





TIMED AUTOMATA ROBUSTNESS ANALYSIS VIA MODEL CHECKING

JAROSLAV BENDÍK ^a, AHMET SENCAN ^b, EBRU AYDIN GOL ^b, AND IVANA ČERNÁ ^c

^a Max Planck Institute for Software Systems, Kaiserslautern, Germany
e-mail address: xbendik@mpi-sws.org

^b Department of Computer Engineering, Middle East Technical University, Ankara, Turkey
e-mail address: {sencan.ahmet, ebrugol}@metu.edu.tr

^c Faculty of Informatics, Masaryk University, Brno, Czech Republic
e-mail address: cerna@fi.muni.cz

ABSTRACT. Timed automata (TA) have been widely adopted as a suitable formalism to model time-critical systems. Furthermore, contemporary model-checking tools allow the designer to check whether a TA complies with a system specification. However, the exact timing constants are often uncertain during the design phase. Consequently, the designer is often able to build a TA with a correct structure, however, the timing constants need to be tuned to satisfy the specification. Moreover, even if the TA initially satisfies the specification, it can be the case that just a slight perturbation during the implementation causes a violation of the specification. Unfortunately, model-checking tools are usually not able to provide any reasonable guidance on how to fix the model in such situations. In this paper, we propose several concepts and techniques to cope with the above mentioned design phase issues when dealing with reachability and safety specifications.

1. INTRODUCTION

Timed automata (TA) [AD94] extend finite automata with a set of real-time variables, called clocks. The clocks enrich the semantics and the constraints on the clocks restrict the behavior of the automaton, which are particularly important in modeling time-critical systems. The examples of TA models of critical systems include scheduling of real-time systems [Feh99, DILS09, GGD⁺07], medical devices [KMPP15, JPAM14], rail-road crossing systems [Wan04] and home-care plans [GBST14].

Model-checking methods allow for verifying whether a given TA meets a given system specification. Contemporary model-checking tools, such as UPPAAL [BDL⁺06] or Imitator [AFKS12], have proved to be practically applicable on various industrial case studies [BDL⁺06, AFMS19, HPW01]. Unfortunately, during the system design phase, the system information is often incomplete. A designer is often able to build a TA with correct structure, i.e., exactly capturing locations and transitions of the modeled system, however the exact clock (timing) constraints that enable/trigger the transitions can be uncertain.

Key words and phrases: Timed Automata, Design, Reachability, Safety.

Consequently, the produced TA often might not meet the specification (i.e., it does not pass the model-checking) and it needs to be fixed. On the other hand, even though the TA meets the specification, some of its constraints might unnecessarily restrict its behavior or the specification might get violated with a few changes over the constraints (e.g. threshold change, constraint removal). In this paper, we present methods to analyze and tune the constraints of the timed automaton to address these issues for reachability and safety specifications. In particular, we identify a minimal set of constraints that needs to be removed to satisfy a reachability specification and we find a maximal set of constraints whose removal do not result in violation of a safety specification. In both cases, we further analyze the thresholds that appear in these constraints.

Tuning TA for a Reachability Specification. In model-checking, if the considered specification declares an existential property, such as reachability, the property has to hold on a trace of the TA. If the property holds, the model checker returns “yes” and generates a witness trace satisfying the property. However, if the property does not hold, the model checker usually returns just “no” and does not provide any additional information that would help the designer to correct the TA. In this paper, we first study the following problem: given a timed automaton \mathcal{A} and a reachability property that is not satisfied by \mathcal{A} , relax clock constraints of \mathcal{A} such that the resultant automaton \mathcal{A}' satisfies the reachability property. Moreover, the goal is to minimize the number of the relaxed clock constraints and, secondary, also to minimize the overall change of the timing constants used in the clock constraints. We propose a two step solution for this problem. In the first step, we identify a *minimal sufficient reduction (MSR)* of \mathcal{A} , i.e., an automaton \mathcal{A}'' that satisfies the reachability property and originates from \mathcal{A} by removing only a minimal necessary set of clock constraints. In the second step, instead of completely removing the clock constraints, we relax the constraint thresholds. We present two methods for this purpose. First, we employ mixed integer linear programming (MILP) to find a minimal relaxation of the constraints that leads to a satisfaction of the reachability property along a witness path. As the second method, we parametrize the identified constraints and use a parameter synthesis tool to find a minimal parameter valuation such that the target set becomes reachable. The second method is guaranteed to find the globally optimal relaxation as it considers all witness paths, while the first one is more efficient. We thoroughly compare both the methods on a case study.

The underlying assumption is that during the design the most suitable timing constants reflecting the system properties are defined. Thus, our goal is to generate a TA satisfying the reachability property by changing a minimum number of timing constants. Some of the constraints of the initial TA can be strict (no relaxation is possible), which can easily be integrated to the proposed solution. Thus, the proposed method can be viewed as a way to handle design uncertainties: develop a TA \mathcal{A} in a best-effort basis and apply our algorithm to find a \mathcal{A}' that is *as close as possible* to \mathcal{A} and satisfies the given reachability property.

Tuning TA for a Safety Specification. On contrary to existential properties, a universal property, e.g., safety or unavailability, needs to hold on each trace of the TA. If a safety specification does not hold for the TA, the model-checker returns “no” and generates a trace along which the property is violated. In recent studies, such traces are used to repair the model in an automated way [KLW19, EYG21]. In the other case, when the safety property holds, the “yes” answer obtained from a model-checker simply states that the TA does not have a trace violating the specification. However, the “yes” answer does not provide further

information on which constraints are effective in avoiding unsafe behaviors and how the verification result changes if some of the constraints are relaxed or removed. The second problem we study aims to provide additional information on a positive verification result for a safety specification described as avoiding a set of “unsafe” locations. In particular, we study the following problem: given a timed automaton \mathcal{A} and a safety property that is satisfied by \mathcal{A} , remove and/or relax the clock constraints of \mathcal{A} such that the resultant automaton \mathcal{A}' still satisfies the safety property. Here, our primary goal is to minimize the number of constraints that need to be left in the TA to prevent reaching the unsafe locations. Equivalently, we maximize the number of constraints that can be removed from the TA while keeping the unsafe locations unreachable. Our secondary goal is to maximize the total change in the timing constants used in the remaining clock constraints, where we consider two scenarios: (1) maximize the total change (as in the reachability case) (2) relax each clock constraint with the same amount and maximize this amount. Again, we present a two step solution to the considered problem. In the first step, we identify a *minimal guarantee (MG)* of \mathcal{A} that is a minimal set of constraints that need to be left in the automaton to ensure that the unsafe locations are still unreachable. In other words the automaton \mathcal{A}'' obtained by removing the constraints that are not in the MG still satisfies the safety specification and removing any additional constraint results in a violation. In the second step, we relax the thresholds of the constraints from the MG (i.e. clock constraints of \mathcal{A}''). For both of the aforementioned relaxation scenarios, we parametrize the constraints of \mathcal{A}'' and employ a parameter synthesis tool.

