Real-time Policy Enforcement with Metric First-Order Temporal Logic

François Hublet, David Basin, and Srđan Krstić

Institute of Information Security, Department of Computer Science, ETH Zürich, Zurich, Switzerland {francois.hublet,basin,srdan.krstic}@inf.ethz.ch

Abstract. Correctness and regulatory compliance of today's software systems are crucial for our safety and security. This can be achieved with policy enforcement: the process of monitoring and possibly modifying system behavior to satisfy a given policy. The enforcer's capabilities determine which policies are enforceable. We study the enforceability of policies specified in metric first-order temporal logic (MFOTL) with enforcers that can cause and suppress different system actions in real time. We consider an expressive safety fragment of MFOTL and show that a policy from that fragment is enforceable if and only if it is equivalent to a policy in a simpler, syntactically defined MFOTL fragment. We then propose an enforcement algorithm for all monitorable policies from the latter fragment, and show that our EnfPoly enforcer outperforms state-of-the-art tools.

1 Introduction

Modern software systems are increasingly complex, ubiquitous and intransparent. In this context, allowing individuals to scrutinize and control the systems that affect their daily lives is an important technical and societal challenge. To achieve this goal it is crucial to develop systems that can *monitor* and *control* other target systems, by enforcing *policies* that describe the acceptable target system's behaviors.

Policy enforcement [54], depicted in Figure 1, is a form of execution monitoring where a system, called an *enforcer*, observes a target system's actions, *detects* attempted policy violations, and reacts to *prevent* them. In contrast, policy monitoring (or runtime verification) [7, 26] provides *monitors* that only passively detect policy violations by the target system. Both problems have offline and online variants: the former considers a trace of recorded target system actions, while the latter observes the target system in real time.

Policy enforcement has been studied in different communities (Section 2) like controller synthesis [1, 46, 47], security [4, 54], and operating systems [51, 52], each defining and solving the problem in a different, specialized context.

Schneider [54] studied the general form of the policy enforcement problem in the context of security. He proposed security automata as enforcers that, when composed

Fig. 1: Policy enforcement

An enforcer E observes actions in a target system S and reacts (e.g., causes or suppresses some actions in S) to ensure policy compliance. S interacts with an environment X, which E cannot control.

with the target system, prevent policy violations by simply terminating it. Schneider, and later others [11], identified classes of policies that were *enforceable* using such an enforcer. As policy enforceability depends on the enforcer's powers over the target system (e.g., its ability to suppress, cause, or delay system's actions), other enforceable policy classes have been suggested [9].

Automata and temporal logic are popular formalisms for policy specification. Existing security policy enforcers typically focus on propositional policies expressed as variants of (security, edit, or timed) automata [27]. In contrast, controller synthesis tools [6] also enforce specifications expressed in LTL [16] or in (fragments of) metric temporal logics [17, 19, 32, 40]. However, automata and propositional temporal logic are limited in their expressiveness: they regard system actions as atomic and thus cannot formulate dependencies between the data values coming from an infinite domain that the actions may carry as parameters. For instance, a data value may contain personally identifiable information, and then each system action that uses the value should be preceded by an action that receives a consent for the value's particular use [5]. To the best of our knowledge, there is no tool that supports enforcement for first-order logic specifications.

In this paper, we consider the online policy enforcement (Section 3) of policies expressed in metric first-order temporal logic (MFOTL) [20], which extends LTL with metric constraints and first-order quantification (Section 4). To enforce MFOTL policies, our enforcer can observe the target system in real time, actively cause or suppress different types of actions, and only observe other actions of the target system. As enforcer must react in real time, policies must be such that their satisfaction does not depend on future information. All actions, caused either by the enforcer or the target system, are instantaneous and tagged with a timestamp.

We therefore consider two "well-behaved" fragments of MFOTL: (1) we study enforceability of $\text{MFOTL}_{\square}^{\mathcal{F}}$, a safety fragment of MFOTL comprising closed formulae of the form $\square\varphi$ ("always φ") where φ's satisfaction does not depend on future information; and (2) we design an efficient enforcement algorithm for *monitorable* and enforceable $\text{MFOTL}_{\square}^{\mathcal{F}}$ formulae. Violations of monitorable formulae [12] can be detected by manipulating only finite sets of satisfying valuations. As these sets are always finite, we can use simple, yet efficient, data structures and reuse the existing, highly-optimized monitoring algorithm for monitorable MFOTL formulae [13].

Overall, we characterize the enforceability of $\text{MFOTL}_{\square}^{\mathcal{F}}$ formulae by an enforcer with the ability to suppress or cause different system actions, and propose and implement an enforcer for monitorable $\text{MFOTL}_{\square}^{\mathcal{F}}$ formulae. We make the following contributions:

– For an enforcer with the ability to suppress or cause (disjoint sets of) actions, we characterize enforceability of $\text{MFOTL}_{\square}^{\mathcal{F}}$ formulae. We show that it is undecidable whether an $\text{MFOTL}_{\square}^{\mathcal{F}}$ formula is enforceable and propose an expressively complete syntactical approximation (Section 5).
– We develop an enforcement algorithm for monitorable $\text{MFOTL}_{\square}^{\mathcal{F}}$ formulae and prove its correctness (Section 6).
– Finally, we describe our enforcer's implementation (Section 7) and evaluate its time and memory usage against other state-of-the-art tools (Section 8).

The proofs of all lemmas and theorems can be found in our extended report [35].

2 Related Work

We group related work into conceptual approaches and those that implement enforcers.

Theory. The policy enforcement problem was studied by Schneider and Erlingsson in the context of security [24, 54]. Schneider defined security automata, a class of Büchi automata, as enforcers. Violations were prevented by terminating the system. Bauer *et al.* [14] extended Schneider's work by considering enforcers that can suppress and cause events. Basin *et al.* [11] distinguished between suppressable and (only-)observable events and refined Schneider's enforceability accordingly, but only discussed enforcement via suppression. Falcone *et al.* [25] later studied the enforcement of propositional timed policies by suppressing and delaying events. Recently, Aceto *et al.* [2] proposed *bidirectional* enforcers that treat input and output system actions differently. We see this distinction as a more refined event type partition (Section 3).

Policy enforcement is closely related to the controller synthesis problem [46], where a controller ($\approx$ enforcer) wants to ensure compliance of a plant ($\approx$ system) with a specification ($\approx$ policy). Specification *realizability* corresponds to enforceability, while controller synthesis (i.e., generating an automaton from a specification) corresponds to generating an enforcement algorithm tailored to a particular specification. Our enforcer does not explicitly generate code for a specific policy, but rather takes the policy directly as input. Early work by Pnueli, Rosner and Abadi [1, 46, 47] studied LTL realizability and LTL (controller) synthesis. More efficient approaches later emerged [29, 39, 53], as well as techniques for timed automata [6] and metric extensions of LTL [17, 19, 32, 40].

Tools. Policy enforcement approaches typically rely on different classes of automata both as enforcers and as policies [11, 14, 21, 24, 25, 27, 28, 41, 42, 45, 49, 50, 54]. A recent survey [26] listed three enforcement tools: GREP [49], Proactive Libraries [50], and TiPEX [45]. Both GREP and TiPEX use timed automata as a specification language, and could thus support propositional temporal logics like MITL [3] via conversion to timed automata [18, 43]. They do not, however, natively support temporal logic.

The state-of-the-art MonPoly tool [13] can detect violations of monitorable MFOTL policies [12]. Other tools for first-order temporal logics include Verimon [8, 38, 55] and DejaVu [30, 31], none of which supports enforcement to prevent violations.

Many controller synthesis tools have been developed for LTL like Lily [36], Unbeast [23], Acacia+ [16] and SSyft [56]. Other tools synthesize controllers for systems described by timed automata to comply with specifications written in TCTL [15, 44], MTL [32], or its fragment $MTL_{0,\infty}$ [40]. BluSTL [22, 48] is a MATLAB toolbox for generating controllers from signal temporal logic (STL) specifications. None of these tools supports first-order logic.

3 Policy Enforcement

We fix a signature $\Sigma = (\mathbb{D}, \mathbb{E}, a)$, containing an infinite set $\mathbb{D}$ of constant symbols, a finite set of *event names* $\mathbb{E}$, and an arity function $a : \mathbb{E} \to \mathbb{N}$. An *event* is a pair $(e, (d_1, \ldots, d_{a(e)})) \in \mathbb{E} \times \mathbb{D}^{a(e)}$ of an event name e and $a(e)$ arguments.

Events model system actions *observable* by the enforcer. While some of these observable events can also be controlled (i.e., suppressed or caused) by the enforcer, others can *only* be observed. To capture these different cases, we partition $\mathbb{E}$ into two

sets: a set of controllable event names, and a set of only-observable event names. Among the controllable event names, we further distinguish between *suppressable* event names $\mathsf{Sup} \subseteq \mathbb{E}$ and *causable* event names $\mathsf{Cau} \subseteq \mathbb{E}$. The set of only-observable event names is $\mathsf{Obs} = (\mathbb{E} \setminus \mathsf{Sup}) \setminus \mathsf{Cau}$. In general, some controllable events might be both suppressable and causable. However, we will assume that no such events exist, i.e. $\mathsf{Sup} \cap \mathsf{Cau} = \emptyset$. Our reason for this will become apparent when we consider MFOTL policy enforcement (Section 6), and we will discuss ways in which this assumption can be relaxed.

Example 1. As a running example, consider the signature $(\mathbb{N}, \{\mathtt{Open}, \mathtt{Close}, \mathtt{Knock}\}, a)$, where $a(\cdot) = 1$, $\mathsf{Sup} = \{\mathtt{Open}\}$, and $\mathsf{Cau} = \{\mathtt{Close}\}$. The target system controls a set of doors indexed by integers, which an enforcer can mechanically close or keep closed, but not hold open. Each door i is equipped with a sensor that causes a $\mathtt{Knock}(i)$ event whenever a human knocks on the door. $\mathtt{Knock}$ events are only-observable ($\mathsf{Obs} = \{\mathtt{Knock}\}$), since they reflect the environment's behavior.

Given a signature Σ, we define the set of *(event) databases* $\mathbb{DB}^*$ as $2^{\{(e,d) \,|\, e \in \mathbb{E}, \, d \in \mathbb{D}^{a(e)}\}}$. Databases represent structures over Σ. We restrict ourselves to considering *automatic* databases, i.e., databases that can be represented by a collection of finite automata [37]. This setup is the most general one used for MFOTL monitoring in [12].

Definition 1 (Automatic event database). *An event database D is* automatic *iff for all $e \in \mathbb{E}$, $D \cap \{(e,d) \,|\, d \in \mathbb{D}^{a(e)}\}$ is a regular set. $\mathbb{DB}$ is the set of automatic event databases.*

Finally, for any $E \subseteq \mathbb{E}$, we denote by $\mathsf{Ev}(E)$ the set of all databases with event names in E only, i.e. $\mathsf{Ev}(E) := \{D \in \mathbb{DB} \mid \forall (e, (d_1, \ldots, d_{a(e)})) \in D. \ e \in E\}$.

