertex in W_0 is emitted, if it exists.

In both cases, κ_φ is used to compute the longest word that can be output, that is the longest word in G for the state reached so far with the current buffer of the enforcement mechanism, or ϵ if this set is empty. The parameters of κ_φ are those which are passed to G , they correspond to the state reached so far by the output of the enforcement mechanism, and its current buffer, respectively.

Some properties are not enforceable (see Example 3.5), but receiving some events may lead to a state from which it is possible to enforce. Therefore, it is possible to define a set of words, called $\text{Pre}(\varphi)$, such that E_φ is sound in $\text{Pre}(\varphi)$, as stated in Proposition 3.14:

⁴ E_φ and store_φ depend on Σ_u and Σ_c , but we did not write it in order to lighten the notations.

Definition 3.11 (Pre). The set of input words $\text{Pre}(\varphi) \subseteq \Sigma^*$ is defined as follows:

$$\text{Pre}(\varphi) = \{\sigma \in \Sigma^* \mid G(\text{Reach}(\sigma|_{\Sigma_u}), \sigma|_{\Sigma_c}) \neq \emptyset\} \cdot \Sigma_c^*$$

Intuitively, $\text{Pre}(\varphi)$ is the set of words in which E_φ is sound. This set is extension-closed, as required by Definition 3.2. In E_φ , using W_0 ensures that once the set G is not empty, it will never be afterwards, no matter what events are received. Thus, $\text{Pre}(\varphi)$ is the set of input words such that the output of E_φ belongs to G . Since E_φ outputs only uncontrollable events until G becomes non-empty, the definition of $\text{Pre}(\varphi)$ considers that the state reached is the one that is reached by emitting only the uncontrollable events of σ , and the corresponding buffer would then be the controllable events of σ .

Example 3.12. Considering the property φ_{ex} as shown in Fig. 2, with the uncontrollable alphabet $\Sigma_u = \{\text{Auth}, \text{LockOff}, \text{LockOn}\}$, $\text{Pre}(\varphi_{\text{ex}}) = \text{Write}^* \cdot \text{Auth} \cdot \Sigma^*$. Indeed, from the initial state q_0 , if an uncontrollable event, say LockOff , is received, then q_3 is reached, which is a non-accepting sink state, and thus any vertex in $\{q_3\} \times \Sigma_c^n \times \{0, 1\}$ will not be in W_0 . In order to reach a vertex in W_0 (i.e. a vertex in $\{q_1, q_2\} \times \Sigma_c^n \times \{0, 1\}$), it is necessary to read Auth . Once Auth is read, q_1 is reached, and from there, all uncontrollable events lead to either q_1 or q_2 . The same holds true from q_2 . Thus, it is possible to stay in the accepting states q_1 and q_2 , by delaying Write events when in q_2 until a LockOff event is received. Consequently, $\{q_1, q_2\} \times \Sigma_c^n \times \{0, 1\} \subseteq W_0$, and thus $\text{Pre}(\varphi_{\text{ex}}) = \text{Write}^* \cdot \text{Auth} \cdot \Sigma^*$, since Write events can be buffered while in state q_0 until event Auth is received, leading to a vertex in $\{q_1\} \times (\text{Write}^* \cap \Sigma_c^n) \times \{0, 1\} \subseteq W_0$.

Considering the property φ_{ex} defined in Fig. 2, we illustrate in Table 1 the enforcement function by showing the evolution of σ_s and σ_c with input $\sigma = \text{Auth} \cdot \text{LockOn} \cdot \text{Write} \cdot \text{LockOff}$.

Table 1: Evolution of $(\sigma, \langle \sigma_s, \sigma_c \rangle) \in \text{store}_{\varphi_{\text{ex}}}$

σ	σ_s	σ_c
ϵ	ϵ	ϵ
Auth	Auth	ϵ
$\text{Auth} \cdot \text{LockOn}$	$\text{Auth} \cdot \text{LockOn}$	ϵ
$\text{Auth} \cdot \text{LockOn} \cdot \text{Write}$	$\text{Auth} \cdot \text{LockOn}$	Write
$\text{Auth} \cdot \text{LockOn} \cdot \text{Write} \cdot \text{LockOff}$	$\text{Auth} \cdot \text{LockOn} \cdot \text{LockOff}$	ϵ
LockOff	Write	

E_φ (as per Definition 3.10) is an enforcement function that is sound with respect to φ in $\text{Pre}(\varphi)$, compliant with respect to Σ_u and Σ_c , and optimal in $\text{Pre}(\varphi)$.

PROPOSITION 3.13. E_φ is an enforcement function as per Definition 3.1.

Sketch of proof. We have to show that for all σ and $\sigma' \in \Sigma^*$, $(\sigma, \sigma_o) \in E_\varphi \wedge (\sigma \cdot \sigma', \sigma'_o) \in E_\varphi \implies \sigma_o \preceq \sigma'_o$. Following the definition of store_φ , this holds provided that $\sigma' \in \Sigma$ (i.e. σ' is a word of size 1). Since \preceq is an order, it follows that the proposition holds for all $\sigma' \in \Sigma'$.

PROPOSITION 3.14. E_φ is sound with respect to φ in $\text{Pre}(\varphi)$, as per Definition 3.2.




Figure 5: Property that can be enforced by blocking all controllable events c , thus outputting only the uncontrollable ones u .

Sketch of proof. We have to show that if $\sigma \in \text{Pre}(\varphi)$, then $(\sigma, \sigma_o) \in E_\varphi \implies \sigma_o \models \varphi$. The proof is made by induction on σ . In the induction step, considering $a \in \Sigma$, we distinguish three cases:

- (1) $\sigma \cdot a \notin \text{Pre}(\varphi)$. Then the proposition holds.
- (2) $\sigma \cdot a \in \text{Pre}(\varphi)$, but $\sigma \notin \text{Pre}(\varphi)$. Then the input reaches $\text{Pre}(\varphi)$, and since it is extension-closed, all extensions of σ also are in $\text{Pre}(\varphi)$, and we prove that the proposition holds considering the definition of $\text{Pre}(\varphi)$.
- (3) $\sigma \in \text{Pre}(\varphi)$ (and thus, $\sigma \cdot a \in \text{Pre}(\varphi)$ since it is extension-closed). Then, we prove that the proposition holds, based on the definition of store_φ , and more precisely on the definition of G , that uses W_0 to ensure that there always exists a compliant output that satisfies φ .

PROPOSITION 3.15. E_φ is compliant, as per Definition 3.3.

Sketch of proof. The proof is made by induction on the input $\sigma \in \Sigma^*$. Considering $\sigma \in \Sigma^*$ and $a \in \Sigma$, the proof is straightforward by considering the different values of $(\sigma \cdot a, \sigma_o) \in \text{store}_{\varphi}$, $(\sigma \cdot a)|_{\Sigma_u}$, and $(\sigma \cdot a)|_{\Sigma_c}$, when $a \in \Sigma_c$ and $a \in \Sigma_u$.

REMARK 1. Notice that for some properties, an enforcement function that would block all controllable events may still be sound and compliant. Consider for instance the property represented in Fig. 5, where c is a controllable event, and u an uncontrollable event. Then, outputting only the event u and buffering all the c events allows us to stay in state q_0 , which is sound since $\{q_0\} \times (c^* \cap \Sigma_c^n) \times \{0, 1\} \subseteq W_0$. This means that an enforcement mechanism that blocks all controllable events would be sound and compliant. Nevertheless, if $c \cdot c$ is received, it can be output to reach state q_2 , which is also accepting and $\{q_2\} \times \Sigma_c^n \times \{0, 1\} \subseteq W_0$. Then it is possible to release more events. Therefore, an enforcement mechanism that would output two c events when they are received would be “better” than the first one blocking all of them, in the sense that its output would be longer (and thus closer to the input).

For any given input $\sigma \in \text{Pre}(\varphi)$, $E_\varphi(\sigma)$ is the longest possible word that ensures soundness and compliance, i.e. controllable events are blocked only when necessary. Thus, E_φ is also optimal in $\text{Pre}(\varphi)$:

PROPOSITION 3.16. E_φ is optimal in $\text{Pre}(\varphi)$, as per Definition 3.4.

Sketch of proof. The proof is made by induction on the input $\sigma \in \Sigma^*$. Once $\sigma \in \text{Pre}(\varphi)$, we know that $(\sigma, \sigma_o) \in E_\varphi \implies \sigma_o \models \varphi$ since E_φ is sound in $\text{Pre}(\varphi)$. E_φ is optimal because, in store_φ , κ_φ provides the longest possible word. If a longer word were output, then either the output would not satisfy φ , or it would lead to a vertex that is not in W_0 , meaning that there would exist an uncontrollable word leading to a non-accepting state and to a vertex that would not be in W_0 . Then, the enforcement mechanism would have to output some controllable events from the buffer to reach an accepting

state, but since the vertex is not in W_0 , there would exist again an uncontrollable word leading to a non-accepting state and a vertex not in W_0 . By iterating, the buffer would become ϵ whereas the output of the enforcement mechanism would be leading to a non-accepting state. Therefore, outputting a longer word would mean that the function is not sound. This means that E_φ is optimal in $\text{Pre}(\varphi)$, since it outputs the longest word that allows us to be both sound and compliant.

3.3 Enforcement Monitors

Enforcement monitors are operational descriptions of enforcement mechanisms. We give a representation of an enforcement mechanism for a property φ as an input/output transition system. The input/output behaviour of the enforcement monitor is the same as the one of the enforcement function E_φ defined in Section 3.2. Enforcement monitors are purposed to ease the implementation of enforcement mechanisms, since they give an operational representation of the enforcement mechanism.