The methods we develop to solve the second problem allows us to relax the TA as much as possible without violating the safety specification. In general, during the design of the automaton, redundant constraints can be added unintentionally to ensure safety. The results of our analysis allows the designer to identify and remove such unnecessary constraints. Furthermore, the constraint constants can be too tight unnecessarily restricting the set of possible behaviors of the automaton. The results obtained in the second step help the designer to relax such constants. On the other hand, if it is not possible to further relax the constraints in the MG, small perturbations results in violation of the specification, in which case, the designer might choose to further restrict some of the constraint constants from the MG. Consequently, the developed method is intended to assist the designer to improve the model that is generated in a best-effort manner.

The proposed approach for tuning a TA for reachability specifications first appeared in [BSGČ21]. This paper extends [BSGČ21] by introducing the minimal guarantee concepts and the corresponding methods to (i) generate MG and (ii) the corresponding relaxations for tuning TA for safety specifications.

Outline. The rest of the paper is organized as follows. Section 2 introduces basic concepts used throughout the paper and formally defines the problems we deal with. Subsequently, in Sections 3 and 4, we describe our approaches for identifying Minimal Sufficient Reductions (MSRs) and Minimal Guarantees (MGs), respectively. In Section 6, we describe how to just relax timing constraints in an MSR instead of completely removing the constraints from the TA. Similarly, in Section 7 we show how to further relax an MG via parameter synthesis. In Section 8, we provide an overview of related work. Finally, we experimentally evaluate the proposed techniques in Section 9, and conclude in Section 10.

2. PRELIMINARIES

2.1. Timed Automata. A *timed automaton* (TA) [Alu99, AD94, LY93] is a finite-state machine extended with a finite set C of real-valued clocks. A *clock* $x \in C$ measures the time spent after its last reset. In a TA, clock constraints are defined for locations (states) and transitions. A *simple clock constraint* is defined as $x - y \sim c$ where $x, y \in C \cup \{0\}$, $\sim \in \{<, \leq\}$ and $c \in \mathbb{Z} \cup \{\infty\}$.¹ Simple clock constraints and constraints obtained by combining these with the conjunction operator (\wedge) are called *clock constraints*. The sets of simple and all clock constraints are denoted by $\Phi_S(C)$ and $\Phi(C)$, respectively. For a clock constraint $\phi \in \Phi(C)$, $\mathcal{S}(\phi)$ denotes the simple constraints from ϕ , e.g., $\mathcal{S}(x - y < 10 \wedge y \leq 20) = \{x - y < 10, y \leq 20\}$. A clock constraint is called *parametric* if the numerical constant (i.e. c) is represented by a parameter. A *clock valuation* $v : C \rightarrow \mathbb{R}_+$ assigns non-negative real values to each clock. The notation $v \models \phi$ denotes that the clock constraint ϕ evaluates to true when each clock x is replaced with $v(x)$. For a clock valuation v and $d \in \mathbb{R}_+$, $v + d$ is the clock valuation obtained by *delaying* each clock by d , i.e., $(v + d)(x) = v(x) + d$ for each $x \in C$. For $\lambda \subseteq C$, $v[\lambda := 0]$ is the clock valuation obtained after *resetting* each clock from λ , i.e., $v[\lambda := 0](x) = 0$ for each $x \in \lambda$ and $v[\lambda := 0](x) = v(x)$ for each $x \in C \setminus \lambda$.

Definition 2.1 (Timed Automata). A *timed automaton* $\mathcal{A} = (L, l_0, C, \Delta, Inv)$ is a tuple, where L is a finite set of locations, $l_0 \in L$ is the initial location, C is a finite set of clocks, $\Delta \subseteq L \times 2^C \times \Phi(C) \times L$ is a finite transition relation, and $Inv : L \rightarrow \Phi(C)$ is an invariant function.

For a transition $e = (l_s, \lambda, \phi, l_t) \in \Delta$, l_s is the source location, l_t is the target location, λ is the set of clocks reset on e and ϕ is the guard (i.e., a clock constraint) tested for enabling e . A Parametric TA (PTA) extends TA by allowing the use of parametric constraints. Given a PTA \mathcal{A} with parameter set P and a parameter valuation $\mathbf{p} : P \rightarrow \mathbb{N}$, a (non-parametric) TA $\mathcal{A}(\mathbf{p})$ is obtained by replacing each parameter $p \in P$ in \mathcal{A} with the corresponding valuation $\mathbf{p}(p)$. The semantics of a TA is given by a labelled transition system (LTS). An LTS is a tuple $\mathcal{T} = (S, s_0, \Sigma, \rightarrow)$, where S is a set of states, $s_0 \in S$ is an initial state, Σ is a set of symbols, and $\rightarrow \subseteq S \times \Sigma \times S$ is a transition relation. A transition $(s, a, s') \in \rightarrow$ is also denoted as $s \xrightarrow{a} s'$.

Definition 2.2 (LTS semantics for TA). Given a TA $\mathcal{A} = (L, l_0, C, \Delta, Inv)$, the labelled transition system $T(\mathcal{A}) = (S, s_0, \Sigma, \rightarrow)$ is defined as follows:

- $S = \{(l, v) \mid l \in L, v \in \mathbb{R}_+^{|C|}, v \models Inv(l)\}$,
- $s_0 = (l_0, \mathbf{0})$, where $\mathbf{0}(x) = 0$ for each $x \in C$,
- $\Sigma = \{act\} \cup \mathbb{R}_+$, and
- the transition relation \rightarrow is defined by the following rules:
 - delay transition: $(l, v) \xrightarrow{d} (l, v + d)$ if $v + d \models Inv(l)$
 - discrete transition: $(l, v) \xrightarrow{act} (l', v')$ if there exists $(l, \lambda, \phi, l') \in \Delta$ such that $v \models \phi$, $v' = v[\lambda := 0]$, and $v' \models Inv(l')$.