Traces are finite or infinite sequences $\sigma = (\tau_i, D_i)_{1 \leq i \leq k}$, $k \in \mathbb{N} \cup \{\infty\}$, where $\tau_i \in \mathbb{N}$ are nondecreasing timestamps, and $D_i \in \mathbb{DB}$ are databases. The smallest timestamp of a trace σ is denoted by $\mathsf{sts}(\sigma) = \tau_1 \in \mathbb{N}$, its largest timestamp is denoted by $\mathsf{lts}(\sigma) = \sup_{1 \leq i \leq k} \tau_i \in \mathbb{N} \cup \{\infty\}$. The empty trace is denoted by ε, the set of traces by $\mathbb{T}$, and the set of finite traces by $\mathbb{T}_f = \{\sigma \in \mathbb{T} \mid |\sigma| < \infty\}$. If σ, σ' are two traces such that σ is finite, $\sigma \cdot \sigma'$ denotes the concatenation of σ and σ'. A *(trace) property* is a subset $P \subseteq \mathbb{T}$. For all $\sigma, \sigma' \in \mathbb{T}$, we write $\sigma \preceq \sigma'$ iff σ is a prefix of σ', and denote by $\mathsf{pre}(\sigma)$ the set of all prefixes of σ. The *limit closure* of a set $A \subseteq \mathbb{T}$, denoted by $\mathsf{cl}(A)$, contains all traces whose finite prefixes are all in A, i.e., $\mathsf{cl}(A) = \{\sigma \in \mathbb{T} \mid \forall \sigma' \in \mathsf{pre}(\sigma). \ |\sigma'| < \infty \Rightarrow \sigma' \in A\}$. The *truncation* of A is $\mathsf{trunc}(A) = \{\sigma \in A \mid \mathsf{pre}(\sigma) \subseteq A\}$, the largest prefix-closed subset of A.

Finite databases $\mathbb{DB}^\dagger \subseteq \mathbb{DB}$ are a specific type of automatic databases. We also consider traces with finite databases $\mathbb{T}^\dagger \subseteq \mathbb{T}$, and finite traces with finite databases $\mathbb{T}_f^\dagger$.

We now extend the definition of enforceability [11] to support causable events.

Definition 2 (Enforceability). *A property $P \subseteq \mathbb{T}$ is* enforceable *iff there is a deterministic Turing machine (TM) $\mathcal{M}$ accepting a set of finite traces S such that*
  (i) $\mathsf{cl}(\mathsf{trunc}(S)) = P$;
 (ii) *$\mathcal{M}$ accepts ε;*
(iii) *For all $\sigma \in \mathsf{trunc}(S)$, $\tau \geq \mathsf{lts}(\sigma)$, and $D \in \mathbb{DB}$, $\mathcal{M}$ halts on $\sigma \cdot ((\tau, D))$; and*
(iv) *For all $\sigma \in \mathsf{trunc}(S)$, $\tau \geq \mathsf{lts}(\sigma)$, and $D \in \mathbb{DB}$, there exists $S \in \mathsf{Ev}(\mathsf{Sup})$ and $C \in \mathsf{Ev}(\mathsf{Cau})$ such that $\mathcal{M}$ accepts $\sigma \cdot ((\tau, (D \setminus S) \cup C))$.*

Properties are sets of infinite traces, while enforcers (that do not know the system's implementation) can only observe finite traces. Hence, an enforceable property must be checked "prefix-wise": a trace is in a property iff an enforcer accepts all of its prefixes. Enforceable properties must hold on the empty trace, i.e., the system must initially comply with the property. For any extension of a (non-violating) prefix, the enforcer must be able to decide on its compliance to the property. Whenever a valid prefix is extended with an additional database, there must exist sets of suppressable and causable events which the enforcer can respectively suppress and cause to ensure satisfaction of the property.

Our notion of enforceability implies safety:

Lemma 1. *Any enforceable property $P \subseteq \mathbb{T}$ is a safety property.*

The converse is not true: a safety property that requires that no `Knock` event ever happens is not enforceable, as `Knock` events are only-observable and cannot be suppressed.

An *enforcer* can be seen as a Turing machine that, given a finite trace, returns a pair of sets of events to be respectively suppressed and caused in the last database of the trace, with the additional requirement that events to be suppressed (resp. caused) should be suppressable (resp. causable) and present (resp. not already present) in this database.

Definition 3 (Enforcer). *An* enforcer *is a computable function $\mu : \mathbb{T}_f \to \mathbb{DB} \times \mathbb{DB}$ such that for all $\sigma \in \mathbb{T}_f$, $\tau \geq \mathsf{lts}(\sigma)$, $D \in \mathbb{DB}$, and $(B,C) = \mu(\sigma \cdot ((\tau, D)))$:*
  (i) For all $(e,d) \in B$, $e \in \mathsf{Sup}$ and $(e,d) \in D$; and
  (ii) For all $(e,d) \in C$, $e \in \mathsf{Cau}$ and $(e,d) \notin D$.

An enforcer μ is correct with respect to a property P if, for all $\sigma \in P$, any trace σ' obtained by adding a single database at the end of σ and then updating it (to some σ'') according to μ ensures $\sigma'' \in P$.

Definition 4 (Correct enforcement). *An enforcer μ is called* correct *with respect to a property $P \subseteq \mathbb{T}$ and a set of databases $\Delta \subseteq \mathbb{DB}$ if for all $\sigma \in P \cap \mathbb{T}_f$, $\tau \geq \mathsf{lts}(\sigma)$, $D \in \Delta$, and $(B,C) = \mu(\sigma \cdot ((\tau, D)))$, we have $\sigma \cdot ((\tau, (D \setminus B) \cup C)) \in P$.*

Transparent enforcers [10] do not to alter traces that belong to the enforced property:

Definition 5 (Transparent enforcement). *An enforcer μ is called* transparent *with respect to a property $P \subseteq \mathbb{T}$ and a set of databases $\Delta \subseteq \mathbb{DB}$ if for all $\sigma \in P \cap \mathbb{T}_f$, $\tau \geq \mathsf{lts}(\sigma)$, $D \in \Delta$, we have $\sigma \cdot ((\tau, D)) \in P \Longrightarrow \mu(\sigma \cdot ((\tau, D))) = (\emptyset, \emptyset)$.*

Given $A \subseteq \mathbb{T}$ and $B, C \subseteq \mathbb{E}$, $\mathsf{extend}(A, B, C)$ is the set of all traces $\sigma \cdot (\tau, D)$ obtained by appending to any trace $\sigma \in A$ the pair $(\tau, D \cup D')$ with $\tau \geq \mathsf{lts}(\sigma)$, $D \in 2^{B \times \mathbb{D}^*}$ and $D' = \{(c,d) \mid c \in C, d \in \mathbb{D}^{a(c)}\}$. Intuitively, set $\mathsf{extend}(A, B, C)$ is obtained from the set A by appending *some* events from B and *all* events from C to A. We have:

Lemma 2. *Let $P \subseteq \mathbb{T}$ such that P is enforceable. Then there exists a correct and transparent enforcer with respect to P and $\mathbb{DB}$.*

$$v,i \models_\sigma r(t_1,...,t_n) \text{ iff } (r,(v(t_1),...,v(t_n))) \in D_i \qquad\qquad\quad v,i \models_\sigma \neg\varphi \quad \text{ iff } v,i \not\models_\sigma \varphi$$
$$v,i \models_\sigma \exists x.\, \varphi \qquad \text{iff } v[x \mapsto d],i \models_\sigma \varphi \text{ for } d \in \mathbb{D} \qquad v,i \models_\sigma \varphi \vee \psi \text{ iff } v,i \models_\sigma \varphi \text{ or } v,i \models_\sigma \psi$$
$$v,i \models_\sigma \bullet_I \varphi \qquad \text{iff } i > 1 \text{ and } v,i-1 \models_\sigma \varphi \text{ and } \tau_i - \tau_{i-1} \in I \quad | \quad v,i \models_\varepsilon \varphi$$
$$v,i \models_\sigma \bigcirc_I \varphi \qquad \text{iff } i+1 \leq |\sigma| \text{ and } v,i+1 \models_\sigma \varphi, \text{ and } \tau_{i+1} - \tau_i \in I$$
$$v,i \models_\sigma \varphi\, \mathsf{S}_I\, \psi \qquad \text{iff } v,j \models_\sigma \psi \text{ for some } j \leq i, \tau_i - \tau_j \in I, \text{ and } v,k \models_\sigma \varphi \text{ for all } k,\, j < k \leq i$$
$$v,i \models_\sigma \varphi\, \mathsf{U}_I\, \psi \qquad \text{iff } v,j \models_\sigma \psi \text{ for some } |\sigma| \geq j \geq i, \tau_j - \tau_i \in I, \text{ and } v,k \models_\sigma \varphi \text{ for all } k,\, j > k \geq i$$

Fig. 2: MFOTL semantics

4   Metric First-Order Temporal Logic

Metric first-order temporal logic (MFOTL) extends first-order logic with the metric temporal operators "previous" ($\bullet_I$), "next" ($\bigcirc_I$), "since" (S_I), and "until" (U_I). We write $\mathbb{I}$ for the set of intervals over $\mathbb{N}$ and $\mathbb{V}$ for a countable set of variables. MFOTL formulae over a signature Σ are defined by the grammar

$$\varphi ::= r(t_1,\ldots,t_{a(r)}) \mid \neg\varphi \mid \varphi \vee \varphi \mid \exists x.\, \varphi \mid \bullet_I \varphi \mid \bigcirc_I \varphi \mid \varphi\, \mathsf{S}_I\, \varphi \mid \varphi\, \mathsf{U}_I\, \varphi,$$

where $t_1,\ldots,t_{a(r)} \in \mathbb{V} \cup \mathbb{D}$, $r \in \mathbb{E}$, and $I \in \mathbb{I}$. We define shorthands $\top := p \vee \neg p$, $\bot := \neg\top$, $\varphi \Rightarrow \psi := \neg\varphi \vee \psi$, and the operators "once" ($\blacklozenge_I \varphi := \top\, \mathsf{S}_I\, \varphi$), "eventually" ($\Diamond_I \varphi := \top\, \mathsf{U}_I\, \varphi$), "always" ($\Box_I \varphi := \neg \Diamond_I \neg\varphi$), and "historically" ($\blacksquare_I \varphi := \neg \blacklozenge_I \neg\varphi$). Temporal operators with no interval have $[0,\infty)$ instead. Predicates are formulae of the form $r(t_1,\ldots,t_{a(r)})$.

We extend the domain of *valuation* $v : \mathbb{V} \to \mathbb{D}$ to $\mathbb{D}$ by setting $v(d) = d$ for all $d \in \mathbb{D}$. We write $v[x \mapsto d]$ for the mapping equal to v, except that $v(x)$ is d. We use $\mathsf{fv}(\varphi)$ for the set of φ's free variables. For $k \in \mathbb{N}$, a trace $\sigma = ((\tau_i, D_i))_{1 \leq i \leq k}$, a timepoint $1 \leq i \leq |\sigma|$, a valuation v, and a formula φ, satisfaction relation $\models$ is defined in Figure 2. Note that $\models$ is well-defined for both finite and infinite traces. We write $v \models_\sigma \varphi$ for $v,1 \models_\sigma \varphi$.

We say that two MFOTL formulae φ and ψ are *equivalent*, written $\varphi \equiv \psi$, iff for all $v, \sigma \in \mathbb{T}$, $1 \leq i \leq |\sigma|$, we have $v,i \models_\sigma \varphi \Leftrightarrow v,i \models_\sigma \psi$.