Definition 3.17 (Enforcement monitor). An enforcement monitor \mathcal{E} for φ is a transition system $\langle C^\mathcal{E}, c_0^\mathcal{E}, \Gamma^\mathcal{E}, \hookrightarrow_\mathcal{E} \rangle$ such that:

- $C^\mathcal{E} = Q \times \Sigma^*$ is the set of configurations.
- $c_0^\mathcal{E} = \langle q_0, \epsilon \rangle$ is the initial configuration.
- $\Gamma^\mathcal{E} = \Sigma^* \times \{\text{dump}(\cdot), \text{pass-uncont}(\cdot), \text{store-cont}(\cdot)\} \times \Sigma^*$ is the alphabet, where the first, second, and third members are an input sequence, an enforcement operation, and an output sequence, respectively.
- $\hookrightarrow_\mathcal{E} \subseteq C^\mathcal{E} \times \Gamma^\mathcal{E} \times C^\mathcal{E}$ is the transition relation, defined as the smallest relation obtained by applying the following rules in order (where $w / \bowtie w'$ stands for $(w, \bowtie, w') \in \Gamma^\mathcal{E}$):
 - **Dump:** $\langle q, a.\sigma_c \rangle \xrightarrow{\epsilon / \text{dump}(a)/a}_\mathcal{E} \langle q', \sigma_c \rangle$, if $a \in \Sigma_c$, $G(q, a.\sigma_c) \neq \emptyset$ and $G(q, a.\sigma_c) \neq \{\epsilon\}$, with $q' = q$ after a ,
 - **Pass-uncont:** $\langle q, \sigma_c \rangle \xrightarrow{a / \text{pass-uncont}(a)/a}_\mathcal{E} \langle q', \sigma_c \rangle$, with $a \in \Sigma_u$ and $q' = q$ after a ,
 - **Store-cont:** $\langle q, \sigma_c \rangle \xrightarrow{a / \text{store-cont}(a)/\epsilon}_\mathcal{E} \langle q, \sigma_c.a \rangle$, with $a \in \Sigma_c$.

In \mathcal{E} , a configuration $c = \langle q, \sigma \rangle$ represents the current state of the enforcement mechanism. The state q is the one reached so far in \mathcal{A}_φ with the output of the monitor. The word of controllable events σ represents the buffer of the monitor, i.e. the controllable events of the input that it has not output yet. Rule **dump** outputs the first event of the buffer if it can ensure soundness afterwards (i.e. if there is a non-empty word in G , that must begin with this event). Rule **pass-uncont** releases an uncontrollable event as soon as it is received. Rule **store-cont** simply adds a controllable event at the end of the buffer. Compared to Section 3.2, the second member of the configuration represents buffer σ_c in the definition of store_φ , whereas σ_s is here represented by state q which is the first member of the configuration, such that $q = \text{Reach}(\sigma_s)$.

PROPOSITION 3.18. *The output of the enforcement monitor \mathcal{E} for input σ is $E_\varphi(\sigma)$.*

In Proposition 3.18, the output of the enforcement monitor is the concatenation of all the outputs of the word labelling the path followed when reading σ .

Sketch of proof. The proof is made by induction on the input $\sigma \in \Sigma^*$. We just consider the rules that can be applied when receiving a new event. If the event is controllable, then rule **store-cont**() can be applied, possibly followed by rule **dump**() applied once or more times. If the event is uncontrollable, then rule **pass-uncont**() can be applied, again possibly followed by rule **dump**() applied once or more times. Since rule **dump**() applies only when there is a non-empty word in G , then this word must begin with the first event of the buffer, and the rule **dump**() can be applied again if there was a word in G of size at least 2, meaning that there is another non-empty word in the new set G ... Thus, the output of all the applications of the rule **dump**() corresponds to the computation of κ_φ in the definition of store_φ , and consequently the outputs of \mathcal{E} and E_φ are the same.

4 Algorithms and Implementation

We describe some of the algorithms that allow us to use a game graph (as per Definition 3.6) to define an enforcement mechanism, and discuss their time complexity. We suppose that the set of winning nodes of the graph is known, as there exist well-known algorithms to compute it (see for example [5]).

4.1 Algorithms

Algorithm 1 computes the set Σ_c^n (see Section 3.2), the set of words that allow to reach a state that is unreachable with all its prefixes from at least one state. Algorithm 1 uses a recursive function described in Algorithm 2. The algorithm builds the words incrementally, adding each possible event to a word in the set, until adding an event does not allow to reach a new state from any state. We make use of arrays of one and two dimensions. Function `arrayInit(m, n)` returns an array of m rows and n columns, when n is not specified, it returns a 1-dimensional array of size m . The returned array is filled with 0.

```

input : An automaton  $\mathcal{A} = \langle Q = \{q_i \mid i \in [1; n]\}, q_0, \Sigma = \Sigma_u \cup \Sigma_c, \delta, F \rangle$ 
output : The set  $\Sigma_c^n$  as defined in Section 3.2

1 reachable  $\leftarrow$  arrayInit( $n, n$ );
2 lasts  $\leftarrow$  arrayInit( $n$ );
3  $\Sigma_c^n \leftarrow \{\epsilon\}$ ;
4 for  $i \leftarrow 1$  to  $n$  do
5   | reachable[ $i, i$ ]  $\leftarrow 1$ ;
6   | lasts[ $i$ ]  $\leftarrow q_i$ ;
7 end
8 foreach  $c \in \Sigma_c$  do
9   |  $\Sigma_c^n \leftarrow \Sigma_c^n \cup \text{compute}\Sigma_c^n\text{Rec}(\mathcal{A}, c, \text{reachable}, \text{lasts})$ ;
10 end

```

Algorithm 1: Algorithm for computing Σ_c^n

Algorithm 1 first initialises Σ_c^n to $\{\epsilon\}$. Then, the two arrays `reachable` and `lasts` are initialised accordingly, i.e. `reachable` is filled with 0, with 1 on the diagonal, and `lasts`[i] is q_i . Words are then added to Σ_c^n by calling the recursive function `computeSigma_c^nRec` described in Algorithm 2. The array `reachable` is an array of size $n \times n$, where $n = |Q|$, such that `reachable`[i, j] is equal to 1 if there is a prefix w' of w such that q_i after $w' = q_j$, and 0 otherwise. The array `lasts` is an array of size n , such that for all $i \in [1; n]$, `lasts`[i] = q_i after w . The function considers recursively all the extensions of w , and add them to Σ_c^n until the array `reachable` stabilises, and then returns the computed Σ_c^n .


```

input : An automaton  $\mathcal{A} = \langle Q = \{q_i \mid i \in [1; n]\}, q_0, \Sigma = \Sigma_u \cup \Sigma_c, \delta, F \rangle$ ,
        reachable, lasts as defined in Algorithm 1,  $w \in \Sigma_c^n$ 
output : The set of all the extensions of  $w$  that belong to  $\Sigma_c^n$ 
1 Function compute $\Sigma_c^n$ Rec ( $\mathcal{A}, w, \text{reachable}, \text{lasts}$ ):
2    $\Sigma_c^n \leftarrow \emptyset$ ;
3   foreach  $c \in \Sigma_c$  do
4     reachableAfter  $\leftarrow$  reachable;
5     for  $i \leftarrow 1$  to  $n$  do
6       let  $j \in [1; n]$  be such that lasts[i] after  $c = q_j$ ;
7       reachableAfter[i,j]  $\leftarrow 1$ ;
8       lastsAfter[i]  $\leftarrow q_j$ ;
9     end
10    if reachableAfter  $\neq$  reachable then
11       $\Sigma_c^n \leftarrow \Sigma_c^n \cup \{w\} \cup$ 
12        compute $\Sigma_c^n$ Rec( $\mathcal{A}, w \cdot c, \text{reachableAfter}, \text{lastsAfter}$ );
13    end
14  end
15  return  $\Sigma_c^n$ ;

```

Algorithm 2: Function compute Σ_c^n Rec

The worst-case complexity of Algorithm 1 in terms of assignments is at most $2n + |\Sigma_c|^{n^2-n}$, where $n = |Q|$.

The following algorithms define the primitives of an enforcement mechanism. A state of the enforcement mechanism is represented by a tuple in $V \times V \times \Sigma_c^* \times \Sigma_c^* \times \Sigma_c^*$, where V is the set of nodes of the graph \mathcal{G} defined in Definition 3.7. The first node is the node reached by the output of the enforcement mechanism (real node). The second node is the strategy node, i.e. the first winning node that can be reached by outputting the first events of the buffer, or the node reached by outputting all the buffer if such a node does not exist. The third member is the buffer composed of the controllable events that have not been output yet. The fourth member is the input, and the fifth is the output.