The notation $s \rightarrow_d s'$ is used to denote a delay transition of duration d followed by a discrete transition from s to s' , i.e., $s \xrightarrow{d} s \xrightarrow{act} s'$. A run ρ of \mathcal{A} is either a finite or an infinite

¹Simple constraints are only defined as upper bounds to simplify the presentation. This definition is not restrictive since $x - y \geq c$ and $x \geq c$ are equivalent to $y - x \leq -c$ and $0 - x \leq -c$, respectively. A similar argument holds for strict inequality ($>$).




FIGURE 1. An illustration of a TA used in Examples 2.3, 2.13 and 2.14.

alternating sequence of delay and discrete transitions, i.e., $\rho = s_0 \rightarrow_{d_0} s_1 \rightarrow_{d_1} s_2 \rightarrow_{d_2} \dots$. The set of all runs of \mathcal{A} is denoted by $[[\mathcal{A}]]$.

A path π of \mathcal{A} is an interleaving sequence of locations and transitions, $\pi = l_0, e_1, l_1, e_2, \dots$, where $e_{i+1} = (l_i, \lambda_{i+1}, \phi_{i+1}, l_{i+1}) \in \Delta$ for each $i \geq 0$. A path $\pi = l_0, e_1, l_1, e_2, \dots$ is *realizable* if there exists a delay sequence d_0, d_1, \dots such that $(l_0, \mathbf{0}) \rightarrow_{d_0} (l_1, v_1) \rightarrow_{d_1} (l_1, v_2) \rightarrow_{d_2} \dots$ is a run of \mathcal{A} and for every $i \geq 1$, the i th discrete transition is taken according to e_i , i.e., $e_i = (l_{i-1}, \lambda_i, \phi_i, l_i)$, $v_{i-1} + d_{i-1} \models \phi_i$, $v_i = (v_{i-1} + d_{i-1})[\lambda_i := 0]$ and $v_i \models \text{Inv}'(l_i)$.

For a TA \mathcal{A} and a subset of its locations $L_T \subseteq L$, L_T is said to be *reachable* on \mathcal{A} if there exists $\rho = (l_0, \mathbf{0}) \rightarrow_{d_0} (l_1, v_1) \rightarrow_{d_1} \dots \rightarrow_{d_{n-1}} (l_n, v_n) \in [[\mathcal{A}]]$ such that $l_n \in L_T$; otherwise, L_T is *unreachable*. In this study, L_T is used to denote the set of *target locations* for reachability specifications and the set of *unsafe locations* for safety specifications. In the latter case, \mathcal{A} is called *safe* if L_T is unreachable; otherwise \mathcal{A} is *unsafe*. The reachability problem, $\text{isReachable}(\mathcal{A}, L_T)$, is decidable and implemented in various verification tools including UPPAAL [BDL⁺06]. The verifier either returns “No” indicating that such a run does not exist, or it generates a run (counter-example) leading from the initial state of \mathcal{A} to a location $v_n \in L_T$.

Example 2.3. In Figure 1, we illustrate a TA with 8 locations: $\{l_0, \dots, l_7\}$, 9 transitions: $\{e_1, \dots, e_9\}$, an initial location l_0 , and a set of unreachable locations $L_T = \{l_4\}$.

2.2. Timed Automata Relaxation. For a timed automaton $\mathcal{A} = (L, l_0, C, \Delta, \text{Inv})$, the set of pairs of transition and associated simple constraints is defined in (2.1) and the set of pairs of location and associated simple constraints is defined in (2.2).

$$\Psi(\Delta) = \{(e, \varphi) \mid e = (l_s, \lambda, \phi, l_t) \in \Delta, \varphi \in \mathcal{S}(\phi)\} \quad (2.1)$$

$$\Psi(\text{Inv}) = \{(l, \varphi) \mid l \in L, \varphi \in \mathcal{S}(\text{Inv}(l))\} \quad (2.2)$$

Definition 2.4 (constraint-relaxation). Let $\phi \in \Phi(C)$ be a constraint over C , $\Theta \subseteq \mathcal{S}(\phi)$ be a subset of its simple constraints and $\mathbf{r} : \Theta \rightarrow \mathbb{N} \cup \{\infty\}$ be a positive valued relaxation

valuation. The relaxed constraint is defined as:

$$R(\phi, \Theta, \mathbf{r}) = \left(\bigwedge_{\varphi \in \mathcal{S}(\phi) \setminus \Theta} \varphi \right) \wedge \left(\bigwedge_{\varphi = x - y \sim c + \mathbf{r}(\varphi)} x - y \sim c + \mathbf{r}(\varphi) \right) \quad (2.3)$$

Intuitively, $R(\phi, \Theta, \mathbf{r})$ relaxes only the thresholds of simple constraints from Θ with respect to \mathbf{r} , e.g., $R(x - y \leq 10 \wedge y < 20, \{y < 20\}, \mathbf{r}) = x - y \leq 10 \wedge y < 23$, where $\mathbf{r}(y < 20) = 3$. Setting a threshold to ∞ implies removing the corresponding simple constraint, e.g., $R(x - y \leq 10 \wedge y < 20, \{y < 20\}, \mathbf{r}) = x - y \leq 10$, where $\mathbf{r}(y < 20) = \infty$. Note that $R(\phi, \Theta, \mathbf{r}) = \phi$ when Θ is empty.

Definition 2.5 ((D, I, \mathbf{r}) -relaxation). Let $\mathcal{A} = (L, l_0, C, \Delta, Inv)$ be a TA, $D \subseteq \Psi(\Delta)$ and $I \subseteq \Psi(Inv)$ be transition and location constraint sets, and $\mathbf{r} : D \cup I \rightarrow \mathbb{N} \cup \{\infty\}$ be a positive valued relaxation valuation. The (D, I, \mathbf{r}) -relaxation of \mathcal{A} , denoted $\mathcal{A}_{\langle D, I, \mathbf{r} \rangle}$, is a TA $\mathcal{A}' = (L', l'_0, C', \Delta', Inv')$ such that:

- $L = L'$, $l_0 = l'_0$, $C = C'$, and
- Δ' originates from Δ by relaxing D via \mathbf{r} . For $e = (l_s, \lambda, \phi, l_t) \in \Delta$, let $D|_e = \{\varphi \mid (e, \varphi) \in D\}$, and let $\mathbf{r}|_e(\varphi) = \mathbf{r}(e, \varphi)$, then $\Delta' = \{(l_s, \lambda, R(\phi, D|_e, \mathbf{r}|_e), l_t) \mid e = (l_s, \lambda, \phi, l_t) \in \Delta\}$
- Inv' originates from Inv by relaxing I via \mathbf{r} . For $l \in L$, let $I|_l = \{\varphi \mid (l, \varphi) \in I\}$, and $\mathbf{r}|_l(\varphi) = \mathbf{r}(l, \varphi)$, then $Inv'(l) = R(Inv(l), I|_l, \mathbf{r}|_l)$.