If φ is closed, i.e., $\mathsf{fv}(\varphi) = \emptyset$, φ's satisfaction does not depend on v. We then write $\models_\sigma \varphi$ as shorthand for $\forall v.\, v \models_\sigma \varphi$. Given a closed formula φ, we denote by $\mathcal{L}(\varphi) \subseteq \mathbb{T}$ the set of all traces that satisfy φ, i.e., $\mathcal{L}(\varphi) := \{\sigma \in \mathbb{T} \mid \models_\sigma \varphi\}$. Finally, we denote by $\mathcal{L}_f(\varphi)$ the set of finite traces in $\mathcal{L}(\varphi)$, i.e., $\mathcal{L}_f(\varphi) = \{\sigma \in \mathcal{L}(\varphi) \mid |\sigma| < \infty\}$. Extending the previous terminology, we say that a formula φ is enforceable iff $\mathcal{L}(\varphi)$ is enforceable.

If the truth value of a formula only depends on the trace content in the past or present, an enforcer can compute satisfactions for each trace prefix, and react timely.

Definition 6 (Future-free formulae). *An MFOTL formula φ is called* future-free *iff for all $\sigma \in \mathbb{T}$, valuation v, and $\sigma' \preceq \sigma$ such that $|\sigma'| = i$, we have $v,i \models_\sigma \varphi \Leftrightarrow v,i \models_{\sigma'} \varphi$.*

For instance, formulae without future operators (U_I, $\bigcirc_I$, $\Diamond_I$, $\Box_I$) are future-free, but also some that have these operators nested in appropriate past operators.

Example 2. The formula $\varphi_1 = \blacklozenge_{[3,4]}(\exists x.\, \mathtt{Close}(x))$ uses no future temporal operators, and is therefore future-free. The formula $\varphi_2 = \blacklozenge_{[3,4]}(\exists x.\, \mathtt{Close}(x) \wedge \Diamond_{[1,2]} \mathtt{Open}(x))$ contains a future operator, but is still future-free, since the future operator $\Diamond_{[1,2]}$ (looking at most 2 time units into the future) is nested in a $\blacklozenge_{[3,4]}$ operator that is always evaluated at least 3 time units in the past. The formula $\varphi_3 = \Diamond_{[1,2]} \mathtt{Open}(x)$ is not future-free: its truth value depends on events happening up to 2 time units in the future.

In the rest of this paper, we consider the fragment MFOTL$_\square^{\mathcal{F}}$ that contains all closed formulae of the form $\square\varphi$, where φ is future-free. Given the correctness of the monitoring algorithm [12] for MFOTL formulae of the form $\square\varphi$, where all future operators in φ have bounded intervals and the fact that future-free formulae are a subset of the algorithm's supported formulae, we have:

Lemma 3. *For any $\varphi \in$ MFOTL$_\square^{\mathcal{F}}$, there exists a TM that decides $\mathcal{L}_f(\varphi)$.*

In fact, the algorithm determines without delay whether a future-free formula is satisfied.

5   MFOTL Enforceability

In this section, we characterize the enforceability of MFOTL$_\square^{\mathcal{F}}$ formulae with an enforcer as described in Section 3. Our first result is negative: a reduction, presented in our extended report [35], shows that the enforceability of MFOTL$_\square^{\mathcal{F}}$ formulae in undecidable.

Theorem 1. *Assume that* Sup *contains at least one event of arity at least 2 and* Obs $\neq \emptyset$. *The set* $\mathcal{E} = \{\varphi \in$ MFOTL$_\square^{\mathcal{F}} \mid \varphi$ *is enforceable*$\}$ *is not computable.*

The proof relies on the undecidability of universal validity in FOL. Therefore, it is sensible to ask whether some syntactical characterization of enforceability can be recovered by reasoning *modulo equivalence of formulae*. Is there a decidable and enforceable fragment of MFOTL$_\square^{\mathcal{F}}$ that contains all enforceable policies modulo equivalence? If so, such a fragment would not only provide a sound approximation of enforceable MFOTL$_\square^{\mathcal{F}}$ formulae, but also an approximation that is *expressively complete*. All enforceable MFOTL$_\square^{\mathcal{F}}$ policies could be expressed using the fragment via an appropriate (manual) rewriting. Rather surprisingly, such a fragment exists. Consider the following:

Definition 7 (GMFOTL). Guarded MFOTL *(GMFOTL) is defined inductively by:*

$$\psi ::= \bot \mid s(t_1,\ldots,t_n) \mid \neg c(t_1,\ldots,t_n) \mid \psi \wedge \varphi \mid \psi \vee \psi \mid \exists x.\, \psi$$

where $s \in$ Sup, $c \in$ Cau, and φ is an MFOTL *formula.*

In GMFOTL, all subformulae (and, in particular, all temporal subformulae) are *guarded* by an instance of a predicate $r(t_1,\ldots,t_n)$ with r being suppressable, or by an instance of a negated predicate $\neg r(t_1 \ldots, t_n)$ with r being causable. In the following, we call such a (possibily negated) predicate a *guard*. The presence of a guard ensures that, when an GMFOTL formula is satisfied with respect to a trace prefix, it can always be made false by suppressing or causing appropriate events in the last database of the prefix.

Example 3. Consider the formula $\varphi_4 = \neg\mathtt{Close}(x) \wedge \psi$, with an arbitrary future-free formula ψ and $\mathsf{fv}(\psi) = \{x\}$. For φ_4 to be satisfied with respect to a trace prefix σ, it must hold for some valuation of x and $\{(\mathtt{Close}, (a)) \mid v, |\sigma| \models_\sigma \psi, v(x) = a\}$ must not be in the last database of σ. Hence, φ_4 can be falsified by causing the appropriate $\mathtt{Close}$ events.

It can be shown that all closed formulae of the form $\square\neg\psi$ with $\psi \in$ GMFOTL and future-free are enforceable. Since enforceability is defined in terms of the language recognized by a given formula, we obtain that all MFOTL$_\square^{\mathcal{F}}$ formulae equivalent to some $\square\neg\psi$, with $\psi \in$ GMFOTL closed and future-free, are enforceable. In fact, the converse is also true: all future-free MFOTL$_\square^{\mathcal{F}}$ formulae are equivalent to a formula of the above form. We have thus obtained an expressively complete fragment of enforceable MFOTL$_\square^{\mathcal{F}}$. Formally:

Theorem 2. *A formula $\Box\varphi \in \text{MFOTL}_{\Box}^{\mathcal{F}}$ is enforceable iff there exists $\psi \in \text{GMFOTL}$ such that $\Box\varphi \equiv \Box\neg\psi$.*

Example 4. Consider the formula $\varphi_5 = \Box\forall x.\, (\text{Open}(x) \Rightarrow \neg\blacklozenge_{[2,5]}\,\text{Open}(x))$. This formula is enforceable: Open events that lead to a violation (i.e., those occurring 2 to 5 time units after a previous Open event with the same argument) can always be suppressed. The formula φ_5 is equivalent to $\Box\neg\psi$ where $\psi \in \text{GMFOTL}$ is

$$(\exists x.\, \text{Open}(x)) \wedge \neg(\forall x.\, (\text{Open}(x) \Rightarrow \neg\blacklozenge_{[2,5]}\,\text{Open}(x))).$$

6   MFOTL Enforcement in the Finite Case

In the previous section, we have presented GMFOTL, a syntactic class of MFOTL that is expressively complete for enforceable $\text{MFOTL}_{\Box}^{\mathcal{F}}$ formulae. Lemma 2 implies the existence of an enforcer for such formulae. However, the naive enforcer constructed in the lemma's proof may be inefficient—in fact, it may cause an infinite number of new events.

In this section, we focus on traces with finite databases and MFOTL formulae from the intersection of enforceable $\text{MFOTL}_{\Box}^{\mathcal{F}}$ formulae with monitorable MFOTL formulae [12]. We show that, in this case, we can exhibit a correct and transparent enforcer that produces only a finite number of events to be suppressed or caused.

6.1   Monitoring MFOTL Formulae

Basin *et al.* [12] describe an algorithm that efficiently monitors monitorable MFOTL formulae. Variants of this algorithm and the fragment it supports are used in several state-of-the-art tools [13, 55]. We now briefly recall the algorithm and some of its properties.

The algorithm encodes each database $D \in \mathbb{DB}^{\dagger}$ as a finite set of tables, one for each event name in the database. The row d is in the table corresponding to the event name e if $(e, d) \in D$. The set of satisfying valuations of a formula can similarly be encoded as a table whose rows represent valuations restricted to the domain of the formula's free variables.

The algorithm computes the table of satisfying valuations for a monitorable MFOTL formula bottom-up, using well-known table operations like join, anti-join, union, and projection. The syntactic monitorable fragment ensures that table operations always produce finite tables. In the rest of the section, we assume that this algorithm is available as a subroutine $\text{SAT}(\varphi, \sigma) = \{v \mid v, |\sigma| \models_{\sigma} \varphi\}$ that returns the set of satisfying valuations of a monitorable MFOTL formula φ with respect to finite trace $\sigma \in \mathbb{T}^{\dagger}$ and timepoint $|\sigma|$.

The monitorable MFOTL fragment [55] also ensures that for any valuation v satisfying a formula φ from the fragment with respect to a finite trace σ and a time point $1 \leq i \leq |\sigma|$, for every $x \in \text{fv}(\varphi)$ the value $v(x) \in \mathbb{D}$ is contained in some event argument in a database in σ or a constant term in φ. Formally:

Lemma 4. *For all monitorable $\varphi \in \text{MFOTL}$, valuation v, trace $\sigma \in \mathbb{T}^{\dagger}$, and timepoint $1 \leq i \leq |\sigma|$, assuming $v, i \models_{\sigma} \varphi$, we have*

$$\forall x \in \text{fv}(\varphi).\ \exists 1 \leq j \leq |\sigma|.\ (e, d) \in D_j, 1 \leq k \leq a(e).\ d_k = v(x) \vee d_k \in \text{cst}(\varphi)$$

where $\text{cst}(\varphi) \subset \mathbb{D}$ denotes the (finite) set of constant terms that appear in φ.

We will use this lemma, as well as the termination of the subroutine SAT [12], to prove the termination of our enforcer.

Algorithm 1 Function enf

---

function enf(φ, σ, v)
    if $\varphi = r(t_1, \ldots, t_n), r \in \mathsf{Sup}$ **then**
        return $(\{(r, (v(t_1), \ldots, v(t_n)))\}, \emptyset)$
    else if $\varphi = \neg r(t_1, \ldots, t_n), r \in \mathsf{Cau}$ **then**
        return $(\emptyset, \{(r, (v(t_1), \ldots, v(t_n)))\})$
    else if $\varphi = \varphi_1 \wedge \varphi_2$ **then**
        return enf(φ_1, σ, v)
    else if $\varphi = \varphi_1 \vee \varphi_2$ **then**
        return FIXPOINT$(\sigma, \mathsf{enf}_{\mathsf{or}, \varphi_1, \varphi_2, v})$
    else if $\varphi = \exists x.\, \varphi_1$ **then**
        return FIXPOINT$(\sigma, \mathsf{enf}_{\mathsf{ex}, \varphi_1, v})$

function $\mathsf{enf}_{\mathsf{or}, \varphi_1, \varphi_2, v}(\sigma)$
    $(D^-, D^+) \leftarrow (\emptyset, \emptyset)$
    if $v \in \mathrm{SAT}(\varphi_1, \sigma)$ **then**
        $(D^-, D^+) \leftarrow (D^-, D^+) \uplus \mathsf{enf}(\varphi_1, \sigma, v)$
    if $v \in \mathrm{SAT}(\varphi_2, \sigma)$ **then**
        $(D^-, D^+) \leftarrow (D^-, D^+) \uplus \mathsf{enf}(\varphi_2, \sigma, v)$
    return (D^-, D^+)

function $\mathsf{enf}_{\mathsf{ex}, \varphi_1, v}(\sigma)$
    $(D^-, D^+) \leftarrow (\emptyset, \emptyset)$
    for $v \in \mathbb{D}$ s.t. $v[x \mapsto v] \in \mathrm{SAT}(\varphi_1, \sigma)$ **do**
        $(D^-, D^+) \leftarrow (D^-, D^+) \uplus \mathsf{enf}(\varphi_1, \sigma, v[x \mapsto v])$
    return (D^-, D^+)

6.2 Enforcer

Given $\sigma \in \mathbb{T}_f$, $\tau \geq \mathsf{lts}(\sigma)$, and $D, D^-, D^+ \in \mathbb{DB}$, we first define the function update as

$$\mathsf{update}(\sigma \cdot ((\tau, D)), (D^-, D^+)) := \sigma \cdot ((\tau, (D \cup D^+) \setminus D^-)).$$

Namely, update returns the trace obtained by adding all events from D^+ and removing all events from D^- in the last database of σ.