The first function is `enforcerInit` (not provided here, due to the lack of space), that returns the initial state of the enforcer, i.e. $\langle \langle q_0, \epsilon \rangle, \langle q_0, \epsilon \rangle, \epsilon, \epsilon, \epsilon \rangle$.

```

input : A state  $\langle \langle q_r, b_r, 0 \rangle, \langle q_s, b_s, 0 \rangle, b, i, o \rangle$  of the enforcer, an event  $e \in \Sigma$ 
output : The state of the enforcer after having received  $e$ 
1 Function enforcerEventReceived ( $\langle \langle q_r, b_r, 0 \rangle, \langle q_s, b_s, 0 \rangle, b, i, o \rangle, e$ ):
2   if  $e \in \Sigma_c$  then
3     if  $b_r \cdot e \in \Sigma_c^n$  then
4        $r \leftarrow \langle q_r, b_r \cdot e, 0 \rangle$ ;
5     else
6        $r \leftarrow \langle q_r, b_r, 0 \rangle$ ;
7     end
8      $w \leftarrow \min_{\preceq} (\{w' \preceq b \mid w'^{-1} \cdot (b_s \cdot e) \in \Sigma_c^n\})$ ;
9      $s \leftarrow \langle q_s \text{ after } w, w^{-1} \cdot (b_s \cdot e), 0 \rangle$ ;
10    return  $\langle r, s, b \cdot e, i \cdot e, o \rangle$ ;
11  else /*  $e \in \Sigma_u$  */
12     $r \leftarrow \langle q_r \text{ after } e, b_r, 0 \rangle$ ;
13     $w \leftarrow \min_{\preceq} (\{w' \preceq b \mid \langle q_r \text{ after } e \text{ after } w', \max_{\preceq} (\{w'' \preceq w'^{-1} \cdot b \mid w'' \in \Sigma_c^n\}, 0) \in W_0\} \cup \{b\}\})$ ;
14     $s \leftarrow \langle q_r \text{ after } e \text{ after } w, \max_{\preceq} (\{w' \preceq w^{-1} \cdot b \mid w' \in \Sigma_c^n\}, 0) \rangle$ ;
15    return  $\langle r, s, b, i \cdot e, o \cdot e \rangle$ ;
16  end
17 end

```

Function enforcerEventReceived(state, event)

Function `enforcerEventReceived` computes the next state of the enforcer after the reception of an event. If the event is controllable, then only the strategy node is updated; if the event is uncontrollable, then it is immediately emitted and the real node is

changed accordingly, then the strategy node is computed from this new node. The time complexity of this function is linear in the size of the buffer.

```

input : A state  $\langle r, s, b, i, o \rangle$  of the enforcer
output : EMIT if it is possible to emit some controllable events, DONTEMIT otherwise
1 Function enforcerGetStrat ( $\langle r, s, b, i, o \rangle$ ):
2   if  $r \neq s$  and  $s \in W_0$  then
3     strat  $\leftarrow$  EMIT;
4   else
5     strat  $\leftarrow$  DONTEMIT;
6   end
7   return strat;
8 end

```

Function enforcerGetStrat(state)

```

input : A state  $\langle \langle q_r, b_r, 0 \rangle, \langle q_s, b_s, 0 \rangle, e \cdot b, i, o \rangle$  of the enforcer, where
         $e \in \Sigma_c$  and  $b \in \Sigma_c^*$ 
output : The state of the enforcer after having emitted the first controllable event
        of its buffer
1 Function enforcerEmit ( $\langle \langle q_r, b_r, 0 \rangle, \langle q_s, b_s, 0 \rangle, e \cdot b, i, o \rangle$ ):
2    $s \leftarrow \langle q_s, b_s, 0 \rangle$ ;
3    $r \leftarrow \langle q_r \text{ after } e, b, 0 \rangle$ ;
4    $w \leftarrow \min_{\preceq} (\{w' \preceq b \mid \langle q_r \text{ after } e \text{ after } w', \max_{\preceq} (\{w'' \preceq w'^{-1} \cdot b \mid w'' \in \Sigma_c^n\}, 0) \in W_0\} \cup \{b\}\})$ ;
5    $s \leftarrow \langle q_r \text{ after } e \text{ after } w, \max_{\preceq} (\{w' \preceq w^{-1} \cdot b \mid w' \in \Sigma_c^n\}, 0) \rangle$ ;
6   return  $\langle r, s, b, i, o \cdot e \rangle$ 
7 end

```

Function enforcerEmit(state)

Function `enforcerGetStrat` returns the strategy to follow. If the strategy node is ahead of the real node, and it is a winning node, then the strategy is to emit the first event of the buffer. Otherwise, the strategy is not to emit. This function has a constant time complexity.

Function `enforcerEmit` emits the first event of the buffer. The real node is updated accordingly, and the strategy node is unchanged except if the real node caught up with it (i.e. they are equal), in which case the strategy node is updated accordingly. The time complexity of this function is linear in the size of the buffer.

```

input : A property  $\varphi$ , described by an automaton  $\mathcal{A}$ , the game graph  $\mathcal{G}$ 
        associated to  $\varphi$ , the input sequence of events, through the function
        read()
output : The output of the enforcer mechanism
1 EM  $\leftarrow$  enforcerInit ( $\mathcal{A}, \mathcal{G}$ );
2 while The input sequence has not been read entirely do
3   e  $\leftarrow$  read();
4   EM  $\leftarrow$  enforcerEventReceived (EM, e);
5   while enforcerGetStrat (EM) = EMIT do
6     EM  $\leftarrow$  enforcerEmit (EM);
7   end
8 end

```

Algorithm 3: Main algorithm to enforce a property

Algorithm 3 describes the main algorithm that uses all these functions to actually enforce a property. It needs one more function: `read()` which returns the next input event. The algorithm first creates an enforcement monitor for the given property with `enforcerInit`, then all the events from the input are read with function `read`, and fed to the enforcer with function

`enforcerEventReceived`. Then, the enforcer emits a maximal number of events from its buffer with `enforcerEmit`, i.e. until `enforcerGetStrat` indicates that it should not emit. Note that emitting events is done by adding the events to the output of the enforcer. The time complexity of this algorithm between two calls to `read` is linear in the size of the buffer, and thus does not depend on the size of the automaton.

4.2 Implementation

We implemented the algorithms in the C programming language. Our tool takes as input a file describing an automaton, and reads the events from its standard input. It outputs information on the evolution of the state of the enforcement mechanism as well as a summary of the execution when it has ended on its standard output. This approach allows us to adapt easily the tool in order to use it with off-the-shelf applications. The tool first creates the game graph from the automaton, then solves the Büchi game on it. Then, the enforcement mechanism is initialised, and then used to compute the controllable events that can be emitted. When the end of the input is reached, the tool displays first the input of the enforcement mechanism, then its output and the remaining events of its buffer at the end of the the execution. Lastly, it displays a verdict indicating whether the execution ended in an accepting state of the automaton.

Performance Analysis. We provide some execution times of our tool, running on the example given in the paper (φ_{ex}). We evaluated it on 20 randomly selected inputs of 20 events each. Table 2 shows the inputs and the corresponding mean times taken by the enforcement mechanism to compute its output after each event. The times are given in nanoseconds. The means have been computed over 100 iterations. The inputs are abbreviated: *w* stands for *Write*, *f* for *LockOff*, *n* for *LockOn*, and *a* for *Auth*. The results have been obtained using a computer running Ubuntu 16.04 with a 3.40 GHz Intel Core i7 CPU with 16GB RAM. In Table 2, all the values have the same order of magnitude, showing that the execution time depends little on the input sequence. Thus, the execution time of our enforcement mechanism is quite stable, which makes it reliable.

5 RELATED WORK AND DISCUSSION

Runtime enforcement was pioneered by the work of Schneider with security automata [15], a runtime mechanism for enforcing safety properties. In [15], monitors are able to stop the execution of the system once a deviation of the property has been detected. Later, Ligatti et al. proposed edit-automata [11], a more powerful model of enforcement monitors able to insert and suppress events from the execution, thus permitting to enforce non-safety properties. Later, Falcone et al. proposed more general models where the monitors can be synthesised from regular properties [10]. Another recent approach by Dolzhenko et al. [6] introduces Mandatory Result Automata (MRAs). MRAs extend edit-automata by refining the input-output relationship of an enforcement mechanism and thus allowing a more precise description of the enforcement abilities of an enforcement mechanism in concrete application scenarios. All these approaches do not consider uncontrollable events.

Basin et al. [2] introduced uncontrollable events for security automata [15]. The approach in [2] allows to enforce safety properties where some of the events in the specification are uncontrollable.

More recently, they proposed a more general approach [1] related to enforcement of security policies with controllable and uncontrollable events. They presented several complexity results and how to synthesise enforcement mechanisms, but they did not provide a tool implementation.

To our knowledge, two runtime enforcement methods using games have been proposed. In [3], Bloem et al. focus on enforcement of safety properties for reactive hardware systems, i.e. systems with boolean signals as inputs and outputs. They propose a method based on a 2-players safety games in order to synthesise a variant of an enforcement monitor called a *safety shield* and present a tool implementing this approach. This shield ensures correctness (i.e. soundness), and minimum interference according to a notion of distance permitting to measure the deviation between the output and the input of the shield. More recently, Wu et al. propose in [16] to improve the algorithm of [3] in the way that it takes the best recovery strategy among all possible ones, permitting to minimise the deviation, especially in case of burst errors. These two approaches are limited to safety properties and do not support uncontrollable events.

6 CONCLUSION AND FUTURE WORK

This paper revisits the work done in [13, 14] and introduces another way to compute the behaviour of the defined enforcement mechanisms in case of uncontrollable events using a Büchi game. Thus, we define enforcement monitors at two levels of abstraction, one is functional and the second operational. As in [13, 14], we consider that some events are uncontrollable, meaning that they are only observable by the enforcement monitor, that must output them when they are received. We introduce a different way to compute the behaviour of the enforcement mechanism using a Büchi game, that is equivalent to the behaviour of the enforcement mechanism described in [13, 14]. Given a property, we build a graph over which we solve a Büchi game representing the behaviour of the enforcement mechanism. Even though this graph should have an infinite number of vertices, it is possible to reduce it to a finite number, which allows us to store it. The behaviour of the enforcement mechanism only depends on the “current” vertex in the graph. When receiving an event, the enforcement mechanism only updates the current vertex by following some path in the graph, instead of computing again the set of winning configurations. This reduces the time spent in computing the behaviour of the enforcement mechanism, which is one of the main inconveniences when using enforcement mechanisms on running systems. Computing and storing the graph allows us to compute offline the behaviour of the enforcement mechanism, providing better performance at runtime.