Intuitively, the TA $\mathcal{A}_{\langle D, I, \mathbf{r} \rangle}$ emerges from \mathcal{A} by relaxing the guards of the transitions from the set D and relaxing invariants of the locations from I with respect to \mathbf{r} .

Proposition 2.6. *Let $\mathcal{A} = (L, l_0, C, \Delta, Inv)$ be a timed automaton, $D \subseteq \Psi(\Delta)$ and $I \subseteq \Psi(Inv)$ be sets of simple guard and invariant constraints, and $\mathbf{r} : D \cup I \rightarrow \mathbb{N} \cup \{\infty\}$ be a relaxation valuation. Then $[[\mathcal{A}]] \subseteq [[\mathcal{A}_{\langle D, I, \mathbf{r} \rangle}]]$.*

Proof. Observe that for a clock constraint $\phi \in \Phi(C)$, a subset of its simple constraints $\Theta \subseteq \mathcal{S}(\phi)$, a relaxation valuation \mathbf{r}' for Θ , and the relaxed constraint $R(\phi, \Theta, \mathbf{r}')$ as in Definition 2.4, it holds that for any clock valuation $v : v \models \phi \implies v \models R(\phi, \Theta, \mathbf{r}')$. Now, consider a run $\rho = (l_0, \mathbf{0}) \xrightarrow{d_0} (l_1, v_1) \xrightarrow{d_1} (l_2, v_2) \xrightarrow{d_2} \dots \in [[\mathcal{A}]]$. Let $\pi = l_0, e_1, l_1, e_2, \dots$ with $e_i = (l_{i-1}, \lambda_i, \phi_i, l_i) \in \Delta$ for each $i \geq 1$ be the path realized as ρ via delay sequence d_0, d_1, \dots . By Definition 2.5 for each $(l, \lambda, \phi, l') \in \Delta$, there is $(l, \lambda, R(\phi, D|_e, \mathbf{r}|_e), l') \in \Delta'$. We define a path induced by π on $\mathcal{A}_{\langle D, I, \mathbf{r} \rangle}$ as:

$$M(\pi) = l_0, (l_0, \lambda_1, R(\phi_1, D|_{e_1}, \mathbf{r}|_{e_1}), l_1), l_1, (l_1, \lambda_2, R(\phi_2, D|_{e_2}, \mathbf{r}|_{e_2}), l_2), \dots \quad (2.4)$$

For each $i = 0, \dots, n-1$ it holds that $v_i \models R(Inv(l_i), D|_{l_i}, \mathbf{r}|_{l_i})$, $v_i + d_i \models R(Inv(l_i), D|_{l_i}, \mathbf{r}|_{l_i})$ and $v_i + d_i \models R(\phi_{i+1}, D|_{e_{i+1}}, \mathbf{r}|_{e_{i+1}})$. Thus $M(\pi)$ is realizable on $\mathcal{A}_{\langle D, I, \mathbf{r} \rangle}$ via the same delay sequence and $\rho \in [[\mathcal{A}_{\langle D, I, \mathbf{r} \rangle}]]$. As $\rho \in [[\mathcal{A}]]$ is arbitrary, we conclude that $[[\mathcal{A}]] \subseteq [[\mathcal{A}_{\langle D, I, \mathbf{r} \rangle}]]$. \square

2.3. Reductions and Guarantees.

Definition 2.7. A *reduction* is a relaxation $\mathcal{A}_{\langle D, I, \mathbf{r} \rangle}$ of \mathcal{A} such that $\mathbf{r}(a) = \infty$ for each $a \in D \cup I$. Moreover, since \mathbf{r} is fixed, we simply denote the reduction by $\mathcal{A}_{\langle D, I \rangle}$.

Intuitively, a reduction $\mathcal{A}_{\langle D, I \rangle}$ effectively removes all the simple constraints $D \cup I$ from \mathcal{A} . Also, note that $\mathcal{A} = \mathcal{A}_{\langle \emptyset, \emptyset \rangle}$. Hereafter, we use two notations for naming a reduction; either we simply use capital letters, e.g., M, N, K to name a reduction, or we

use the notation $\mathcal{A}_{\langle D, I \rangle}$ to also specify the sets D, I of simple clock constraints. Given a reduction $N = \mathcal{A}_{\langle D, I \rangle}$, $|N|$ denotes the cardinality $|D \cup I|$. Furthermore, $\mathcal{R}_{\mathcal{A}}$ denotes the set of all reductions of \mathcal{A} . We define a partial order relation \sqsubseteq on $\mathcal{R}_{\mathcal{A}}$ as $\mathcal{A}_{\langle D, I \rangle} \sqsubseteq \mathcal{A}_{\langle D', I' \rangle}$ iff $D \cup I \subseteq D' \cup I'$. Similarly, we write $\mathcal{A}_{\langle D, I \rangle} \sqsubset \mathcal{A}_{\langle D', I' \rangle}$ iff $D \cup I \subsetneq D' \cup I'$. We say that a reduction $\mathcal{A}_{\langle D, I \rangle}$ is a *sufficient reduction* (w.r.t. \mathcal{A} and L_T) iff L_T is reachable on $\mathcal{A}_{\langle D, I \rangle}$; otherwise, $\mathcal{A}_{\langle D, I \rangle}$ is an *insufficient reduction*. Crucially, observe that the property of being a sufficient reduction is monotone w.r.t. the partial order:

Proposition 2.8. *Let $\mathcal{A}_{\langle D, I \rangle}$ and $\mathcal{A}_{\langle D', I' \rangle}$ be reductions such that $\mathcal{A}_{\langle D, I \rangle} \sqsubseteq \mathcal{A}_{\langle D', I' \rangle}$. If $\mathcal{A}_{\langle D, I \rangle}$ is sufficient then $\mathcal{A}_{\langle D', I' \rangle}$ is also sufficient.*

Proof. Note that $\mathcal{A}_{\langle D', I' \rangle}$ is a $(D' \setminus D, I' \setminus I)$ -reduction of $\mathcal{A}_{\langle D, I \rangle}$. By Proposition 2.6, $[[\mathcal{A}_{\langle D, I \rangle}]] \subseteq [[\mathcal{A}_{\langle D', I' \rangle}]]$, i.e., the run of $\mathcal{A}_{\langle D, I \rangle}$ that witnesses the reachability of L_T is also a run of $\mathcal{A}_{\langle D', I' \rangle}$. \square

Definition 2.9 (MSR). A sufficient reduction $\mathcal{A}_{\langle D, I \rangle}$ is a *minimal sufficient reduction* (MSR) iff there is no $c \in D \cup I$ such that the reduction $\mathcal{A}_{\langle D \setminus \{c\}, I \setminus \{c\} \rangle}$ is sufficient. Equivalently, due to Proposition 2.8, $\mathcal{A}_{\langle D, I \rangle}$ is an MSR iff there is no sufficient reduction $\mathcal{A}_{\langle D', I' \rangle}$ such that $\mathcal{A}_{\langle D', I' \rangle} \sqsubset \mathcal{A}_{\langle D, I \rangle}$.