For any $\sigma \in \mathbb{T}_f$ and enforcer μ, we define $\ell_\mu(\sigma) \in \mathbb{T}_f$ as the limit of the sequence $(u_i)_{i \in \mathbb{N}} \in \mathbb{T}_f^{\mathbb{N}}$ defined by $u_0 = \sigma$ and for all $i \in \mathbb{N}$, $u_{i+1} = \mathsf{update}(u_i, \mu(u_i))$. This limit is always well-defined [35], and if $u_{i+1} = u_i$ for some $i \in \mathbb{N}$, we have $\ell_\mu(\sigma) = u_i$. This allows us to define a routine FIXPOINT(σ, μ) that iteratively computes $u_0, u_1, \ldots, u_i, \ldots$, returns $\ell_\mu(\sigma) = u_i$ as soon as $(u_i)_{i \in \mathbb{N}}$ reaches a fixpoint $u_{i+1} = u_i$, and does not terminate otherwise. We will later show that, in our setup, this procedure always terminates.

Our enforcer relies on the function enf described in Algorithm 1, which takes as an input a future-free and monitorable GMFOTL formula φ, a finite trace σ, and a valuation v such that $v, |\sigma| \models_\sigma \varphi$, and returns a pair of sets of events to be respectively suppressed and caused at the last timepoint in σ in order to obtain some new trace σ' such that $v, |\sigma| \not\models_{\sigma'} \varphi$. For notational convenience, we denote by $\uplus$ the elementwise union of pairs of sets $(A, B) \uplus (C, D) = (A \cup C, B \cup D)$.

The intuition behind enf is as follows. If the formula φ is reduced to an atom $r(t_1, \ldots, t_n)$ or $\neg r(t_1, \ldots, t_n)$, we can make it false by suppressing or causing a single event. If φ is of the form $\varphi_1 \wedge \varphi_2$ with $\varphi \in \mathrm{GMFOTL}$, it is sufficient to make φ_1 false to make φ false: enf looks for events to be suppressed or caused in φ_1.

For formulae of the form $\varphi_1 \vee \varphi_2$, additional care is needed. At first glance, the strategy used for $\wedge$ seems applicable, modulo a simple case distinction: if both φ_1 and φ_2 are satisfied by a given pair of a trace and a valuation, we need to find events to suppress or cause in *both* subformulae; if only one conjunct is satisfied, we look for events to suppress or cause in this subformula only. But such a one-step strategy is insufficient.

Example 5. Consider the formula $\varphi_6 = \mathtt{Open}(1) \vee (\neg\mathtt{Close}(2) \wedge \neg\mathtt{Open}(1)) \in \mathrm{GMFOTL}$ and the trace $\sigma_6 = ((0, \{(\mathtt{Open}, (1))\}))$. Only the left disjunct is satisfied. Hence, applying the above strategy would produce the trace $\sigma_6' = ((0, \emptyset))$, which again satisfies

φ_6 as it satisfies the right disjunct now. Hence, after having suppressed $(\mathtt{Open}, (1))$ we must check for satisfaction of φ_6 again, and, if necessary, select additional events to be suppressed or caused, here causing $(\mathtt{Close}, (2))$ suffices. This results in the trace $\sigma_6'' = ((0, \{(\mathtt{Close}, (2))\}))$, which now does not satisfy φ_6.

The above iterative approach, which performs a fixpoint computation, is formalized as a call to $\text{FIXPOINT}(\sigma, \mathsf{enf}_{\mathsf{or}, \varphi_1, \varphi_2, v})$, where $\mathsf{enf}_{\mathsf{or}, \varphi_1, \varphi_2, v}$ performs the above case distinction for a fixed valuation v satisfying $\varphi_1 \vee \varphi_2$.

The same problem arises with existentially quantified formulae of the form $\exists x.\ \varphi_1$. For fixed v, function $\mathsf{enf}_{\mathsf{ex}, \varphi_1, v}$ identifies events that must be suppressed or caused to prevent the satisfaction of φ_1 using any valuation v' extending v, and a call to $\text{FIXPOINT}(\sigma, \mathsf{enf}_{\mathsf{ex}, \varphi_1, v})$ computes the corresponding fixpoint.

Finally, for any closed, monitorable and future-free $\varphi \in \text{GMFOTL}$, we define our tentative enforcer for $\square \neg \varphi$ as

$$\hat{\mu}_\varphi(\rho) = \begin{cases} \mathsf{enf}(\varphi, \rho, \emptyset) & \text{if } |\sigma| \models_\rho \varphi \\ (\emptyset, \emptyset) & \text{otherwise.} \end{cases}$$

Example 6. Consider the GMFOTL monitorable formula

$$\varphi_7 = \underbrace{(\exists x.\ \mathtt{Open}(x) \wedge \blacklozenge_{[0,5]} \mathtt{Close}(x))}_{\varphi_7^1} \vee \underbrace{(\exists y.\ \neg\mathtt{Close}(y) \wedge \neg\mathtt{Close}(y) \, \mathsf{S}_{[5,\infty)} \, \mathtt{Open}(y))}_{\varphi_7^2},$$

which is satisfied whenever an $(\mathtt{Open}, (x))$ event follows a $(\mathtt{Close}, (x))$ within 5 time units for some $x \in \mathbb{D}$, or there is a $(\mathtt{Close}, (y))$ event for some $y \in \mathbb{D}$ that is not followed by any $(\mathtt{Close}, (y))$ event within 5 time units. Consider the following trace:

$$\sigma_7 = ((0, \{(\mathtt{Open}, (1))\}), (1, \{(\mathtt{Close}, (2))\}), (5, \{(\mathtt{Open}, (2))\})).$$

We have $\models_{\sigma_7} \varphi_7$: events $(\mathtt{Close}, (2))$ and $(\mathtt{Open}, (2))$ at timestamps 1 and 5 satisfy the left disjunct, while the $(\mathtt{Open}, (1))$ event at timestamp 0 and the lack of a $(\mathtt{Close}, (1))$ event between timestamps 0 and 5 satisfies the right disjunct. As φ_7 is closed, the set of valuations satisfying it is $\{\emptyset\}$, where $\emptyset$ denotes the empty application. We compute $\mathsf{enf}(\varphi_7, \sigma_7, \emptyset) = \text{FIXPOINT}(\sigma_7, \mathsf{enf}_{\mathsf{or}, \varphi_7^1, \varphi_7^2, \emptyset})$.

Since σ_7 satisfies both φ_7^1 and φ_7^2, we get:

$$\begin{aligned}
\mathsf{enf}_{\mathsf{or}, \varphi_7^1, \varphi_7^2, \emptyset}(\sigma_7) &= \mathsf{enf}(\varphi_7^1, \sigma_7, \emptyset) \uplus \mathsf{enf}(\varphi_7^2, \sigma_7, \emptyset) \\
&= \mathsf{enf}(\mathtt{Open}(x) \wedge \blacklozenge_{[0,5]} \mathtt{Close}(x), \sigma_7, \{x \mapsto 2\}) \uplus \\
&\quad \mathsf{enf}(\neg\mathtt{Close}(y) \wedge \neg\mathtt{Close}(y) \, \mathsf{S}_{[5,\infty)} \, \mathtt{Open}(y), \sigma_7, \{y \mapsto 1\}) \\
&= \mathsf{enf}(\mathtt{Open}(x), \sigma_7, \{\{x \mapsto 2\}\}) \uplus \mathsf{enf}(\neg\mathtt{Close}(y), \sigma_7, \{\{y \mapsto 1\}\}) \\
&= (\{(\mathtt{Open}, (2))\}, \emptyset) \uplus (\emptyset, \{(\mathtt{Close}, (1))\}) \\
&= (\{(\mathtt{Open}, (2))\}, \{(\mathtt{Close}, (1))\}).
\end{aligned}$$

We then update σ_7:

$$\begin{aligned}
\sigma_7' &= \mathsf{update}(\sigma_7, \mathsf{enf}_{\mathsf{or}, \varphi_7^1, \varphi_7^2, \emptyset}(\sigma_7)) \\
&= ((\{0, \mathtt{Open}, (1)\}), (1, \{\mathtt{Close}, (2)\}), (5, \{\mathtt{Close}, (1)\}))
\end{aligned}$$

and check that $\sigma'_7 = \mathsf{update}(\sigma'_7, \mathsf{enf}_{\mathsf{or}, \varphi_7^1, \varphi_7^2, \emptyset}(\sigma'_7))$, i.e., that $\not\models_{\sigma'_7} \varphi_7$.

Hence, we finally get $\hat{\mu}_{\varphi_7}(\sigma_7) = \mathsf{enf}(\varphi_7, \sigma_7, \emptyset) = (\{(\mathtt{Open}, (2))\}, \{(\mathtt{Close}, (1))\})$.

6.3   Correctness and Transparency

For any monitorable, future-free and closed $\varphi \in \mathrm{GMFOTL}$ and finite $\sigma \in \mathbb{T}^\dagger$, the enforcer $\hat{\mu}_\varphi$ always terminates. Termination is a consequence of Lemma 4 above; the corresponding proofs are given in our extended report [35]. Having established termination, we can prove that our enforcer is correct and transparent:

Theorem 3. *Let $\varphi \in \mathrm{GMFOTL}$ be closed, monitorable and future-free. Then $\hat{\mu}_\varphi$ is a correct and transparent enforcer with respect to $\mathcal{L}(\Box \neg \varphi) \cap \mathbb{T}^\dagger$ and $\mathbb{DB}^\dagger$.*

At this point, it is worth reflecting on the effect that the assumption $\mathsf{Sup} \cap \mathsf{Cau} = \emptyset$ has on the correctness of our enforcer. In general, dropping this assumption results in some non-enforceable formula being equivalent to some formula $\Box \neg \psi$ with $\psi \in \mathrm{GMFOTL}$; thus, Theorem 2 no longer holds. For example, a formula such as $\varphi_7 = \Box \neg(\mathtt{C} \vee \neg \mathtt{C})$ where $\mathtt{C} \in \mathsf{Sup} \cap \mathsf{Cau}$ and $a(\mathtt{C}) = 0$ is not enforceable: given an initially empty trace—on which, by convention, φ_7 is satisfied—adding any first timepoint makes the formula unsatisfiable, since $\neg(\mathtt{C} \vee \neg \mathtt{C}) \equiv \bot$. This rules out enforceability, which requires that appending only-observable events to a valid trace does not lead to a violation.