We also provide a new way to describe soundness, compliance and optimality with a global view of the system based on inputs and outputs, and we present an implementation showing the effectiveness of the approach.

As a future work, we intend to investigate the runtime enforcement of timed properties [12] using timed games. We consider using the method introduced in this paper to enforce timed properties, as done in [9] without uncontrollable events and with suppression; and in [13] with uncontrollable events and without suppression. The

Table 2: Table of the mean execution times of our tool for different inputs

Input	Times																			
fnawwwaawawawfnwvnn	724	515	355	275	318	287	677	471	321	253	592	233	453	230	500	437	231	241	555	607
ffnnwfaawnwfaaffnnawnf	707	534	364	423	251	345	435	222	314	285	794	389	337	320	328	321	319	244	453	969
nnwnfnannnawwfaanwnf	726	507	257	470	336	282	424	278	268	428	305	218	290	617	372	312	316	241	436	495
aafaaawnffwnfwfaana	700	498	394	414	281	320	545	292	320	427	322	288	279	282	273	275	274	301	383	314
nwanfnwafnaawfaaww	728	307	505	460	310	294	412	216	299	317	425	285	277	276	279	610	368	310	238	379
awfwawawaawfaannwna	689	617	464	459	318	432	364	287	307	426	285	276	317	274	267	293	287	217	437	318
nwaanfnfnwfnanwf	692	304	472	452	313	275	417	308	277	442	300	281	216	293	610	370	317	312	243	562
aaffwaaanaanaanwfwan	677	502	392	429	468	346	354	273	314	413	271	278	264	263	278	221	450	324	377	318
wannafnnwffaafwfwfn	74	1309	432	421	293	444	353	272	222	621	288	274	275	275	326	289	272	309	380	315
nfwafnwnwnfwnaafa	710	302	492	221	452	317	351	347	591	256	586	226	453	224	453	428	402	396	458	516
awwnafannwfwfnwn	680	620	405	461	344	342	405	281	316	432	223	254	571	300	283	303	283	383	286	365
wfnwvanwnwnfffwvvnf	77	1010	501	246	426	629	454	336	420	240	617	414	401	386	380	379	230	244	531	663
aanannffnaafwnaawffn	690	506	379	513	326	333	388	293	306	421	292	284	470	287	280	278	224	474	394	320
fafwaaawannfafaafnn	735	503	374	363	422	283	405	221	288	323	428	282	285	279	276	272	268	291	315	379
nnnwvwwvwnfnwnwnvvnf	712	503	335	283	314	286	262	329	788	255	668	459	230	478	228	492	229	515	565	635
aaaawwffaffwvvnfnfa	679	512	345	409	483	369	430	315	309	434	294	279	295	287	290	287	287	311	387	320
nfnwnaanwanwnwvfwv	711	529	258	478	330	290	422	278	280	217	462	278	276	599	254	240	501	420	232	368
ffnawwafwfaanaanwnvf	704	519	352	418	240	230	365	394	278	630	513	325	328	305	297	318	237	236	497	528
wanawvwnavnfffvfnwv	106	1320	447	438	226	247	324	352	714	458	388	475	1077	312	282	346	374	432	324	204
annffnvnwvfwvfnvnn	678	527	342	439	299	331	359	276	306	598	348	278	355	283	283	282	272	306	380	308

graph would be bigger due to the presence of clocks in the automaton, but performance should be improved. The gain in computation would be higher, and this would even be more interesting in this setting since it might allow timed properties to be enforced on devices with little computational power.

Acknowledgement

The work reported in this article has been done in the context of the COST Action ARVI IC1402, supported by COST (European Cooperation in Science and Technology). It was also partially supported by French Region Nouvelle Aquitaine, and Bordeaux INP.

REFERENCES

[1] David Basin, Vincent Jugé, Felix Klaedtke, and Eugen Zălinescu. 2013. Enforceable Security Policies Revisited. *ACM Trans. Inf. Syst. Secur.* 16, 1, Article 3 (June 2013), 26 pages. DOI : <http://dx.doi.org/10.1145/2487222.2487225>

[2] David Basin, Felix Klaedtke, and Eugen Zălinescu. 2011. Algorithms for Monitoring Real-Time Properties. In *Proceedings of the 2nd International Conference on Runtime Verification (RV 2011) (Lecture Notes in Computer Science)*, Sarfraz Khurshid and Koushik Sen (Eds.), Vol. 7186. Springer-Verlag, 260–275.

[3] Roderick Bloem, Bettina Könighofer, Robert Könighofer, and Chao Wang. 2015. Shield Synthesis: - Runtime Enforcement for Reactive Systems. In *Tools and Algorithms for the Construction and Analysis of Systems - 21st International Conference, TACAS 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings.* 533–548.

[4] Peter Bulychov, Thomas Chatain, Alexandre David, and Kim G. Larsen. 2009. *Efficient on-the-fly Algorithm for Checking Alternating Timed Simulation*. Springer Berlin Heidelberg, Berlin, Heidelberg, 73–87. DOI : http://dx.doi.org/10.1007/978-3-642-04368-0_8

[5] Krishnendu Chatterjee and Monika Henzinger. 2012. An O(N2) Time Algorithm for Alternating BÜChi Games. In *Proceedings of the Twenty-third Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '12)*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1386–1399. <http://dl.acm.org/citation.cfm?id=2095116.2095225>

[6] Egor Dolzhenko, Jay Ligatti, and Srikar Reddy. 2015. Modeling Runtime Enforcement with Mandatory Results Automata. *International Journal of Information Security* 14, 1 (Feb. 2015), 47–60. <http://dx.doi.org/10.1007/s10207-014-0239-8>

[7] Yliès Falcone. 2010. You Should Better Enforce Than Verify. In *Runtime Verification - First International Conference, RV 2010, St. Julians, Malta, November*

1-4, 2010. *Proceedings (Lecture Notes in Computer Science)*, Howard Barringer, Yliès Falcone, Bernd Finkbeiner, Klaus Havelund, Insup Lee, Gordon J. Pace, Grigore Rosu, Oleg Sokolsky, and Nikolai Tillmann (Eds.), Vol. 6418. Springer, 89–105. DOI : http://dx.doi.org/10.1007/978-3-642-16612-9_9

[8] Yliès Falcone, Jean-Claude Fernandez, and Laurent Mounier. 2012. What can you verify and enforce at runtime? *International Journal on Software Tools for Technology Transfer* 14, 3 (2012), 349–382. DOI : <http://dx.doi.org/10.1007/s10009-011-0196-8>

[9] Yliès Falcone, Thierry Jéron, Hervé Marchand, and Srinivas Pinisetty. 2016. Runtime enforcement of regular timed properties by suppressing and delaying events. *Systems & Control Letters* 123 (2016), 2–41. DOI : <http://dx.doi.org/10.1016/j.scl.2016.02.008>

[10] Yliès Falcone, Laurent Mounier, Jean-Claude Fernandez, and Jean-Luc Richier. 2011. Runtime enforcement monitors: composition, synthesis, and enforcement abilities. *Formal Methods in System Design* 38, 3 (2011), 223–262. DOI : <http://dx.doi.org/10.1007/s10703-011-0114-4>

[11] Jay Ligatti, Lujo Bauer, and David Walker. 2009. Run-Time Enforcement of Nonsafety Policies. *ACM Trans. Inf. Syst. Secur.* 12, 3, Article 19 (Jan. 2009), 41 pages.

[12] Srinivas Pinisetty, Yliès Falcone, Thierry Jéron, Hervé Marchand, Antoine Rollet, and Omer Nguena-Timo. 2014. Runtime enforcement of timed properties revisited. *Formal Methods in System Design* 45, 3 (2014), 381–422. DOI : <http://dx.doi.org/10.1007/s10703-014-0215-y>

[13] Matthieu Renard, Yliès Falcone, Antoine Rollet, Thierry Jéron, and Hervé Marchand. To appear. Optimal Enforcement of (Timed) Properties with Uncontrollable Events. *Mathematical Structures in Computer Science (MSSCS)* (To appear).

[14] Matthieu Renard, Yliès Falcone, Antoine Rollet, Srinivas Pinisetty, Thierry Jéron, and Hervé Marchand. 2015. Enforcement of (Timed) Properties with Uncontrollable Events. In *Theoretical Aspects of Computing - ICTAC 2015 (Lecture Notes in Computer Science)*, Martin Leucker, Camilo Rueda, and Frank D. Valencia (Eds.), Vol. 9399. Springer International Publishing, 542–560. DOI : http://dx.doi.org/10.1007/978-3-319-25150-9_31

[15] Fred B. Schneider. 2000. Enforceable Security Policies. *ACM Trans. Inf. Syst. Secur.* 3, 1 (Feb. 2000), 30–50. DOI : <http://dx.doi.org/10.1145/353323.353382>

[16] Meng Wu, Haibo Zeng, and Chao Wang. 2016. Synthesizing Runtime Enforcer of Safety Properties under Burst Error. In *8th NASA Formal Methods Symposium NFM16*. Minneapolis, USA.

A PROOFS OF SECTION 3

In all this section, we will use the notations from Section 3, meaning that φ is a property whose associated automaton is $\mathcal{A}_\varphi = \langle Q, q_0, \Sigma, \rightarrow, F \rangle$. In some proofs, we also use notations from Definition 3.10. We also use a functional notation of enforcement functions, i.e. for an enforcement function E and a word $\sigma \in \Sigma^*$, we write $E(\sigma) = \sigma'$ whenever $(\sigma, \sigma') \in E$.

PROPOSITION 3.13. E_φ is an enforcement function as per Definition 3.1.