Definition 2.10 (MIR). An insufficient reduction $\mathcal{A}_{\langle D, I \rangle}$ is a *maximal insufficient reduction* (MIR) iff there is no $c \in (\Psi(\Delta) \cup \Psi(Inv)) \setminus (D \cup I)$ such that the reduction $\mathcal{A}_{\langle D', I' \rangle}$ with $D' \cup I' = D \cup I \cup \{c\}$ is insufficient. Equivalently, due to Proposition 2.8, $\mathcal{A}_{\langle D, I \rangle}$ is an MIR iff there is no insufficient reduction $\mathcal{A}_{\langle D'', I'' \rangle}$ such that $\mathcal{A}_{\langle D, I \rangle} \sqsubset \mathcal{A}_{\langle D'', I'' \rangle}$.

Intuitively, an MSR represents a minimal set of constraints that need to be removed from \mathcal{A} to make the target location(s) L_T reachable, whereas an MIR represents a maximal set of constraints whose removal does not make the target location(s) reachable.

Recall that a reduction $\mathcal{A}_{\langle D, I \rangle}$ is determined by $D \subseteq \Psi(\Delta)$ and $I \subseteq \Psi(Inv)$. Consequently, $|\mathcal{R}_{\mathcal{A}}| = 2^{|\Psi(\Delta) \cup \Psi(Inv)|}$ (i.e., there are exponentially many reductions w.r.t. $|\Psi(\Delta) \cup \Psi(Inv)|$). Moreover, there can be up to $\binom{k}{k/2}$ MSRs (MIRs) where $k = |\Psi(\Delta) \cup \Psi(Inv)|$.² Also note, that the *minimality* (*maximality*) of a reduction does not mean a *minimum* (*maximum*) number of simple clock constraints that are removed by the reduction; there can exist two MSRs (MIRs), M and N , such that $|M| < |N|$. We call an MSR M a *minimum MSR* if there is no MSR M' with $|M'| < |M|$. Similarly, an MIR M is a *maximum MIR* if there is no MIR M' with $|M'| > |M|$. Note that there can be also up to $\binom{k}{k/2}$ minimum MSRs and up to $\binom{k}{k/2}$ maximum MIRs.

In some applications, instead of thinking about an MIR of \mathcal{A} , i.e. a maximal set of simple clock constraints whose removal does not make the target location(s) reachable, it might be more natural to think about the complement of an MIR, i.e. a *minimal set of simple clock constraints* that need to be left in \mathcal{A} to ensure that the target location is still unreachable. We define this complementary notion as a *minimal guarantee*:

Definition 2.11 (MG). Given a reduction $\mathcal{A}_{\langle D, I \rangle}$, the set $D \cup I$ of simple clock constraints constitutes a *guarantee* (for \mathcal{A}) iff the reduction $\mathcal{A}_{\langle \Psi(\Delta) \setminus D, \Psi(Inv) \setminus I \rangle}$ is insufficient. Furthermore, a guarantee $D \cup I$ is a *minimal guarantee* (MG) iff for every $c \in D \cup I$ the reduction

²There are $\binom{k}{k/2}$ pair-wise incomparable elements of $\mathcal{R}_{\mathcal{A}}$ w.r.t. \sqsubset (see Sperner's theorem [Spe28]) and all of them can be MSRs (or MIRs).




FIGURE 2. An illustration of the set of all TA reductions from Example 2.14. We denote individual reductions of \mathcal{A} using a bit-vector representation; for instance, 1101 represents the reduction $\mathcal{A}_{\langle D, I \rangle}$ where $D \cup I = \{c_1, c_2, c_4\}$. The reductions with a red dashed border are the insufficient reductions, and the reductions with solid green border are sufficient reductions. The MRSes and MIRs are filled with a background color.

$\mathcal{A}_{\langle \Psi(\Delta) \setminus (D \setminus \{c\}), \Psi(Inv) \setminus (I \setminus \{c\}) \rangle}$ is sufficient. Equivalently, due to Proposition 2.8, $D \cup I$ is an MG iff there is no guarantee $D' \cup I'$ such that $D' \cup I' \subsetneq D \cup I$.

Observation 2.12. A reduction $\mathcal{A}_{\langle D, I \rangle}$ is insufficient iff the set $(\Psi(\Delta) \cup \Psi(Inv)) \setminus (D \cup I)$ is a guarantee. Furthermore, $\mathcal{A}_{\langle D, I \rangle}$ is an MIR iff $(\Psi(\Delta) \cup \Psi(Inv)) \setminus (D \cup I)$ is an MG.

Note that due to technical reasons, we define the concept of an MG as a *set* $D \cup I$ of *simple clock constraints*, whereas the concept of an MIR is defined as a *reduction* $\mathcal{A}_{\langle D', I' \rangle}$ (i.e., a TA) that is determined by a set $D' \cup I'$ of simple clock constraints.

Example 2.13. Assume the TA \mathcal{A} and $L_T = \{l_4\}$ from Example 2.3 (Fig. 1). There are 24 MSRs and 4 of them are minimum. For example, $\mathcal{A}_{\langle D, I \rangle}$ with $D = \{(e_5, x \geq 25)\}$ and $I = \{(l_3, u \leq 26)\}$ is a minimum MSR, and $\mathcal{A}_{\langle D', I' \rangle}$ with $D' = \{(e_9, y \leq 15), (e_7, z \leq 15)\}$ and $I' = \{(l_6, x \leq 10)\}$ is a non-minimum MSR. There are 40 MGs (and hence MIRs) and 21 of them are minimum. For instance, $D \cup I$ with $D = \{(e_5, z \geq 9), (e_8, u \geq 35), (e_1, x \geq 9), (e_4, z \geq 9)\}$ and $I = \{(l_5, x \leq 10), (l_6, x \leq 10), (l_0, z \leq 10), (l_3, u \leq 26)\}$ is a non-minimum MG, and $D' \cup I'$ with $D' = \{(e_7, x \geq 9), (e_9, y \leq 15), (e_8, x \geq 9), (e_5, x \geq 25)\}$ and $I' = \{(l_1, x \leq 14), (l_3, z \leq 10)\}$ is a minimum MG.