To understand why we need to assume $\mathsf{Sup} \cap \mathsf{Cau} = \emptyset$ for the above algorithm to be correct, consider the behavior of $\hat{\mu}_{\varphi_7}$ for the (non-enforceable) formula φ_7 above on the trace $\sigma_7 = ((\{\mathtt{C}\}, 0))$. The enforcer calls $\mathrm{FIXPOINT}(\sigma_7, \mathsf{enf}_{\mathsf{or}, \mathtt{C}, \neg \mathtt{C}, \emptyset})$, which itself calls $\mathsf{enf}_{\mathsf{or}, \mathtt{C}, \neg \mathtt{C}, \emptyset}(\sigma_7)$. This routine determines that only the left disjunct $\mathtt{C}$ is satisfied, and returns the actions $(D^-, D^+) = (\{\mathtt{C}\}, \emptyset)$. We get $\sigma'_7 = ((0, \emptyset))$ and call $\mathsf{enf}_{\mathsf{or}, \mathtt{C}, \neg \mathtt{C}}(\sigma_7)$ again to find a fixpoint. Now, the second disjunct is not satisfied, leading to the actions $(D^-, D^+) = (\emptyset, \{\mathtt{C}\})$ and to the updated trace $\sigma''_7 = ((0, \{\mathtt{C}\})) = \sigma_7$. The same process repeats indefinitely.

When $\mathsf{Sup} \cap \mathsf{Cau} = \emptyset$, such a behavior is avoided. Since only suppressable events are suppressed and causable events caused, and since suppressable and causable events are disjoint, the algorithm will never try to suppress (resp. cause) an event that it has previously caused (resp. suppressed). Hence, the sets of caused and suppressed events can only grow during the fixpoint computation. This ensures termination, as any new iteration except the last one must compute at least one new event to cause or suppress.

Note that the assumption $\mathsf{Sup} \cap \mathsf{Cau} = \emptyset$ can be relaxed if we additionally require each suppressable *and* causable event to appear only with, or only without, a negation in the formula. In the definition of enf, each element from $\mathsf{Sup} \cap \mathsf{Cau}$ can then be considered to belong to Sup or Cau only.

7   Implementation

We have implemented our enforcer in the EnfPoly tool [34], which extends the MonPoly tool [13] with ca. 500 lines of OCaml code. Users can specify suppressable and causable events by adding "–" or "+" after the corresponding event description in the signature.

Example 7. The example signature Σ can be specified as:
```
Open(int)- Close(int)+ Knock(int)
```

Events that are both enforceable and causable can be specified, e.g. as `SomeE+-`. In this case, for each formula to be enforced, a simple constraint-solving procedure is used to determine whether each such event can be considered only enforceable or only causable in the context of the current formula.

Strictly relative-past MFOTL Note that Algorithm 1 takes as input a monitorable and enforceable MFOTL$_\square^\mathcal{F}$ formula. Monitorability and enforcability can be syntactically approximated, but determining whether an MFOTL formula is future-free is undecidable [35]. Therefore, we have also developed a syntactical approximation of future-free formulae, called *strictly relative-past* formulae, which EnfPoly uses in practice. We formally define the fragment in our extended report [35]. Intuitively, all formulae that use only past temporal operators (i.e. *past-only MFOTL*) are strictly relative-past. Additionally, the strictly relative-past fragment contains many non-past formulae, for which one can statically verify that they do not depend on the future. For example, $\varphi_8 = \blacklozenge_{[5,+\infty)}(\texttt{Close}(2) \, \mathsf{U}_{[0,5)} \, \texttt{Open}(3))$ is strictly relative-past, but not past-only. Observe that the intervals of the temporal operators of φ_8 ensure that its truth value does not depend on future events: the evaluation of φ_8 at timestamp τ uses `Close` events from timestamps $\leq \tau - 5$, and `Open` from timestamps $< \tau - 5 + 5 = \tau$, which all lie in the past.

To enforce a formula of the form $\square \neg \varphi$, EnfPoly checks if φ is closed, in GMFOTL, and strict relative-past. Associative and commutative rewriting is used to relax the GMFOTL membership conditions in conjuncts. Then, the enforcement loop starts. At every timepoint, the enforcer reacts either with `OK`, if there is no violation, or with a set of events to cause and a set of events to suppress, otherwise.

Example 8. The output of EnfPoly when enforcing formulae $\square \neg \varphi_6$ and $\square \neg \varphi_7$ (from Examples 5 and 6) on traces σ_6 and σ_7, respectively, is shown in the table below.

Formula: $\square \neg \varphi_6$, Trace: σ_6	Formula: $\square \neg \varphi_7$, Trace: σ_7
`@0 Open(1);`	`@0 Open(1);`
`[Enforcer] Suppress: Open(1)`	`[Enforcer] OK.`
`[Enforcer] Cause: Close(2)`	`@1 Close(2);`
`[Enforcer] OK.`	`[Enforcer] OK.`
	`@5 Open(2);`
	`[Enforcer] Suppress: Open(2)`
	`[Enforcer] Cause: Close(1)`
	`[Enforcer] OK.`

Timestamped databases (prefixed with `@`) of a trace are incrementally input to EnfPoly, while its output (prefixed with `[Enforcer]`) is shown chronologically interleved with the input. When enforcing $\square \neg \varphi_6$ on σ_6, the enforcer immediately reacts to the the first database $\{(\texttt{Open}, (1))\}$ at timestamp 0 with two actions: it suppresses the event $(\texttt{Open}, (1))$ and causes the event $(\texttt{Close}, (2))$. Finally, it indicates that it has finished enforcing the formula by emitting `OK`. For $\square \neg \varphi_7$, EnfPoly processes three timestamped databases. The first two do not violate the policy and hence there is no reaction other than `OK` from the enforcer. The third database causes a violation and the enforcer suppresses event $(\texttt{Open}, (2))$ and causes event $(\texttt{Close}, (1))$ to satisfy the policy.

(a) Policy χ_1 (b) Policy χ_2 (c) Policy χ_3

$\chi_1 = \Box \neg (\mathbf{r} \wedge ((\blacklozenge_{[1,5]}\, \mathbf{r}) \vee (\bullet_{[0,0]}\, \blacklozenge_{[0,0]}\, \mathbf{r}))))$

$\chi_2 = \Box \neg ((\mathbf{g} \wedge \blacklozenge_{[0,6)}(\mathbf{r} \wedge (\neg \bullet \blacklozenge \mathbf{r}))) \vee (\mathbf{r} \wedge (\bullet(\neg \mathbf{g}\, \mathsf{S}\,(\mathbf{r} \wedge (\neg \bullet \blacklozenge \mathbf{r})))))))$

$\chi_3 = \Box \neg (\mathbf{r} \wedge ((\neg \mathbf{r}\, \mathsf{S}_{(20,\infty)}\, \mathbf{g}) \vee (\neg \mathbf{r}\, \mathsf{S}_{[0,15)}\, \mathbf{g}) \vee (\neg \mathbf{g}\, \mathsf{S}\, \mathbf{r}))) \vee \mathbf{g} \wedge \neg \mathbf{r}\, \mathsf{S}\, \mathbf{g})$

$\mathsf{Sup} = \{\mathbf{r}\}, \mathsf{Obs} = \{\mathbf{a}\}$

$\mathsf{Sup} = \{\mathbf{r},\mathbf{g}\}, \mathsf{Obs} = \emptyset$

$\mathsf{Sup} = \{\mathbf{r},\mathbf{g}\}, \mathsf{Obs} = \{\mathbf{a}\}$
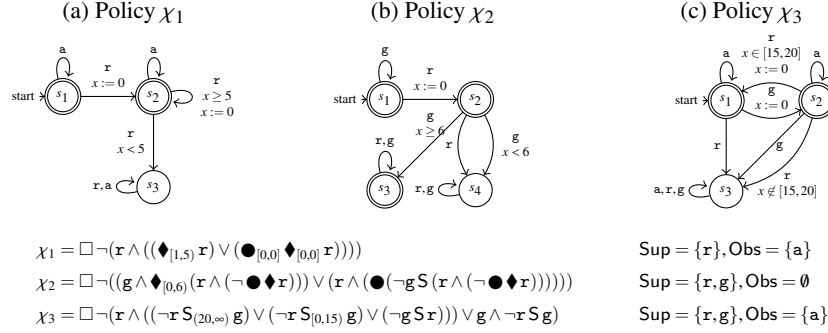
Fig. 3: Policies used to compare EnfPoly to GREP

8 Evaluation

We now compare our enforcer with other state-of-the-art tools. As our tool is the first one to support the enforcement of first-order temporal policies, comparison is only possible with (1) *propositional* temporal enforcers or (2) first-order temporal *monitors*.

Note that when there are no causable events in the signature, online monitoring tools can be used as online enforcers in the following way. First, before the events of every timepoint are sent to the monitor, save the monitor's internal state. Then, have the monitor process the timepoint. If the monitor does not detect a violation, save the monitor's state again and proceed with the next timepoint. If a violation is detected, restore the previous saved state and re-read *only the only-observable events* from the timepoint that led to a violation, suppressing all suppressable events from the last timepoint. When the formula to monitor is enforceable and there are no causable events in the signature, this construction always provides a valid enforcer. This approach has been used recently [33] to perform MFOTL enforcement with MonPoly.

Our evaluation aims to answer the following research questions:

RQ1. Does EnfPoly show better performance than existing propositional enforcers?

RQ2a. Given an MFOTL formula, how much overhead does EnfPoly's enforcement cause compared to MonPoly's monitoring of the same formula?

RQ2b. Does EnfPoly show better performance in enforcing formulae over a signature with no causable events than MonPoly adapted to be an online enforcer?

For RQ1, we focus on runtime enforcement tools, which use a setup similar to ours in terms of enforcement capabilities. We compare EnfPoly to GREP [49]. The tool GREP, along with TiPEX and Proactive Libraries, is one of three tools referenced in a recent survey paper [26]. GREP has been shown to outperform TiPEX by up to two orders of magnitude [49], and, unlike Proactive Libraries, it comes with an publicly available implementation. For RQ2, we compare EnfPoly to MonPoly [13].

In all experiments, we measure the enforcers' memory using Python's `psutil`. We also measure enforcers' total runtime, as well as their latency, i.e., the time spent waiting for an enforcer to compute its output, which we normalize by the number of events in the trace. The speedup of our tool with respect to a tool t is computed as the difference between t's and our tool's runtime divided by t's runtime. All experiments are run on an 2.4 GHz Intel Core i5-1135G7 QuadCore CPU with 32 GB RAM.