PROOF. Let us consider $\sigma \in \Sigma^*$, and $\sigma' \in \Sigma^*$. If $\sigma' = \epsilon$, then $E_\varphi(\sigma) = E_\varphi(\sigma.\sigma') \preceq E_\varphi(\sigma.\sigma')$. Otherwise, let us consider $(\sigma_s, \sigma_c) = \text{store}_\varphi(\sigma)$, $a = \sigma'(1)$, and $(\sigma_t, \sigma_d) = \text{store}_\varphi(\sigma.a)$. Then, if $a \in \Sigma_u$, $\sigma_t = \sigma_s.a.\sigma'_s$, where σ'_s is defined in Definition 3.10, meaning that $\sigma_s \preceq \sigma_t$. If $a \in \Sigma_c$, then $\sigma_t = \sigma_s.\sigma''_s$, where σ''_s is defined in Definition 3.10, thus again, $\sigma_s \preceq \sigma_t$. In both cases, $E_\varphi(\sigma) = \sigma_s \preceq \sigma_t = E_\varphi(\sigma.a)$. Since the order \preceq is transitive, this means that $E_\varphi(\sigma) \preceq E_\varphi(\sigma.a) \preceq E_\varphi(\sigma.a.\sigma'(2)) \preceq \dots \preceq E_\varphi(\sigma.\sigma')$. Thus E_φ is an enforcement function. \square

LEMMA A.1. $\forall q \in Q, \forall w \in \Sigma_c^n$,
 $\langle \langle q, w, 1 \rangle \in W_0 \wedge \langle q, w, 1 \rangle, \langle q', w', l \rangle \in E \rangle \implies \langle q', w', l \rangle \in W_0$.

PROOF. W_0 is the winning set of the Büchi game for P_0 . \square

LEMMA A.2. $\forall q \in Q, \forall \sigma \in \Sigma_c^n$,
 $\langle q, \sigma, 0 \rangle \in W_0 \implies G(q, \sigma) \neq \emptyset$.

PROOF. Let us consider $q \in Q$ and $\sigma \in \Sigma_c^n$ such that $\langle q, \sigma, 0 \rangle \in W_0$. Then, since $\langle q, \sigma, 0 \rangle$ is a vertex that belongs to P_0 that is winning (since it is in W_0), this means that there is a winning strategy for P_0 in the Büchi game. Thus, there is a path in \mathcal{G} that allows us P_0 to reach a Büchi state, that is a state in $F \times \Sigma_c^n \times \{0, 1\}$, whatever the strategy of P_1 is. The strategy of P_0 is to follow vertices that are only in W_0 until it finally reaches a Büchi state. The construction of W_0 ensures that this is possible. Now, the only edges that leave a vertex belonging to P_0 are the ones corresponding to the action of emitting the first of the stored controllable events, or not emitting it and let P_1 play. Thus, if $\langle q, \sigma, 0 \rangle \in W_0$, this means that there is a path in the graph that leads to a state in $F \times \Sigma_c^n \times \{0\}$ such that all the vertices along the path belong to P_0 and are in W_0 . This holds because there is a path in W_0 to such a state, and if a state of the path belongs to P_1 , then the strategy of P_1 could be to go back to the previous vertex belonging to P_0 , and thus there could be an infinite loop in these two vertices, meaning that they are in $F \times \Sigma_c^n \times \{0, 1\}$ or that from the previous state belonging to P_0 , emitting the first stored controllable event is a winning strategy. Thus, there exists $w \preceq \sigma$ such that q after $w \in F$ and $\langle q \text{ after } w, w^{-1}.\sigma, 0 \rangle \in W_0$. Now, since Σ_c^n is finite, it is possible to choose w such that $\langle q \text{ after } w, w^{-1}.\sigma, 1 \rangle \in W_0$, because otherwise, the only possible strategy would be to emit from every vertex, but it is not possible from vertices whose second member is ϵ , and then the only possible strategy would lead to a vertex not in W_0 , meaning that the original vertex would not be in W_0 , which is absurd. Thus, $G(q, \sigma) \neq \emptyset$. \square

LEMMA A.3. $\forall \sigma \in \Sigma^*$,
 $(\sigma \notin \text{Pre}(\varphi) \wedge (\sigma_s, \sigma_c) = \text{store}_\varphi(\sigma)) \implies (\sigma_s = \sigma|_{\Sigma_u} \wedge \sigma_c = \sigma|_{\Sigma_c})$.

PROOF. For $\sigma \in \Sigma^*$, let $P(\sigma)$ be the predicate “ $(\sigma \notin \text{Pre}(\varphi) \wedge (\sigma_s, \sigma_c) = \text{store}_\varphi(\sigma)) \implies (\sigma_s = \sigma|_{\Sigma_u} \wedge \sigma_c = \sigma|_{\Sigma_c})$ ”. Let us show by induction that $P(\sigma)$ holds for any $\sigma \in \Sigma^*$.

- *Induction basis:* $\text{store}_\varphi(\epsilon) = (\epsilon, \epsilon)$, and since $\epsilon|_{\Sigma_u} = \epsilon|_{\Sigma_c} = \epsilon$, $P(\epsilon)$ holds.
- *Induction step:* let us suppose that for $\sigma \in \Sigma^*$, $P(\sigma)$ holds. Let us then consider $a \in \Sigma$, $(\sigma_s, \sigma_b) = \text{store}_\varphi(\sigma)$, and $(\sigma_t, \sigma_d) = \text{store}_\varphi(\sigma.a)$. Then, if $\sigma.a \in \text{Pre}(\varphi)$, $P(\sigma.a)$ holds. Let us now consider that $\sigma.a \notin \text{Pre}(\varphi)$. Then, since $\text{Pre}(\varphi)$ is extension-closed, $\sigma \notin \text{Pre}(\varphi)$, and thus, by induction hypothesis, $\sigma_s = \sigma|_{\Sigma_u}$ and $\sigma_c = \sigma|_{\Sigma_c}$. We consider two cases:
 - Case 1: $a \in \Sigma_u$, then $\sigma_t = \sigma_s.a.\sigma'_s$, with $\sigma'_s \in G(\text{Reach}(\sigma_s.a), \sigma_c) \cup \{\epsilon\}$. Since $\sigma.a \notin \text{Pre}(\varphi)$, $G(\text{Reach}((\sigma.a)|_{\Sigma_u}), (\sigma.a)|_{\Sigma_c}) = \emptyset$. Moreover, since $a \in \Sigma_u$, $(\sigma.a)|_{\Sigma_u} = \sigma|_{\Sigma_u}.a = \sigma_s.a$ and $(\sigma.a)|_{\Sigma_c} = \sigma|_{\Sigma_c} = \sigma_c$, thus $G(\text{Reach}(\sigma_s.a), \sigma_c) = \emptyset$. It follows that $\sigma'_s \in \{\epsilon\}$, meaning that $\sigma_t = \sigma_s.a = \sigma|_{\Sigma_u}.a = (\sigma.a)|_{\Sigma_u}$, and $\sigma_d = \sigma_s'^{-1}.\sigma_c = \sigma_c = \sigma|_{\Sigma_c} = (\sigma.a)|_{\Sigma_c}$.
 - Case 2: $a \in \Sigma_c$, then $\sigma_t = \sigma_s.\sigma''_s$, with $\sigma''_s \in G(\sigma_s, \sigma_c.a) \cup \{\epsilon\}$. Since $\sigma.a \notin \text{Pre}(\varphi)$, $G(\text{Reach}((\sigma.a)|_{\Sigma_u}), (\sigma.a)|_{\Sigma_c}) = \emptyset$. Moreover, since $a \in \Sigma_c$, $(\sigma.a)|_{\Sigma_u} = \sigma|_{\Sigma_u} = \sigma_s$ and $(\sigma.a)|_{\Sigma_c} = \sigma|_{\Sigma_c}.a = \sigma_c.a$. Thus, $G(\text{Reach}(\sigma_s), \sigma_c.a) = \emptyset$, meaning that $\sigma''_s = \epsilon$. Thus, $\sigma_t = \sigma_s = \sigma|_{\Sigma_u} = (\sigma.a)|_{\Sigma_u}$ and $\sigma_d = \sigma_s'^{-1}.\sigma_c.a = \sigma_c.a = \sigma|_{\Sigma_c}.a = (\sigma.a)|_{\Sigma_c}$.

In both cases, $P(\sigma.a)$ holds.

By induction on $\sigma \in \Sigma^*$, for all $\sigma \in \Sigma^*$, if $\sigma \notin \text{Pre}(\varphi)$ and $(\sigma_s, \sigma_c) = \text{store}_\varphi(\sigma)$, then $\sigma_s = \sigma|_{\Sigma_u}$ and $\sigma_c = \sigma|_{\Sigma_c}$. \square

PROPOSITION 3.14. E_φ is sound with respect to φ in $\text{Pre}(\varphi)$, as per Definition 3.2.

PROOF. Let $P(\sigma)$ be the predicate: “ $(\sigma \in \text{Pre}(\varphi) \wedge (\sigma_s, \sigma_c) = \text{store}_\varphi(\sigma)) \implies (E_\varphi(\sigma) \models \varphi \wedge \langle \text{Reach}(\sigma_s), \max_{\preceq}(\{w \preceq \sigma_c \mid w \in \Sigma_c^n\}), 1 \rangle \in W_0)$ ”. Let us prove by induction that for any $\sigma \in \Sigma^*$, $P(\sigma)$ holds.