Finally, note that in some situations, we might not want to include the whole set $\Psi(\Delta) \cup \Psi(Inv)$ of all simple clock constraints in the analysis but rather just its subset (e.g., because some simple clock constraints simply could not be modified). Our definitions of (D, I, \mathbf{r}) -relaxations, reductions, and guarantees, can be naturally extended also to work with just a subset of $\Psi(\Delta) \cup \Psi(Inv)$. We illustrate this on a simple example.

Example 2.14. Assume that only 4 of the simple clock constraints from the TA in Fig. 1 can be removed/relaxed and the other simple clock constraints represent physical limitations that can not be changed. The tunable simple clock constraints are $c_1 = x \geq 9$, $c_2 = z \geq 9$, $c_3 = x \leq 14$ and $c_4 = u \leq 26$ that appear on edge e_1 , edge e_5 , location l_1 , and location l_3 , respectively. Note that these constraints are highlighted using green color in Fig. 1. If we restrict our analysis only to those four constraints, then there are 2 MSR: $\{c_3, c_4\}$ and $\{c_1, c_2, c_3\}$, and three MGs: $\{c_3\}$, $\{c_1, c_4\}$ and $\{c_2, c_4\}$. We provide a power-set illustration of this example in Fig. 2

Algorithm 1: Minimum MSR Extraction Scheme

```

1  $N \leftarrow \mathcal{A}_{\langle \Psi(\Delta), \Psi(Inv) \rangle}; \mathcal{I} \leftarrow \emptyset; \mathcal{S} \leftarrow \emptyset$ 
2 while  $N \neq \text{null}$  do
3    $M, \mathcal{I} \leftarrow \text{shrink}(N, \mathcal{I})$  // Algorithm 2
4    $\mathcal{S} \leftarrow \mathcal{S} \cup \{X \mid M \sqsubseteq X\}$ 
5    $\mathcal{I} \leftarrow \mathcal{I} \cup \{Y \mid Y \sqsupseteq M\}$ 
6    $N, \mathcal{I}, \mathcal{S} \leftarrow \text{findSSeed}(M, \mathcal{M}_{msr}, \mathcal{I}, \mathcal{S})$  // Algorithm 3
7 return  $M$ 

```

2.4. Problem Formulations. In this paper, we are mainly concerned with the following two problems. The first problem and the proposed solution were presented in our conference paper [BSGČ21].

Problem 2.15. Given a TA $\mathcal{A} = (L, l_0, C, \Delta, Inv)$ and a set of target locations $L_T \subset L$ that is unreachable on \mathcal{A} , find a *minimal* (D, I, \mathbf{r}) -relaxation $\mathcal{A}_{\langle D, I, \mathbf{r} \rangle}$ of \mathcal{A} such that L_T is reachable on $\mathcal{A}_{\langle D, I, \mathbf{r} \rangle}$. In particular, the goal is to identify a (D, I, \mathbf{r}) -relaxation that minimizes the number $|D \cup I|$ of relaxed constraints, and, secondly, we tend to minimize the overall change of the clock constraints $\sum_{c \in D \cup I} \mathbf{r}(c)$.

Our solution to Problem 2.15 is described in detail in Sections 3 and 6. Briefly, we solve Problem 2.15 in two steps. First, we identify a minimum MSR $\mathcal{A}_{\langle D, I \rangle}$ for \mathcal{A} , i.e., a minimal set $D \cup I$ of simple clock constraints whose removal from \mathcal{A} makes the target locations L_T reachable. Second, instead of completely removing the constraints, we turn the MSR $\mathcal{A}_{\langle D, I \rangle}$ into the resultant (D, I, \mathbf{r}) -relaxation. To construct the (D, I, \mathbf{r}) -relaxation, we propose two alternative approaches: (1) an approach based on Mixed Integer Linear Programming (MILP) and (2) an approach based on parameter synthesis for PTA.

Problem 2.16. Given a TA $\mathcal{A} = (L, l_0, C, \Delta, Inv)$ and a set of target locations $L_T \subset L$ that is unreachable on \mathcal{A} , find a *maximal* (D, I, \mathbf{r}) -relaxation $\mathcal{A}_{\langle D, I, \mathbf{r} \rangle}$ of \mathcal{A} such that L_T is still unreachable on $\mathcal{A}_{\langle D, I, \mathbf{r} \rangle}$. In particular, the goal is to identify a (D, I, \mathbf{r}) -relaxation that maximizes the number $|\{c \in D \cup I \mid \mathbf{r}(c) = \infty\}|$ of constraints that are completely removed, and, secondary, maximizes the overall change of the clock constraints $\sum_{c \in \{c' \in D \cup I \mid \mathbf{r}(c') \neq \infty\}} \mathbf{r}(c)$ that are not completely removed.

Our solution to Problem 2.16 is presented in Sections 4 and 7. Briefly, we first identify a minimum MG $D' \cup I'$ for \mathcal{A} , i.e., a minimal set $D' \cup I'$ of simple clock constraints that need to be left in \mathcal{A} to ensure that the target location L_T is still unreachable. Subsequently, we employ parameter synthesis to further relax (as much as possible) the constraints $D' \cup I'$ that are left in the system. For both of the considered problems, we assume that there is a path from the initial state to the target set L_T (unrealizable since L_T is not reachable). Thus, the target set can become reachable via constraint removals/relaxations.

3. FINDING MINIMAL SUFFICIENT REDUCTIONS

In this section, we gradually describe our approach for finding a minimum minimal sufficient reduction.