(a) Runtime performance for various choices of p, fixing $N = 25$, $n = 10$, $L = 5000$

Policy χ_1

Policy χ_2

Policy χ_3
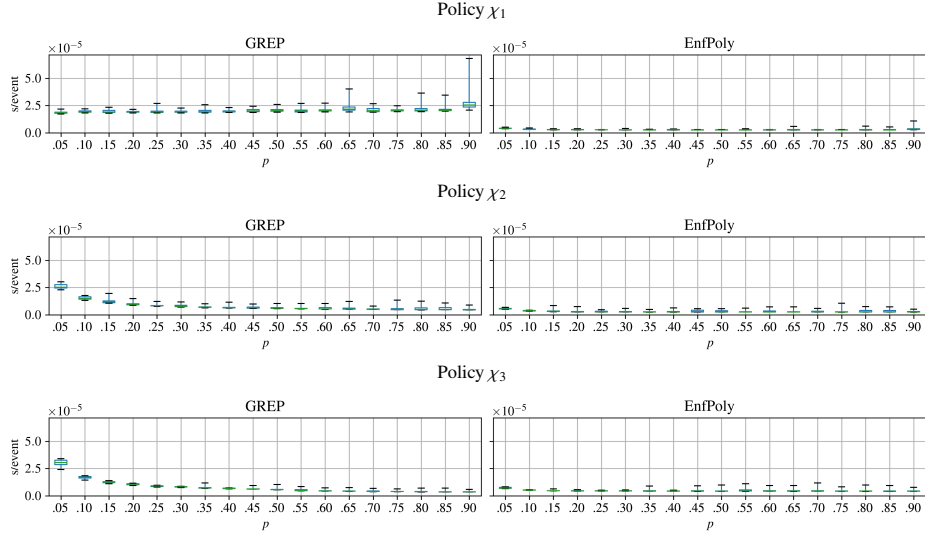
(b) Runtime and memory over time for $N = 1000$ executions, fixing $n = 10$, $L = 5000$, $p = 0.1$
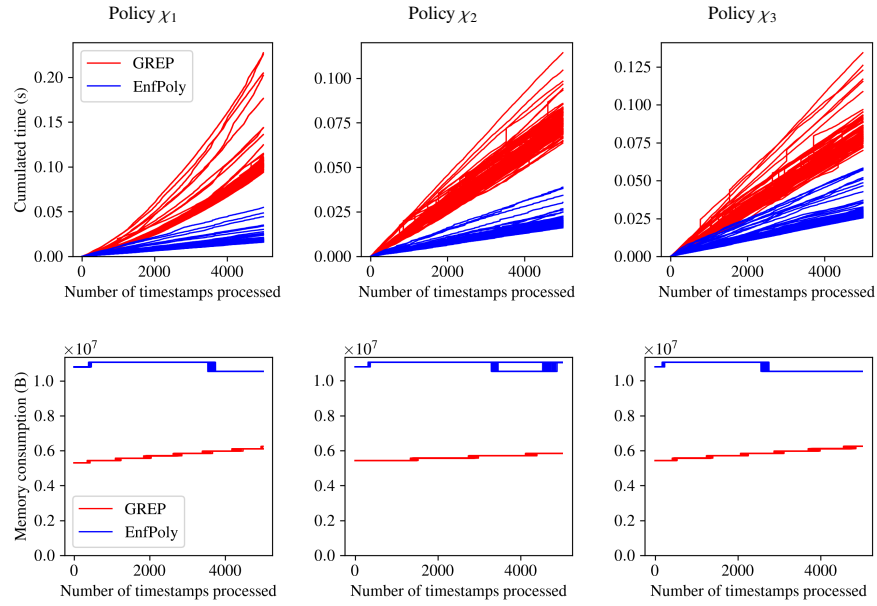

Fig. 4: Runtime and memory consumption of EnfPoly and GREP

EnfPoly vs GREP (RQ1). To compare the performance of the two enforcers, we consider the three policies presented on Figure 3, which are slight adaptations of the three benchmark examples used in [45, 49] to evaluate GREP and TiPEX. The original benchmark policies were not enforceable according to Definition 2. To enable enforceability, some previously non-accepting states were made accepting. As GREP takes as input policies specified as timed automata, we provide both an automaton and an MFOTL$_\square^{\mathcal{F}}$ definition for each formula. These specifications are equivalent on traces with at most one event per database. We generate such random traces of length $L \cdot n = 50\,000$ with

- $L = 5\,000$ unique timestamps from $\{1, \ldots, L\}$;
- timestamps τ_i equal to $\lceil \frac{i}{n} \rceil$ for timepoint $i \in \{1 \ldots L \cdot n\}$, where $n = 10$;
- each timepoint containing an event with probability p and no event otherwise; and
- event names sampled uniformly from $\{a, r\}$ for χ_1 and from $\{a, g, r\}$ for χ_2 and χ_3.

For GREP, the duration of a time unit is set to 1 ms. GREP's and EnfPoly's code is instrumented to report the latency of processing inputs (i.e., excluding communication costs). Communication costs were excluded since GREP and EnfPoly receive inputs in a different format (one timepoint per line for EnfPoly, several timepoints per line for GREP). The experiment is repeated $N = 25$ times for various values of p to measure the effect of the *event rate* (i.e., the number of events per time unit) on the enforcers' performance. Note that as the signatures of χ_1, χ_2, and χ_3 contain at most three event names, we can keep the maximal number of events per timestamp small, fixing $n = 10$ and varying p only. GREP is run in online mode with the "fast" option (flag `-f`) activated.

For formulae χ_1 and χ_2, EnfPoly is faster than GREP on average for all values of p, with a speedup between 40% and 90%. For χ_3, GREP outperforms EnfPoly by up to 20% for $p \geq 0.55$, but underperforms it for $p < 0.55$. The corresponding summary figures are presented in Figure 4a. Numerical data is given in Table 1 in the Appendix.

Additionally, in Figure 4b, we plot the cumulated latency and the memory consumption over time for $N = 100$ individual executions of both EnfPoly and GREP. The memory consumption of our tool is constant over time, while GREP's is linear. GREP also displays quadratic latency for policy χ_1, while EnfPoly's latency is constant in all three cases, resulting in linear cumulative latency.

EnfPoly vs MonPoly (RQ2). For RQ2a, we compare the runtime of EnfPoly with the runtime of MonPoly (used as a monitor) on the same traces and formulae. For RQ2b, we repeat this experiment using MonPoly as an enforcer, in the way described above.

In both cases, we generate random enforceable and monitorable MFOTL formulae and random traces over a signature $(\text{int}, \mathbb{E}, a)$ with $\mathbb{E} = \text{Sup} = \{\text{A}, \text{B}, \text{C}\}$ and $a(\cdot) = 1$. The random formula generator has a configurable maximal depth d and samples bounds of temporal operator intervals uniformly from $\{(i, j) \in \{0, \ldots, I\}^2 \mid i \leq j\}$. Random traces of length $1\,000$ are generated with timestamps $1, 2, \ldots, L$ with $L = 1\,000$ with no repetitions. The number of events in a database is sampled according to the binomial distribution with n trials and success probability p, while event names are sampled uniformly from $\mathbb{E}$. Finally, event's arguments are sampled uniformly from $\{1, \ldots, A\}$.

Given parameters $n, A, d, I \in \mathbb{N}$ and $p \in [0, 1]$, both tools are executed on pairs of independently generated random traces and enforceable and monitorable MFOTL$_\square^{\mathcal{F}}$ formulae with the same combinations of parameters, repeated $N = 25$ times.

For all values of the parameters, enforcement with EnfPoly adds up to 50% runtime overhead on top of the costs of monitoring with MonPoly, and does not affect memory consumption. On the other hand, using EnfPoly for enforcement is still 4 to 20 times faster than using MonPoly as an enforcer, working in the way described above, and with a comparable memory consumption. Most of the overhead of MonPoly used as an enforcer is due to loading and saving the (complete) monitor state at each iteration, which EnfPoly avoids. Average runtime costs are under 0.1 ms per event, with most averages under 10 μs. In individual executions, both tools display constant time and memory consumption. Detailed numerical results can be found in Table 1 in the Appendix (for RQ2b), as well as in our extended report [35] (for RQ2a).

Discussion The above experiments show that EnfPoly, despite supporting a much larger specification language, displays a runtime and memory performance at least as good as GREP's. Our enforcer's performance is less sensitive to the choice of the input formula and consumes a constant amount of memory over time. Compared to using MonPoly as an MFOTL enforcer, EnfPoly provides a speedup of one order of magnitude. Runtime and memory consumption per event processed is stable or decreasing when more events occur simultaneously, and is not affected by longer trace sizes.

9   Conclusion

We have presented both the theory and practice of enforcing metric first-order temporal logic (MFOTL) formulae with disjoint sets of causable and suppressable events. We have characterized enforceability for MFOTL for such enforcers and proposed an efficient enforcement algorithm. Our enforcer EnfPoly extends the MonPoly monitoring tool and it is the first tool for first-order temporal logic enforcement. We have evaluated EnfPoly and showed that although it supports a more expressive language it can still outperform state-of-the-art enforcers.

As future work, we plan to generalize our approach to allow events that are both suppressable and causable. Currently, it remains open whether enforceability can be characterized syntactically modulo equivalence (as in Theorem 2) when this assumption is lifted. But even if no such characterization exists, in practice one could develop enforcement algorithms for larger (syntactical) fragments of enforceable policies.

References

1. Abadi, M., Lamport, L., Wolper, P.: Realizable and unrealizable specifications of reactive systems. In: Ausiello, G., Dezani-Ciancaglini, M., Ronchi Della Rocca, S. (eds.) International Colloquium on Automata, Languages and Programming (ICALP). LNCS, vol. 372, pp. 1–17. Springer (1989)

2. Aceto, L., Cassar, I., Francalanza, A., Ingólfsdóttir, A.: On bidirectional runtime enforcement. In: Peters, K., Willemse, T. (eds.) International Conference Formal Techniques for Distributed Objects, Components, and Systems (FORTE). LNCS, vol. 12719, pp. 3–21. Springer (2021)

3. Alur, R., Feder, T., Henzinger, T.: The benefits of relaxing punctuality. Journal of the ACM **43**(1), 116–146 (1996)

4. Ames, S.R., Gasser, M., Schell, R.R.: Security kernel design and implementation: An introduction. Computer **16**(7), 14–22 (1983)

5. Arfelt, E., Basin, D., Debois, S.: Monitoring the GDPR. In: European Symposium on Research in Computer Security. pp. 681–699. Springer (2019)

6. Asarin, E., Maler, O., Pnueli, A.: Symbolic controller synthesis for discrete and timed systems. In: Antsaklis, P., Kohn, W., Nerode, A., Sastry, S. (eds.) International Workshop on Hybrid Systems. LNCS, vol. 999, pp. 1–20. Springer (1994)

7. Bartocci, E., Falcone, Y.: Lectures on Runtime Verification, LNCS, vol. 10457. Springer (2018)

8. Basin, D., Dardinier, T., Heimes, L., Krstić, S., Raszyk, M., Schneider, J., Traytel, D.: A formally verified, optimized monitor for metric first-order dynamic logic. In: Peltier, N., Sofronie-Stokkermans, V. (eds.) International Joint Conference on Automated Reasoning (IJCAR). LNCS, vol. 12166, pp. 432–453. Springer (2020)

9. Basin, D., Debois, S., Hildebrandt, T.: In the Nick of Time: Proactive prevention of obligation violations. In: Computer Security Foundations Symposium (CSF). pp. 120–134. IEEE (2016)

10. Basin, D., Debois, S., Hildebrandt, T.: Proactive enforcement of provisions and obligations. Journal of Computer Security (2021)

11. Basin, D., Jugé, V., Klaedtke, F., Zălinescu, E.: Enforceable security policies revisited. ACM Trans. Inf. Syst. Secur. **16**(1), 1–26 (2013)

12. Basin, D., Klaedtke, F., Müller, S., Zălinescu, E.: Monitoring metric first-order temporal properties. Journal of the ACM **62**(2), 1–45 (2015)