- *Induction basis:* if $\epsilon \in \text{Pre}(\varphi)$, then following the definition of $\text{Pre}(\varphi)$, $G(\text{Reach}(\epsilon), \epsilon) \neq \emptyset$. Thus $\epsilon \in G(\text{Reach}(\epsilon), \epsilon)$ (since ϵ is the only word satisfying $\epsilon \preceq \epsilon$). This means that $\text{Reach}(\epsilon)$ after $\epsilon = \text{Reach}(\epsilon) \in F$. Considering that $\text{store}_\varphi(\epsilon) = (\epsilon, \epsilon)$, it follows that $E_\varphi(\epsilon) = \epsilon$, and thus, $E_\varphi(\epsilon) \models \varphi$. Moreover, since $\epsilon \in G(\text{Reach}(\epsilon), \epsilon)$, $\langle \text{Reach}(\epsilon)$ after $\epsilon, \max_{\preceq}(\{w \preceq \epsilon^{-1}.\epsilon \mid w \in \Sigma_c^n\}), 1 \rangle = \langle \text{Reach}(\epsilon), \max_{\preceq}(\{w \preceq \epsilon^{-1}.\epsilon \mid w \in \Sigma_c^n\}), 1 \rangle \in W_0$. Thus $P(\epsilon)$ holds.
- *Induction step:* Suppose now that, for $\sigma \in \Sigma^*$, $P(\sigma)$ holds. Let us consider $a \in \Sigma$, $(\sigma_s, \sigma_c) = \text{store}_\varphi(\sigma)$, and $(\sigma_t, \sigma_d) = \text{store}_\varphi(\sigma.a)$. Let us prove that $P(\sigma.a)$ holds. We consider three different cases:
 - Case 1: $(\sigma.a) \notin \text{Pre}(\varphi)$. Then $P(\sigma.a)$ holds.
 - Case 2: $(\sigma.a) \in \text{Pre}(\varphi) \wedge \sigma \notin \text{Pre}(\varphi)$. Then, since $\text{Pre}(\varphi)$ is extension-closed, it follows that $\sigma.a \in \{w \in \Sigma^* \mid G(\text{Reach}(w|_{\Sigma_u}), w|_{\Sigma_c}) \neq \emptyset\}$, meaning that $G(\text{Reach}((\sigma.a)|_{\Sigma_u}), (\sigma.a)|_{\Sigma_c}) \neq \emptyset$. Moreover, since $\sigma \notin \text{Pre}(\varphi)$, following lemma A.3, $\sigma_s = \sigma|_{\Sigma_u}$ and $\sigma_c = \sigma|_{\Sigma_c}$. Now, we consider two cases:
 - If $a \in \Sigma_u$, then $(\sigma.a)|_{\Sigma_u} = \sigma|_{\Sigma_u}.a = \sigma_s.a$, and $(\sigma.a)|_{\Sigma_c} = \sigma|_{\Sigma_c} = \sigma_c$, thus $G(\text{Reach}(\sigma_s.a), \sigma_c) \neq \emptyset$, meaning that $\sigma'_s = (\sigma_s.a)^{-1}.\sigma_t \in G(\text{Reach}(\sigma_s.a), \sigma_c)$.

Thus, following the definition of G , $\text{Reach}(\sigma_s.a)$ after $\sigma'_s = \text{Reach}(\sigma_s.a.\sigma'_s) = \text{Reach}(\sigma_t) \in F$, and $\langle \text{Reach}(\sigma_s.a) \text{ after } \sigma'_s, \max_{\preceq}(\{w \preceq \sigma_s'^{-1}.\sigma_c \mid w \in \Sigma_c^n\}), 1 \rangle = \langle \text{Reach}(\sigma_t), \max_{\preceq}(\{w \preceq \sigma_d \mid w \in \Sigma_c^n\}), 1 \rangle \in W_0$. Since $\text{Reach}(\sigma_t) \in F$, $E_\varphi(\sigma.a) \models \sigma_t \models \varphi$. This means that $P(\sigma.a)$ holds.

- If $a \in \Sigma_c$, then $(\sigma.a)|_{\Sigma_u} = \sigma|_{\Sigma_u} = \sigma_s$, and $(\sigma.a)|_{\Sigma_c} = \sigma|_{\Sigma_c}.a = \sigma_c.a$. Thus, $G(\text{Reach}(\sigma_s), \sigma_c.a) \neq \emptyset$, meaning that $\sigma_s'' = \sigma_s^{-1}.\sigma_t \in G(\text{Reach}(\sigma_s), \sigma_c.a)$. As in the case where $a \in \Sigma_u$, it follows that $\langle \text{Reach}(\sigma_t), \max_{\preceq}(\{w \preceq \sigma_d \mid w \in \Sigma_c^n\}), 1 \rangle \in W_0$ and thus $E_\varphi(\sigma.a) \models \varphi$. This means that $P(\sigma.a)$ holds.

Thus, if $\sigma.a \in \text{Pre}(\varphi)$ but $\sigma \notin \text{Pre}(\varphi)$, $P(\sigma.a)$ holds.

Case 3: $\sigma \in \text{Pre}(\varphi)$ (and then $(\sigma.a) \in \text{Pre}(\varphi)$ since $\text{Pre}(\varphi)$ is extension-closed). Then, by induction hypothesis, $P(\sigma)$ holds, meaning that $E_\varphi(\sigma) \models \varphi$ and $\langle \text{Reach}(\sigma_s), \max_{\preceq}(\{w \preceq \sigma_c \mid w \in \Sigma_c^n\}), 1 \rangle \in W_0$. Let us note $\sigma_c^m = \max_{\preceq}(\{w \preceq \sigma_c \mid w \in \Sigma_c^n\})$. Again, we consider two cases:

- If $a \in \Sigma_u$, then, since $\langle \text{Reach}(\sigma_s), \sigma_c^m, 1 \rangle \in W_0$, following lemma A.1, since $\langle \text{Reach}(\sigma_s), \sigma_c^m, 1 \rangle, \langle \text{Reach}(\sigma_s) \text{ after } a, \sigma_c^m, 0 \rangle \in E_3 \subseteq E$, $\langle \text{Reach}(\sigma_s) \text{ after } a, \sigma_c^m, 0 \rangle = \langle \text{Reach}(\sigma_s.a), \sigma_c^m, 0 \rangle \in W_0$. Following lemma A.2, this means that $G(\text{Reach}(\sigma_s.a), \sigma_c) \neq \emptyset$, thus $\sigma'_s = (\sigma_s.a)^{-1}.\sigma_t \in G(\text{Reach}(\sigma_s.a), \sigma_c)$. It follows that $\text{Reach}(\sigma_s.a.\sigma'_s) = \text{Reach}(\sigma_t) \in F$, and that $\langle \text{Reach}(\sigma_s.a) \text{ after } a, \max_{\preceq}(\{w \preceq \sigma_s'^{-1}.\sigma_c \mid w \in \Sigma_c^n\}), 1 \rangle = \langle \text{Reach}(\sigma_t), \max_{\preceq}(\{w \preceq \sigma_d \mid w \in \Sigma_c^n\}), 1 \rangle \in W_0$. Thus, $E_\varphi(\sigma.a) = \sigma_t \models \varphi$. Thus $P(\sigma.a)$ holds.
- If $a \in \Sigma_c$, then, since $\langle \text{Reach}(\sigma_s), \sigma_c^m, 1 \rangle \in W_0$ and $\langle \text{Reach}(\sigma_s), \sigma_c^m, 1 \rangle, \langle \text{Reach}(\sigma_s), \max_{\preceq}(\{w \preceq \sigma_c.a \mid w \in \Sigma_c^n\}), 0 \rangle \in E_4 \cup E_5 \subseteq E$, following lemma A.2, this means that $G(\text{Reach}(\sigma_s), \sigma_c.a) \neq \emptyset$. Thus, $\sigma_s'' = \sigma_s^{-1}.\sigma_t \in G(\text{Reach}(\sigma_s), \sigma_c.a)$. As in the previous case, this means that $E_\varphi(\sigma.a) = \sigma_t \models \varphi$ and $\langle \text{Reach}(\sigma_t), \max_{\preceq}(\{w \preceq \sigma_d \mid w \in \Sigma_c^n\}), 1 \rangle \in W_0$. Thus $P(\sigma.a)$ holds.

Thus, if $\sigma \in \text{Pre}(\varphi)$, $P(\sigma.a)$ holds.

In all cases, $P(\sigma.a)$ holds. Thus, $P(\sigma) \implies P(\sigma.a)$.

By induction on σ , $\forall \sigma \in \Sigma^*$, $(\sigma \in \text{Pre}(\varphi) \wedge (\sigma_s, \sigma_b) = \text{store}_\varphi(\sigma)) \implies (E_\varphi(\sigma) \models \varphi \wedge \langle \text{Reach}(\sigma_s), \sigma_c, 1 \rangle \in W_0)$. In particular, for all $\sigma \in \Sigma^*$, $(\sigma \in \text{Pre}(\varphi)) \implies (E_\varphi(\sigma) \models \varphi)$. This means that E_φ is sound with respect to φ in $\text{Pre}(\varphi)$. \square

PROPOSITION 3.15. E_φ is compliant, as per Definition 3.3.

PROOF. For $\sigma \in \Sigma^*$, let $P(\sigma)$ be the predicate: “ $((\sigma_s, \sigma_c) = \text{store}_\varphi(\sigma)) \implies (\sigma_s|_{\Sigma_c}.\sigma_c = \sigma|_{\Sigma_c} \wedge \sigma_s|_{\Sigma_u} = \sigma|_{\Sigma_u})$ ”. Let us prove that for all $\sigma \in \Sigma^*$, $P(\sigma)$ holds.