Algorithm 2: $\text{shrink}(\mathcal{A}_{\langle D, I \rangle}, \mathcal{I})$

```

1  $X \leftarrow \emptyset$ 
2 while  $(D \cup I) \neq X$  do
3    $c \leftarrow$  pick a simple clock constraint from  $(D \cup I) \setminus X$ 
4   if  $\mathcal{A}_{\langle D \setminus \{c\}, I \setminus \{c\} \rangle} \notin \mathcal{I}$  and  $\mathcal{A}_{\langle D \setminus \{c\}, I \setminus \{c\} \rangle}$  is sufficient then
5      $\rho \leftarrow$  a witness run of the sufficiency of  $\mathcal{A}_{\langle D \setminus \{c\}, I \setminus \{c\} \rangle}$ 
6      $\mathcal{A}_{\langle D, I \rangle} \leftarrow$  the reduction core of  $\mathcal{A}_{\langle D \setminus \{c\}, I \setminus \{c\} \rangle}$  w.r.t.  $\rho$ 
7   else
8      $X \leftarrow X \cup \{c\}$ 
9      $\mathcal{I} \leftarrow \mathcal{I} \cup \{N \in \mathcal{R}_{\mathcal{A}} \mid N \sqsubseteq \mathcal{A}_{\langle D \setminus \{c\}, I \setminus \{c\} \rangle}\}$ 
10 return  $\mathcal{A}_{\langle D, I \rangle}, \mathcal{I}$ 

```

3.1. Base Scheme For Computing a Minimum MSR. Algorithm 1 shows a high-level scheme of our approach for computing a minimum MSR. The algorithm iteratively identifies an ordered set of MSRs, $|M_1| > |M_2| > \dots > |M_k|$, such that the last MSR M_k is a minimum MSR. Each of the MSRs, say M_i , is identified in two steps. First, the algorithm finds an *s-seed*³, i.e., a reduction N_i such that N_i is sufficient and $|N_i| < |M_{i-1}|$. Second, the algorithm *shrinks* N_i into an MSR M_i such that $M_i \sqsubseteq N_i$ (and thus $|M_i| \leq |N_i|$). The initial s-seed N_1 is $\mathcal{A}_{\langle \Psi(\Delta), \Psi(\text{Inv}) \rangle}$, i.e., the reduction that removes all simple clock constraints (which makes all locations of \mathcal{A} trivially reachable). Once there is no sufficient reduction N_i with $|N_i| < |M_{i-1}|$, we know that $M_{i-1} = M_k$ is a minimum MSR.

Note that the algorithm also maintains two auxiliary sets, \mathcal{I} and \mathcal{S} , to store all identified insufficient and sufficient reductions, respectively. In particular, whenever we identify a new MSR M_i , we add every reduction X such that $M_i \sqsubseteq X$ to \mathcal{S} since, by Proposition 2.8, every such X is sufficient. Dually, since M_i is an MSR, then every reduction Y such that $Y \sqsubset M_i$ is necessarily insufficient, and hence we add it to \mathcal{I} . The sets \mathcal{I} and \mathcal{S} are used during the process of finding and shrinking an s-seed which we describe below.

3.2. Shrinking an S-Seed. Our approach for shrinking an s-seed N into an MSR M is based on two concepts: a *critical simple clock constraint* and a *reduction core*.

Definition 3.1 (critical constraint). Given a sufficient reduction $\mathcal{A}_{\langle D, I \rangle}$, a simple clock constraint c is *critical* for $\mathcal{A}_{\langle D, I \rangle}$ iff $\mathcal{A}_{\langle D \setminus \{c\}, I \setminus \{c\} \rangle}$ is insufficient.

Proposition 3.2. *If $c \in D \cup I$ is critical for a sufficient reduction $\mathcal{A}_{\langle D, I \rangle}$ then c is critical for every sufficient reduction $\mathcal{A}_{\langle D', I' \rangle}$ such that $\mathcal{A}_{\langle D', I' \rangle} \sqsubseteq \mathcal{A}_{\langle D, I \rangle}$. Moreover, by Definitions 2.9 and 3.1, $\mathcal{A}_{\langle D, I \rangle}$ is an MSR iff every $c \in D \cup I$ is critical for $\mathcal{A}_{\langle D, I \rangle}$.*

Proof. By contradiction, assume that c is critical for $\mathcal{A}_{\langle D, I \rangle}$ but not for $\mathcal{A}_{\langle D', I' \rangle}$, i.e., $\mathcal{A}_{\langle D \setminus \{c\}, I \setminus \{c\} \rangle}$ is insufficient and $\mathcal{A}_{\langle D' \setminus \{c\}, I' \setminus \{c\} \rangle}$ is sufficient. As $\mathcal{A}_{\langle D', I' \rangle} \sqsubseteq \mathcal{A}_{\langle D, I \rangle}$, we have $\mathcal{A}_{\langle D' \setminus \{c\}, I' \setminus \{c\} \rangle} \sqsubseteq \mathcal{A}_{\langle D \setminus \{c\}, I \setminus \{c\} \rangle}$. By Proposition 2.8, if the reduction $\mathcal{A}_{\langle D' \setminus \{c\}, I' \setminus \{c\} \rangle}$ is sufficient then $\mathcal{A}_{\langle D \setminus \{c\}, I \setminus \{c\} \rangle}$ is also sufficient. \square

³Note that the initial “s” in “s-seed” stands for “sufficient”. Later, in Section 4, we dually introduce “i-seeds” as “insufficient” reductions.

Definition 3.3 (reduction core). Let $\mathcal{A}_{\langle D, I \rangle}$ be a sufficient reduction, ρ a witness run of the sufficiency (i.e., reachability of L_T on $\mathcal{A}_{\langle D, I \rangle}$), and π the path corresponding to ρ . Furthermore, let $\pi' = l_0, e_1, \dots, e_n, l_n$ be the path corresponding to π on the original TA \mathcal{A} (i.e., $\pi = M(\pi')$ (2.4)). The *reduction core* of $\mathcal{A}_{\langle D, I \rangle}$ w.r.t. ρ is the reduction $\mathcal{A}_{\langle D', I' \rangle}$ where $D' = \{(e, \varphi) \mid (e, \varphi) \in D \wedge e = e_i \text{ for some } 1 \leq i \leq n\}$ and $I' = \{(l, \varphi) \mid (l, \varphi) \in I \wedge l = l_i \text{ for some } 0 \leq l \leq n\}$.

Intuitively, the reduction core of $\mathcal{A}_{\langle D, I \rangle}$ w.r.t. ρ removes from \mathcal{A} only the simple clock constraints that appear on the witness path.