13. Basin, D., Klaedtke, F., Zalinescu, E.: The MonPoly monitoring tool. In: Reger, G., Havelund, K. (eds.) International Workshop on Competitions, Usability, Benchmarks, Evaluation, and Standardisation for Runtime Verification Tools (RV-CuBES). Kalpa, vol. 3, pp. 19–28 (2017)

14. Bauer, L., Ligatti, J., Walker, D.: More enforceable security policies. In: Workshop on Foundations of Computer Security (FCS). Citeseer (2002)

15. Behrmann, G., Cougnard, A., David, A., Fleury, E., Larsen, K., Lime, D.: UPPAAL-Tiga: Time for playing games! In: Damm, W., Hermanns, H. (eds.) International Conference Computer Aided Verification (CAV). LNCS, vol. 4590, pp. 121–125. Springer (2007)

16. Bohy, A., Bruyère, V., Filiot, E., Jin, N., Raskin, J.: Acacia+, a tool for LTL synthesis. In: Madhusudan, P., Seshia, S.A. (eds.) International Conference Computer Aided Verification (CAV). LNCS, vol. 7358, pp. 652–657. Springer (2012)

17. Bouyer, P., Bozzelli, L., Chevalier, F.: Controller synthesis for MTL specifications. In: Baier, C., Hermanns, H. (eds.) International Conference on Concurrency Theory (CONCUR). LNCS, vol. 4137, pp. 450–464. Springer (2006)

18. Brihaye, T., Geeraerts, G., Ho, H.M., Monmege, B.: MightyL: A compositional translation from MITL to timed automata. In: International Conference on Computer Aided Verification (CAV). pp. 421–440. Springer (2017)

19. Bulychev, P., David, A., Larsen, K., Li, G.: Efficient controller synthesis for a fragment of $\text{MTL}_{0,\infty}$. Acta Informatica **51**(3-4), 165–192 (2014)

20. Chomicki, J.: Efficient checking of temporal integrity constraints using bounded history encoding. ACM Trans. Database Syst. **20**(2), 149–186 (Jun 1995)

21. Dolzhenko, E., Ligatti, J., Reddy, S.: Modeling runtime enforcement with mandatory results automata. International Journal of Information Security **14**(1), 47–60 (2015)

22. Donzé, A., Raman, V.: BluSTL: controller synthesis from signal temporal logic specifications. In: Frehse, G., Althoff, M. (eds.) International Workshop on Applied veRification for Continuous & Hybrid systems (ARCH@CPSWeek). EPiC, vol. 34, pp. 160–168. EasyChair (2015)

23. Ehlers, R.: Unbeast: Symbolic bounded synthesis. In: Abdulla, P.A., Leino, K.R.M. (eds.) International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS). LNCS, vol. 6605, pp. 272–275. Springer (2011)

24. Erlingsson, Ú., Schneider, F.: SASI enforcement of security policies: a retrospective. In: Kienzle, D., Zurko, M.E., Greenwald, S., Serbau, C. (eds.) Workshop on New Security Paradigms. pp. 87–95. ACM (1999)

25. Falcone, Y., Jéron, T., Marchand, H., Pinisetty, S.: Runtime enforcement of regular timed properties by suppressing and delaying events. Sci. of Comp. Programming **123**, 2–41 (2016)

26. Falcone, Y., Krstić, S., Reger, G., Traytel, D.: A taxonomy for classifying runtime verification tools. Int. J. Softw. Tools Technol. Transf. **23**(2), 255–284 (2021)

27. Falcone, Y., Mounier, L., Fernandez, J., Richier, J.: Runtime enforcement monitors: composition, synthesis, and enforcement abilities. Form. Methods Syst. Des. **38**(3), 223–262 (2011)

28. Falcone, Y., Pinisetty, S.: On the runtime enforcement of timed properties. In: International Conference on Runtime Verification (RV). pp. 48–69. Springer (2019)

29. Filiot, E., Jin, N., Raskin, J.: Antichains and compositional algorithms for LTL synthesis. Form. Methods Syst. Des. **39**(3), 261–296 (2011)

30. Havelund, K., Peled, D., Ulus, D.: DejaVu: A monitoring tool for first-order temporal logic. In: Workshop on Monitoring and Testing of Cyber-Physical Systems (MT-CPS). pp. 12–13. IEEE (2018)

31. Havelund, K., Peled, D., Ulus, D.: First-order temporal logic monitoring with BDDs. Form. Methods Syst. Des. **56**(1), 1–21 (2020)

32. Hofmann, T., Schupp, S.: TACoS: A tool for MTL controller synthesis. In: Calinescu, R., Pasareanu, C.S. (eds.) International Conference on Software Engineering and Formal Methods (SEFM). LNCS, vol. 13085, pp. 372–379. Springer (2021)

33. Hublet, F.: The Databank Model. Master's thesis, ETH Zürich (2021)

34. Hublet, F., Basin, D., Krstić, S.: EnfPoly's development repository. `https://gitlab.ethz.ch/fhublet/mfotl-enforcement` (2022)

35. Hublet, F., Basin, D., Krstić, S.: Real-time policy enforcement with metric first-order temporal logic. Tech. rep., ETH Zürich (2022), `https://gitlab.ethz.ch/fhublet/mfotl-enforcement/-/blob/main/paper/extended.pdf`, Extended Report

36. Jobstmann, B., Bloem, R.: Optimizations for LTL synthesis. In: International Conference Formal Methods in Computer-Aided Design (FMCAD). pp. 117–124. IEEE (2006)

37. Khoussainov, B., Nerode, A.: Automatic presentations of structures. In: International Workshop on Logic and Computational Complexity. pp. 367–392. Springer (1994)

38. Krstić, S., Schneider, J.: A benchmark generator for online first-order monitoring. In: Deshmukh, J., Ničković, D. (eds.) International Conference on Runtime Verification (RV). vol. 12399, pp. 482–494. Springer (2020)

39. Kupferman, O., Vardi, M.Y.: Safraless decision procedures. In: Symposium on Foundations of Computer Science (FOCS). pp. 531–542. IEEE (2005)

40. Li, G., Jensen, P., Larsen, K., Legay, A., Poulsen, D.: Practical controller synthesis for $mtl_{0,\infty}$. In: Erdogmus, H., Havelund, K. (eds.) ACM SIGSOFT International SPIN Symposium on Model Checking of Software. pp. 102–111. ACM (2017)

41. Ligatti, J., Bauer, L., Walker, D.: Enforcing non-safety security policies with program monitors. In: European Symposium on Research in Computer Security. pp. 355–373. Springer (2005)

42. Ligatti, J., Bauer, L., Walker, D.: Run-time enforcement of nonsafety policies. ACM Trans. Inf. Syst. Secur. **12**(3), 1–41 (2009)

43. Maler, O., Ničković, D., Pnueli, A.: From MITL to timed automata. In: International conference on formal modeling and analysis of timed systems. pp. 274–289. Springer (2006)

44. Peter, H., Ehlers, R., Mattmüller, R.: Synthia: Verification and synthesis for timed automata. In: Gopalakrishnan, G., Qadeer, S. (eds.) International Conference on Computer Aided Verification (CAV). LNCS, vol. 6806, pp. 649–655. Springer (2011)

45. Pinisetty, S., Falcone, Y., Jéron, T., Marchand, H.: TiPEX: A tool chain for timed property enforcement during execution. In: International Conference on Runtime Verification (RV). pp. 306–320. Springer (2015)
46. Pnueli, A., Rosner, R.: On the synthesis of a reactive module. In: ACM Symposium on Principles of Programming Languages (POPL). pp. 179–190. ACM (1989)
47. Pnueli, A., Rosner, R.: On the synthesis of an asynchronous reactive module. In: Ausiello, G., Dezani-Ciancaglini, M., Rocca, S.R.D. (eds.) International Colloquium on Automata, Languages and Programming (ICALP). LNCS, vol. 372, pp. 652–671. Springer (1989)
48. Raman, V., Donzé, A., Sadigh, D., Murray, R., Seshia, S.: Reactive synthesis from signal temporal logic specifications. In: Girard, A., Sankaranarayanan, S. (eds.) International Conference on Hybrid Systems: Computation & Control (HSCC). pp. 239–248. ACM (2015)
49. Renard, M., Rollet, A., Falcone, Y.: GREP: games for the runtime enforcement of properties. In: Yevtushenko, N., Cavalli, A., Yenigün, H. (eds.) International Conference on Testing Software and Systems (ICTSS). LNCS, vol. 10533, pp. 259–275. Springer (2017)
50. Riganelli, O., Micucci, D., Mariani, L.: Policy enforcement with proactive libraries. In: International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS). pp. 182–192. IEEE (2017)
51. Rushby, J.: Design and verification of secure systems. In: Howard, J., Reed, D. (eds.) Symposium on Operating System Principles (SOSP). pp. 12–21. ACM (1981)
52. Rushby, J.: Kernels for safety. Safe and Secure Computing Systems pp. 210–220 (1989)
53. Schewe, S., Finkbeiner, B.: Bounded synthesis. In: Namjoshi, K.S., Yoneda, T., Higashino, T., Okamura, Y. (eds.) International Symposium Automated Technology for Verification and Analysis (ATVA). LNCS, vol. 4762, pp. 474–488. Springer (2007)
54. Schneider, F.: Enforceable security policies. ACM Trans. Inf. Syst. Secur. **3**(1), 30–50 (2000)
55. Schneider, J., Basin, D., Krstić, S., Traytel, D.: A formally verified monitor for metric first-order temporal logic. In: Finkbeiner, B., Mariani, L. (eds.) International Conference on Runtime Verification (RV). LNCS, vol. 11757, pp. 310–328. Springer (2019)
56. Zhu, S., Tabajara, L., Li, J., Pu, G., Vardi, M.: A symbolic approach to safety LTL synthesis. In: Strichman, O., Tzoref-Brill, R. (eds.) International Haifa Verification Conference (HVC). LNCS, vol. 10629, pp. 147–162. Springer (2017)

A   Evaluation Data

Table 1 shows the raw evaluation data produced by our experiments. The table on the left contains the data obtained when answering RQ1, while the data in the table on the right is obtained when answering RQ2. In the former we use three policies χ_1, χ_2, and χ_3, while in the latter we generate random enforceable and monitorable MFOTL formulae.

Parameter d is the depth of the generated random formulae, while I defines the sample space for the bounds of temporal operator intervals: $\{(i,j) \in \{0,\dots,I\}^2 \mid i \leq j\}$.

Random traces have length $L \cdot n$ with timestamps $1, 2, \dots, L$, each repeated n times. Event names are sampled uniformly from $\mathbb{E} = \{\mathtt{A}, \mathtt{B}, \mathtt{C}\}$, while their arguments are sampled uniformly from $\{1, \dots, A\}$. The number of events in a database is sampled according to the binomial distribution with n trials and success probability p.

Given parameters $n, A, d, I \in \mathbb{N}$ and $p \in [0, 1]$, both tools are executed on pairs of independently generated random traces and enforceable and monitorable MFOTL$_\square^{\mathcal{F}}$ formulae with the same combinations of parameters repeated N times.