- *Induction basis:* $\text{store}_\varphi(\epsilon) = (\epsilon, \epsilon)$, and $\epsilon|_{\Sigma_c} = \epsilon|_{\Sigma_c}.\epsilon$, and $\epsilon|_{\Sigma_u} = \epsilon|_{\Sigma_u}$. Thus $P(\epsilon)$ holds.
- *Induction step:* Let us suppose that for $\sigma \in \Sigma^*$, $P(\sigma)$ holds. Let us consider $(\sigma_s, \sigma_c) = \text{store}_\varphi(\sigma)$, $a \in \Sigma$, and $(\sigma_t, \sigma_d) = \text{store}_\varphi(\sigma.a)$. Let us prove that $P(\sigma.a)$ holds.

Case 1: $a \in \Sigma_u$. Then, $\sigma_t = \sigma_s.a.\sigma'_s$, where σ'_s is defined in Definition 3.10, and $\sigma_t.\sigma_d = \sigma_s.a.\sigma_c$. Therefore, $\sigma_t|_{\Sigma_c}.\sigma_d =$

$(\sigma_t.\sigma_d)|_{\Sigma_c}$, since $\sigma_d \in \Sigma_c^*$. Thus, $\sigma_t|_{\Sigma_c}.\sigma_d = \sigma_s|_{\Sigma_c}.\sigma_c$. Since $P(\sigma)$ holds, $\sigma_t|_{\Sigma_c}.\sigma_d = \sigma|_{\Sigma_c} = (\sigma.a)|_{\Sigma_c}$.

Moreover, since $\sigma'_s \in \Sigma_c^*$, $\sigma_t|_{\Sigma_u} = \sigma_s|_{\Sigma_u}.a$. Since $P(\sigma)$ holds, this means that $\sigma_t|_{\Sigma_u} = \sigma|_{\Sigma_u}.a = (\sigma.a)|_{\Sigma_u}$.

Thus $P(\sigma.a)$ holds.

Case 2: $a \in \Sigma_c$. Then $\sigma_t = \sigma_s.\sigma_s''$, where σ_s'' is defined in Definition 3.10, and $\sigma_t.\sigma_d = \sigma_s.\sigma_c.a$. Therefore, $\sigma_t|_{\Sigma_c}.\sigma_d = (\sigma_t.\sigma_d)|_{\Sigma_c} = (\sigma_s.\sigma_c.a)|_{\Sigma_c} = \sigma_s|_{\Sigma_c}.\sigma_c.a$. Since $P(\sigma)$ holds, this means that $\sigma_t|_{\Sigma_c}.\sigma_d = \sigma|_{\Sigma_c}.a = (\sigma.a)|_{\Sigma_c}$.

Moreover, since $\sigma_s'' \in \Sigma_c^*$, $\sigma_t|_{\Sigma_u} = \sigma_s|_{\Sigma_u}$. Since $P(\sigma)$ holds, this means that $\sigma_t|_{\Sigma_u} = \sigma|_{\Sigma_u} = (\sigma.a)|_{\Sigma_u}$.

Thus $P(\sigma.a)$ holds.

In both cases, $P(\sigma.a)$ holds.

Thus, for all $\sigma \in \Sigma^*$, for all $a \in \Sigma$, $P(\sigma) \implies P(\sigma.a)$.

By induction on σ , for all $\sigma \in \Sigma^*$, $((\sigma_s, \sigma_c) = \text{store}_\varphi(\sigma)) \implies (\sigma_s|_{\Sigma_c}.\sigma_c = \sigma|_{\Sigma_c} \wedge \sigma_s|_{\Sigma_u} = \sigma|_{\Sigma_u})$.

Moreover, if $\sigma \in \Sigma^*$, $u \in \Sigma_u$, $(\sigma_s, \sigma_c) = \text{store}_\varphi(\sigma)$, and $(\sigma_t, \sigma_d) = \text{store}_\varphi(\sigma.u)$, then $\sigma_t = \sigma_s.u.\sigma'_s$, where σ'_s is defined in Definition 3.10. Thus $\sigma_s.u \preceq \sigma_t$, and since $\sigma_s = E_\varphi(\sigma)$, and $\sigma_t = E_\varphi(\sigma.u)$, it follows that $E_\varphi(\sigma).u \preceq E_\varphi(\sigma.u)$.

Thus, for all $\sigma \in \Sigma^*$, $E_\varphi(\sigma)|_{\Sigma_c} \preceq \sigma|_{\Sigma_c} \wedge E_\varphi(\sigma)|_{\Sigma_u} = \sigma|_{\Sigma_u} \wedge \forall u \in \Sigma_u, E_\varphi(\sigma).u \preceq E_\varphi(\sigma.u)$, meaning that E_φ is compliant. \square

PROPOSITION 3.16. E_φ is optimal in $\text{Pre}(\varphi)$, as per Definition 3.4.

PROOF. Let E be an enforcement function such that $\text{compliant}(E, \Sigma_c, \Sigma_u)$ holds, and let us consider $\sigma \in \text{Pre}(\varphi)$ and $a \in \Sigma$ such that $E(\sigma) = E_\varphi(\sigma)$ and $|E(\sigma.a)| > |E_\varphi(\sigma.a)|$. Let us also consider $(\sigma_s, \sigma_c) = \text{store}_\varphi(\sigma)$. Let us show that there exists $\sigma_u \in \Sigma_u^*$ such that $E(\sigma.a.\sigma_u) \not\models \varphi$. We consider two cases:

Case 1: $a \in \Sigma_u$. Then, since E is compliant, and $E(\sigma) = E_\varphi(\sigma) = \sigma_s$, there exists $\sigma_{s1} \preceq \sigma_c$ such that $E(\sigma.a) = E(\sigma).a.\sigma_{s1} = \sigma_s.a.\sigma_{s1}$. Moreover, there exists $\sigma'_s \preceq \sigma_c$ such that $E_\varphi(\sigma.a) = E_\varphi(\sigma).a.\sigma'_s = \sigma_s.a.\sigma'_s$. Since $|E(\sigma.a)| > |E_\varphi(\sigma.a)|$, $|\sigma_{s1}| > |\sigma'_s|$. Considering that $\sigma'_s = \max_{\preceq}(G(\text{Reach}(\sigma_s.a), \sigma_c) \cup \{\epsilon\})$, it follows that $\sigma_{s1} \notin G(\text{Reach}(\sigma_s.a), \sigma_c)$. Following the definition of G , this means that either $\sigma_{s1} \not\preceq \sigma_c$; $\text{Reach}(\sigma_s.a)$ after $\sigma_{s1} \notin F$; or that $\langle \text{Reach}(\sigma_s.a) \text{ after } \sigma_{s1}, \sigma_{s1}^{-1}.\sigma_c, 1 \rangle \notin W_0$. Since E' is compliant, $\sigma_{s1} \preceq \sigma_c$, thus at least one of the two last conditions holds. If $\text{Reach}(\sigma_s.a)$ after $\sigma_{s1} = \text{Reach}(\sigma_s.a.\sigma_{s1}) = \text{Reach}(E(\sigma.a)) \notin F$, then $E(\sigma.a) \not\models \varphi$. Otherwise, $\langle \text{Reach}(\sigma_s.a) \text{ after } \sigma_{s1}, \sigma_{s1}^{-1}.\sigma_c, 1 \rangle \notin W_0$. Then, $\langle \text{Reach}(\sigma_s.a) \text{ after } \sigma_{s1}, \sigma_{s1}^{-1}.\sigma_c, 1 \rangle \in W_1$, meaning that P_1 has a winning strategy. Since receiving controllable events only helps P_0 to win, this means that there exists an uncontrollable event $u \in \Sigma_u$ such that $\langle \text{Reach}(\sigma_s.a.\sigma_{s1}) \text{ after } u, \sigma_{s1}^{-1}.\sigma_c, 0 \rangle \in W_1$. Then, since W_1 is the set of winning vertices for P_1 , $\langle \text{Reach}(E(\sigma.a.u)), E(\sigma.a.u)|_{\Sigma_c}^{-1}.\sigma_c, 1 \rangle \in W_1$. Then again, there exists an uncontrollable event u' such that the output of E after receiving it reaches a vertex in W_1 again. In the end, it is possible to reach a vertex that is not a Büchi state (i.e. in $F \times \Sigma_c^n \times \{0, 1\}$), and that is in W_1 . Thus, there exists $\sigma_u \in \Sigma_u^*$ such that $\text{Reach}(E(\sigma.a.\sigma_u)) \notin F$, meaning that $E(\sigma.a.\sigma_u) \not\models \varphi$.

Case 2: $a \in \Sigma_c$. The proof is the same as in the case where $a \in \Sigma_u$, by replacing occurrences of “ $\sigma_s.a$ ” by “ σ_s ”, and occurrences of “ σ_c ” by “ $\sigma_c.a$ ”.

Thus, if E is an enforcement function such that there exists $\sigma \in \text{Pre}(\varphi)$, and $a \in \Sigma$ such that $\text{compliant}(E, \Sigma_u, \Sigma_c)$, $E(\sigma) = E_\varphi(\sigma)$, and $|E(\sigma.a)| > |E_\varphi(\sigma.a)|$, then there exists $\sigma_u \in \Sigma_u^*$ such that $E(\sigma.a.\sigma_u) \not\models \varphi$.