Table 1: Mean runtime performance (standard deviation) for various parameter values

RQ1. EnfPoly vs. GREP: $N = 25$, $n = 10$, $L = 5000$

	χ_1				χ_2				χ_3		
p	GREP s/event	EnfPoly s/event	Speedup %	p	GREP s/event	EnfPoly s/event	Speedup %	p	GREP s/event	EnfPoly s/event	Speedup %
.05	1.86e-05 (1.01e-06)	4.16e-06 (4.14e-07)	77.7% (2.0%)	.05	2.60e-05 (2.01e-06)	5.56e-06 (4.45e-07)	78.5% (2.3%)	.05	3.03e-05 (2.51e-06)	7.18e-06 (4.70e-07)	76.1% (2.5%)
.10	1.95e-05 (1.05e-06)	3.21e-06 (3.48e-07)	83.6% (1.6%)	.10	1.55e-05 (1.25e-06)	3.63e-06 (3.14e-07)	76.3% (3.1%)	.10	1.68e-05 (1.09e-06)	5.20e-06 (1.64e-07)	69.0% (2.4%)
.15	1.94e-05 (1.18e-06)	2.90e-06 (2.34e-07)	85.1% (1.1%)	.15	1.24e-05 (1.97e-06)	3.58e-06 (1.47e-06)	71.5% (8.1%)	.15	1.25e-05 (7.15e-07)	4.76e-06 (3.80e-07)	61.8% (4.0%)
.20	1.93e-05 (7.45e-07)	2.82e-06 (2.79e-07)	85.3% (1.5%)	.20	1.01e-05 (1.53e-06)	3.18e-06 (1.09e-06)	68.7% (6.7%)	.20	1.05e-05 (5.64e-07)	4.50e-06 (2.33e-07)	57.0% (3.5%)
.25	1.99e-05 (2.02e-06)	2.72e-06 (1.33e-07)	86.2% (1.4%)	.25	8.66e-06 (1.07e-06)	3.00e-06 (6.34e-07)	65.5% (4.9%)	.25	8.87e-06 (4.87e-07)	4.54e-06 (3.08e-07)	48.7% (4.2%)
.30	1.95e-05 (1.07e-06)	2.74e-06 (2.90e-07)	86.0% (1.5%)	.30	8.37e-06 (1.34e-06)	3.10e-06 (9.00e-07)	63.3% (6.1%)	.30	8.14e-06 (3.74e-07)	4.49e-06 (2.35e-07)	44.7% (4.2%)
.35	1.98e-05 (1.51e-06)	2.66e-06 (1.94e-07)	86.5% (1.2%)	.35	7.28e-06 (8.33e-07)	2.83e-06 (6.38e-07)	61.3% (5.6%)	.35	7.60e-06 (1.14e-06)	4.83e-06 (1.32e-06)	37.0% (8.8%)
.40	1.99e-05 (1.10e-06)	2.77e-06 (2.92e-07)	86.0% (1.8%)	.40	6.88e-06 (1.25e-06)	3.04e-06 (9.51e-07)	56.4% (6.6%)	.40	6.76e-06 (2.77e-07)	4.34e-06 (2.30e-07)	35.7% (3.6%)
.45	2.04e-05 (1.51e-06)	2.63e-06 (1.16e-07)	87.1% (1.0%)	.45	6.78e-06 (1.02e-06)	3.29e-06 (9.20e-07)	51.9% (8.5%)	.45	6.53e-06 (8.48e-07)	4.68e-06 (1.09e-06)	28.7% (9.0%)
.50	2.08e-05 (1.78e-06)	2.65e-06 (1.71e-07)	87.2% (1.1%)	.50	6.57e-06 (1.18e-06)	3.17e-06 (9.69e-07)	52.2% (8.5%)	.50	6.00e-06 (1.08e-06)	4.76e-06 (1.44e-06)	21.3% (12.9%)
.55	2.06e-05 (2.12e-06)	2.69e-06 (2.80e-07)	86.8% (1.8%)	.55	6.05e-06 (1.09e-06)	2.95e-06 (9.52e-07)	51.8% (8.7%)	.55	5.34e-06 (1.28e-06)	5.43e-06 (1.88e-06)	-0.4% (15.6%)
.60	2.10e-05 (1.75e-06)	2.68e-06 (1.86e-07)	87.2% (1.2%)	.60	6.19e-06 (1.57e-06)	3.17e-06 (1.21e-06)	49.6% (7.0%)	.60	4.65e-06 (7.65e-07)	4.76e-06 (1.23e-06)	-1.9% (14.6%)
.65	2.32e-05 (4.84e-06)	2.98e-06 (7.76e-07)	87.2% (1.4%)	.65	6.06e-06 (1.88e-06)	3.29e-06 (1.46e-06)	46.5% (13.2%)	.65	4.46e-06 (7.68e-07)	4.74e-06 (1.38e-06)	-5.6% (17.7%)
.70	2.11e-05 (2.04e-06)	2.70e-06 (1.76e-07)	87.1% (1.5%)	.70	5.51e-06 (8.15e-07)	3.05e-06 (8.90e-07)	45.2% (9.0%)	.70	4.33e-06 (7.80e-07)	5.09e-06 (1.82e-06)	-15.6% (21.6%)
.75	2.06e-05 (1.23e-06)	2.62e-06 (1.33e-07)	87.3% (1.0%)	.75	5.63e-06 (1.87e-06)	3.32e-06 (1.90e-06)	41.4% (27.7%)	.75	4.16e-06 (6.60e-07)	4.61e-06 (1.16e-06)	-10.3% (14.5%)
.80	2.22e-05 (4.37e-06)	2.92e-06 (3.19e-07)	86.9% (1.4%)	.80	5.77e-06 (1.94e-06)	3.50e-06 (1.58e-06)	40.2% (13.6%)	.80	4.05e-06 (7.03e-07)	4.67e-06 (1.38e-06)	-14.5% (19.4%)
.85	2.19e-05 (3.68e-06)	2.94e-06 (7.48e-07)	86.6% (2.0%)	.85	5.69e-06 (1.58e-06)	3.42e-06 (1.37e-06)	40.8% (12.2%)	.85	3.91e-06 (8.13e-07)	4.60e-06 (1.47e-06)	-16.5% (16.1%)
.90	2.74e-05 (8.95e-06)	3.74e-06 (1.68e-06)	86.5% (2.6%)	.90	5.01e-06 (1.06e-06)	2.87e-06 (7.66e-07)	42.7% (8.9%)	.90	3.77e-06 (5.52e-07)	4.61e-06 (1.01e-06)	-21.8% (12.1%)

RQ2b. EnfPoly vs. MonPoly used as an enforcer: $N = 25$, $n = 10$, $L = 1000$

$d = 5$, $A = 16$, $n = 10$, $p = .50$

I	MonPoly s/event	EnfPoly s/event	Speedup %
1	1.90e-04 (2.17e-04)	3.36e-06 (5.58e-07)	88.5% (11.8%)
5	3.00e-04 (2.59e-04)	4.06e-06 (9.41e-07)	90.4% (12.2%)
10	2.98e-04 (2.63e-04)	4.08e-06 (1.16e-06)	90.1% (12.6%)
20	2.74e-04 (2.62e-04)	4.42e-06 (1.56e-06)	89.3% (12.7%)
50	2.68e-04 (2.49e-04)	4.71e-06 (2.59e-06)	90.3% (10.3%)
100	2.05e-04 (2.52e-04)	7.16e-06 (1.22e-06)	85.8% (15.6%)
200	2.17e-04 (2.31e-04)	5.65e-06 (4.12e-06)	88.6% (10.1%)
500	2.56e-04 (2.72e-04)	1.82e-05 (5.21e-05)	85.0% (14.1%)

$d = 5$, $I = 50$, $n = 10$, $p = .50$

A	MonPoly s/event	EnfPoly s/event	Speedup %
2	1.34e-04 (1.96e-04)	3.32e-06 (5.24e-07)	83.3% (10.8%)
4	1.92e-04 (2.60e-04)	3.54e-06 (7.41e-07)	85.4% (11.5%)
8	2.19e-04 (2.65e-04)	4.09e-06 (9.88e-07)	86.2% (11.8%)
16	2.68e-04 (2.49e-04)	4.71e-06 (2.59e-06)	90.3% (10.3%)
32	1.72e-04 (2.50e-04)	5.41e-06 (4.38e-06)	84.1% (10.4%)
64	3.46e-04 (2.66e-04)	6.69e-06 (4.97e-06)	90.2% (11.8%)
128	1.75e-04 (2.17e-04)	5.50e-06 (4.46e-06)	84.8% (12.0%)
256	1.69e-04 (2.41e-04)	5.37e-06 (4.13e-06)	82.7% (11.8%)

$I = 50$, $A = 16$, $n = 10$, $p = .50$

d	MonPoly s/event	EnfPoly s/event	Speedup %
2	3.44e-04 (2.50e-04)	2.60e-06 (3.12e-07)	94.3% (9.8%)
3	3.15e-04 (2.17e-04)	3.31e-06 (1.14e-06)	95.1% (8.6%)
4	2.28e-04 (2.55e-04)	3.34e-06 (7.80e-07)	88.2% (11.7%)
5	2.68e-04 (2.49e-04)	4.71e-06 (2.59e-06)	90.3% (10.3%)
6	2.59e-04 (2.79e-04)	5.47e-06 (3.38e-06)	88.0% (10.5%)
7	1.91e-04 (2.39e-04)	8.94e-06 (4.74e-06)	84.3% (11.4%)
8	2.48e-04 (2.81e-04)	1.56e-05 (1.10e-05)	82.7% (12.5%)

$d = 5$, $I = 50$, $A = 16$, $n = 10$

p	MonPoly s/event	EnfPoly s/event	Speedup %
0.00	1.17e-04 (1.27e-04)	3.31e-06 (1.19e-06)	87.7% (11.5%)
0.01	7.80e-04 (5.14e-04)	1.06e-04 (1.31e-05)	82.1% (7.5%)
0.05	6.46e-04 (7.35e-04)	2.47e-05 (5.20e-06)	87.0% (10.5%)
0.10	4.02e-04 (4.48e-04)	1.23e-05 (2.22e-06)	88.1% (9.8%)
0.25	3.36e-04 (3.82e-04)	8.33e-06 (9.35e-06)	87.6% (11.0%)
0.50	2.68e-04 (2.49e-04)	4.71e-06 (2.59e-06)	90.3% (10.3%)
0.75	2.35e-04 (1.94e-04)	3.62e-06 (1.57e-06)	90.1% (11.5%)
0.88	1.95e-04 (1.56e-04)	3.22e-06 (7.72e-07)	90.8% (11.2%)

$d = 5$, $I = 50$, $A = 16$, $p = .50$

n	MonPoly s/event	EnfPoly s/event	Speedup %
1	3.39e-04 (5.57e-04)	2.32e-05 (2.82e-06)	82.4% (7.8%)
2	4.61e-04 (6.81e-04)	1.35e-05 (3.74e-06)	86.0% (10.7%)
5	2.37e-04 (3.83e-04)	6.92e-06 (2.78e-06)	83.3% (10.6%)
10	2.68e-04 (2.49e-04)	4.71e-06 (2.59e-06)	90.3% (10.3%)
20	1.38e-04 (1.42e-04)	2.98e-06 (1.07e-06)	86.5% (13.1%)
50	5.62e-05 (5.80e-05)	2.57e-06 (1.53e-06)	80.3% (17.4%)
100	2.89e-05 (2.95e-05)	2.45e-06 (1.74e-06)	73.4% (21.4%)
200	1.64e-05 (1.47e-05)	2.52e-06 (2.98e-06)	68.3% (22.8%)