This means that E_φ is optimal in $\text{Pre}(\varphi)$. \square

PROPOSITION 3.18. *The output of the enforcement monitor \mathcal{E} for input σ is $E_\varphi(\sigma)$.*

PROOF. Let us introduce some notation for this proof: for a word $w \in \Gamma^{\mathcal{E}^*}$, we note $\text{input}(w) = \Pi_1(w(1)).\Pi_1(w(2)) \dots \Pi_1(w(|w|))$, the word obtained by concatenating the first members (the inputs) of w . In a similar way, we note $\text{output}(w) = \Pi_3(w(1)) \dots \Pi_3(w(|w|))$, the word obtained by concatenating all the third members (outputs) of w . Since all configurations are not reachable from $c_0^{\mathcal{E}}$, for $w \in \Gamma^{\mathcal{E}^*}$, we note $\text{Reach}(w) = c$ whenever $c_0^{\mathcal{E}} \xrightarrow{w}_{\mathcal{E}} c$, and $\text{Reach}(w) = \perp$ if such a c does not exist. We also define the Rules function as follows:

$$\text{Rules} : \begin{cases} \Sigma^* \rightarrow \Gamma^{\mathcal{E}^*} \\ \sigma \mapsto \max_{\preceq} (\{w \in \Gamma^{\mathcal{E}^*} \mid \text{input}(w) = \sigma \wedge \text{Reach}(w) \neq \perp\}) \end{cases}$$

For a word $\sigma \in \Sigma^*$, $\text{Rules}(\sigma)$ is the trace of the longest valid run in \mathcal{E} , i.e. the sequence of all the rules that can be applied with input σ . We then extend the definition of output to words in Σ^* : for $\sigma \in \Sigma^*$, $\text{output}(\sigma) = \text{output}(\text{Rules}(\sigma))$. We also note ϵ the empty word of Σ^* , and $\epsilon^{\mathcal{E}}$ the empty word of $\Gamma^{\mathcal{E}^*}$.

For $\sigma \in \Sigma^*$, let $P(\sigma)$ be the predicate: “ $E_\varphi(\sigma) = \text{output}(\sigma) \wedge ((\sigma_s, \sigma_c) = \text{store}_\varphi(\sigma) \wedge \text{Reach}(\text{Rules}(\sigma)) = \langle q, \sigma_c^{\mathcal{E}} \rangle) \implies (q = \text{Reach}(\sigma_s) \wedge \sigma_c = \sigma_c^{\mathcal{E}})$ ”.

Let us prove by induction that for all $\sigma \in \Sigma^*$, $P(\sigma)$ holds.

- *Induction basis:* $E_\varphi(\epsilon) = \epsilon = \text{output}(\epsilon)$. Moreover, $\text{store}_\varphi(\epsilon) = (\epsilon, \epsilon)$, and $\text{Reach}(\epsilon^{\mathcal{E}}) = c_0^{\mathcal{E}}$. Therefore, as $c_0^{\mathcal{E}} = \langle q_0, \epsilon \rangle$, $P(\epsilon)$ holds, because $\text{Reach}(\epsilon) = q_0$.
- *Induction step:* Let us suppose now that for some $\sigma \in \Sigma^*$, $P(\sigma)$ holds. Let us consider $(\sigma_s, \sigma_c) = \text{store}_\varphi(\sigma)$, $q = \text{Reach}(\sigma_s)$, $a \in \Sigma$, and $(\sigma_t, \sigma_d) = \text{store}_\varphi(\sigma.a)$. Let us prove that $P(\sigma.a)$ holds.

Since $P(\sigma)$ holds, $\text{Reach}(\text{Rules}(\sigma)) = \langle q, \sigma_c \rangle$ and $\sigma_s = \text{output}(\sigma)$. We consider two cases:

Case 1: $a \in \Sigma_u$. Then, considering $\sigma'_s = (\sigma_s.a)^{-1}.\sigma_t$, $\sigma_t = \sigma_s.a.\sigma'_s$. Since $a \in \Sigma_u$, rule *pass-uncont* can be applied: let us consider $q' = q$ after a . Then, $\langle q, \sigma_c \rangle \xrightarrow{a/\text{pass-uncont}(a)/a}_{\mathcal{E}} \langle q', \sigma_c \rangle$. Then, if $\sigma'_s = \epsilon$, $G(q', \sigma_c) = \emptyset$ or $G(q', \sigma_c) = \{\epsilon\}$, meaning that no other rule can be applied, and thus $P(\sigma.a)$ would hold. Otherwise, $\sigma'_s \neq \epsilon$, and thus $\sigma'_s \in G(q', \sigma_c)$, meaning that $G(q', \sigma_c) \neq \emptyset$ and $G(q', \sigma_c) \neq \{\epsilon\}$, thus rule *dump*($\sigma_c(1)$) can be applied. Since $\sigma'_s \preceq \sigma_c$, $\sigma'_s(1) = \sigma_c(1)$, thus if $q_1 = q'$ after $\sigma_c(1)$, $q_1 = q'$ after $\sigma'_s(1)$. If $\sigma'_s(1)^{-1}.\sigma'_s \neq \epsilon$, then $\sigma'_s(1)^{-1}.\sigma'_s \in G(q_1, \sigma_c(1)^{-1}.\sigma_c)$, meaning that rule *dump* can be applied again. Rule *dump* can actually be applied $|\sigma'_s|$ times, since for all $w \preceq \sigma'_s$, if $w \neq \sigma'_s$, then $w^{-1}.\sigma'_s \neq \epsilon$ and $w^{-1}.\sigma'_s \in$

$G(q'$ after $w, w^{-1}.\sigma_c)$. Thus, after rule *dump* has been applied $|\sigma'_s|$ times, the configuration reached is $\langle q'$ after $\sigma'_s, \sigma'_s^{-1}.\sigma_c \rangle$. Moreover, the output produced by all the rules *dump* is σ'_s . Since no rule can be applied after the $|\sigma'_s|$ applications of the rule *dump*, $\text{output}(\sigma.a) = \text{output}(\sigma).a.\sigma'_s = \sigma_t$, and $\text{Reach}(\text{Rules}(\sigma.a)) = \langle q'$ after $\sigma'_s, \sigma'_s^{-1}.\sigma_c \rangle = \langle q$ after a after $\sigma'_s, \sigma_d \rangle = \langle \text{Reach}(\sigma_s)$ after a after $\sigma'_s, \sigma_d \rangle = \langle \text{Reach}(\sigma_s.a.\sigma'_s), \sigma_d \rangle = \langle \text{Reach}(\sigma_t), \sigma_d \rangle$. Thus, if $a \in \Sigma_u$, $P(\sigma.a)$ holds.

Case 2: $a \in \Sigma_c$. Then, considering $\sigma''_s = \sigma_s^{-1}.\sigma_t$, $\sigma_t = \sigma_s.\sigma''_s$. Since $a \in \Sigma_c$, it is possible to apply the *store-cont* rule, and $\langle q, \sigma_c \rangle$ after $a/\text{store-cont}(a)/\epsilon = \langle q, \sigma_c.a \rangle$. Then as in the case where $a \in \Sigma_u$, rule *dump* can be applied $|\sigma''_s|$ times, meaning that the configuration reached would then be $\langle q$ after $(\sigma_c.a)(1) \dots (\sigma_c.a)(|\sigma''_s|)$, $(\sigma_c.a)(|\sigma''_s| + 1) \dots (\sigma_c.a)(|\sigma''_s| + 2) \dots (\sigma_c.a)(|\sigma_c.a|)$. Since $\sigma''_s \preceq \sigma_c.a$, $(\sigma_c.a)(1) \dots (\sigma_c.a)(|\sigma''_s|) = \sigma''_s$, thus $\text{Reach}(\text{Rules}(\sigma.a)) = \langle q$ after $\sigma''_s, \sigma''_s^{-1}.\sigma_c \rangle = \langle \text{Reach}(\sigma_t), \sigma_d \rangle$. Moreover, $\text{output}(\sigma.a) = \text{output}(\sigma).\sigma''_s = \sigma_s.\sigma''_s = \sigma_t = E_\varphi(\sigma.a)$.

Thus, if $a \in \Sigma_c$, $P(\sigma.a)$ holds.

This means that $P(\sigma) \implies P(\sigma.a)$.

Thus, by induction on σ , for all $\sigma \in \Sigma^*$, $P(\sigma)$ holds. In particular, for all $\sigma \in \Sigma^*$, $E_\varphi(\sigma) = \text{output}(\sigma)$. \square

B PROOFS OF SECTION 4

B.1 Proofs of termination and complexity

Algorithm 1 terminates because the array `reachable` contains only 0 and 1, and in `compute $\epsilon\Sigma_c^n\text{Rec}$` , all the subsequent arrays contain more 1 than the previous ones. Since `reachable` is finite, there exists only a finite number of arrays that can be considered in this way, and thus the array stabilises at some point. Since Σ_c is finite, the algorithm terminates.

Let us now compute the worst-case complexity of Algorithm 1 in terms of assignments. The loop starting in line 4 makes $2n$ assignments, and the one starting in line 8 makes $|\Sigma_c|$ assignments, and calls $|\Sigma_c|$ times function `compute $\epsilon\Sigma_c^n\text{Rec}$` . Let us compute the worst-case complexity of function `compute $\epsilon\Sigma_c^n\text{Rec}$` . The loop starting in line 5 makes $2n$ assignments, since it is possible to compute j in constant time (the automaton is deterministic and complete). The assignment in line 4 makes n^2 assignments, and there is another assignment in line 10. Thus, the loop starting in line 3 makes $|\Sigma_c| * (n^2 + 2n + 1)$ assignments, which is the complexity of the function. Now, since the first parameter of all the recursive calls is always the same, and the complexity depends only on it, the number of calls made to `compute $\epsilon\Sigma_c^n\text{Rec}$` only depends on the depth of the call stack. The algorithm terminates when `reachable` stabilises, which must occur in at most $n^2 - n$ calls, since this is the number of 0s present in `reachable` at the beginning. Thus, the depth of the call stack is bounded by $n^2 - n$. Thus, the worst-case complexity of Algorithm 1 is at most $2n + |\Sigma_c|^{n^2 - n}$. Note that this complexity is computed considering that initialisations have a complexity of 0.

The complexities of the other algorithms are straightforward.