Proposition 3.4. *Let $\mathcal{A}_{\langle D, I \rangle}$ be a sufficient reduction, ρ the witness of reachability of L_T on $\mathcal{A}_{\langle D, I \rangle}$, and $\mathcal{A}_{\langle D', I' \rangle}$ the reduction core of $\mathcal{A}_{\langle D, I \rangle}$ w.r.t. ρ . Then $\mathcal{A}_{\langle D', I' \rangle}$ is a sufficient reduction and $\mathcal{A}_{\langle D', I' \rangle} \sqsubseteq \mathcal{A}_{\langle D, I \rangle}$.*

Proof. By Definition 3.3, $D' \subseteq D$ and $I' \subseteq I$, thus $\mathcal{A}_{\langle D', I' \rangle} \sqsubseteq \mathcal{A}_{\langle D, I \rangle}$. As for the sufficiency of $\mathcal{A}_{\langle D', I' \rangle}$, we only sketch the proof. Intuitively, both $\mathcal{A}_{\langle D, I \rangle}$ and $\mathcal{A}_{\langle D', I' \rangle}$ originate from \mathcal{A} by only removing some simple clock constraints ($D \cup I$, and $D' \cup I'$, respectively), i.e., the graph structure of $\mathcal{A}_{\langle D, I \rangle}$ and $\mathcal{A}_{\langle D', I' \rangle}$ is the same, however, some corresponding paths of $\mathcal{A}_{\langle D, I \rangle}$ and $\mathcal{A}_{\langle D', I' \rangle}$ differ in the constraints that appear on the paths. By Definition 3.3, the path π that corresponds to the witness run ρ of $\mathcal{A}_{\langle D, I \rangle}$ is also a path of $\mathcal{A}_{\langle D', I' \rangle}$. Since realizability of a path depends only on the constraints along the path, if π is realizable on $\mathcal{A}_{\langle D, I \rangle}$ then π is also realizable on $\mathcal{A}_{\langle D', I' \rangle}$. \square

Our approach for shrinking a sufficient reduction N is shown in Algorithm 2. The algorithm iteratively maintains a sufficient reduction $\mathcal{A}_{\langle D, I \rangle}$ and a set X of known critical constraints for $\mathcal{A}_{\langle D, I \rangle}$. Initially, $\mathcal{A}_{\langle D, I \rangle} = N$ and $X = \emptyset$. In each iteration, the algorithm picks a simple clock constraint $c \in (D \cup I) \setminus X$ and checks the reduction $\mathcal{A}_{\langle D \setminus \{c\}, I \setminus \{c\} \rangle}$ for sufficiency. If $\mathcal{A}_{\langle D \setminus \{c\}, I \setminus \{c\} \rangle}$ is insufficient, the algorithm adds c to X . Otherwise, if $\mathcal{A}_{\langle D \setminus \{c\}, I \setminus \{c\} \rangle}$ is sufficient, the algorithm obtains a witness run ρ of the sufficiency from the verifier and reduces $\mathcal{A}_{\langle D, I \rangle}$ to the corresponding reduction core. The algorithm terminates when $(D \cup I) = X$. An invariant of the algorithm is that every $c \in X$ is critical for $\mathcal{A}_{\langle D, I \rangle}$. Thus, when $(D \cup I) = X$, $\mathcal{A}_{\langle D, I \rangle}$ is an MSR (Proposition 3.2).

Note that the algorithm also uses the set \mathcal{I} of known insufficient reductions. In particular, before calling a verifier to check a reduction for sufficiency (line 4), the algorithm first checks (in a lazy manner) whether the reduction is already known to be insufficient. Also, whenever the algorithm determines a reduction $\mathcal{A}_{\langle D \setminus \{c\}, I \setminus \{c\} \rangle}$ to be insufficient, it adds $\mathcal{A}_{\langle D \setminus \{c\}, I \setminus \{c\} \rangle}$ and every N , $N \sqsubseteq \mathcal{A}_{\langle D \setminus \{c\}, I \setminus \{c\} \rangle}$, to \mathcal{I} (by Proposition 2.8, every such N is also insufficient).

Finally, note that the algorithm does not add any reduction to the set \mathcal{S} even though it can identify some sufficient reductions during its computation. The reason is that every such identified reduction is larger (w.r.t. \sqsubseteq) than the resultant MSR, and hence all these sufficient reductions are added to \mathcal{S} in the main procedure (Algorithm 1) after the shrinking.

3.3. Finding an S-Seed. We now describe the procedure `findSSeed` that, given the latest identified MSR M , identifies an s-seed, i.e., a sufficient reduction N such that $|N| < |M|$, or returns `null` if there is no s-seed. Let us denote by `CAND` the set of all *candidates* on an s-seed, i.e., $\text{CAND} = \{N \in \mathcal{R}_{\mathcal{A}} \mid |N| < |M|\}$. A brute-force approach would be to check individual reductions in `CAND` for sufficiency until a sufficient one is found, however, this can be practically intractable since $|\text{CAND}| = \sum_{i=1}^{|M|} \binom{|\Psi(\Delta) \cup \Psi(Iw)|}{i-1}$.

Algorithm 3: findSSeed($M, \mathcal{I}, \mathcal{S}$)

```

1 while  $\{N \in \mathcal{R}_{\mathcal{A}} \mid N \notin \mathcal{I} \wedge |N| = |M| - 1\} \neq \emptyset$  do
2    $N \leftarrow$  pick from  $\{N \in \mathcal{R}_{\mathcal{A}} \mid N \notin \mathcal{I} \wedge |N| = |M| - 1\}$ 
3   if  $N$  is sufficient then return  $N, \mathcal{I}, \mathcal{S}$ 
4   else
5      $E, \mathcal{S} \leftarrow$  enlarge( $N, \mathcal{S}$ ) // Algorithm 5
6      $\mathcal{I} \leftarrow \mathcal{I} \cup \{N' \in \mathcal{R}_{\mathcal{A}} \mid N' \sqsubseteq E\}$ 
7 return null,  $\mathcal{I}, \mathcal{S}$ 

```

We provide two observations to prune the set **CAND** of candidates that need to be tested for being an s-seed. The first observation exploits the set \mathcal{I} of already known insufficient reductions: no $N \in \mathcal{I}$ can be an s-seed. The second observation is stated below:

Observation 3.5. For every sufficient reduction $N \in \mathbf{CAND}$ there exists a sufficient reduction $N' \in \mathbf{CAND}$ such that $N \sqsubseteq N'$ and $|N'| = |M| - 1$.

Proof. If $|N| = |M| - 1$, then $N = N'$. For the other case, when $|N| < |M| - 1$, let $N = \mathcal{A}_{\langle D^N, I^N \rangle}$ and $M = \mathcal{A}_{\langle D^M, I^M \rangle}$. We construct $N' = \mathcal{A}_{\langle D^{N'}, I^{N'} \rangle}$ by adding arbitrary $(|M| - |N|) - 1$ simple clock constraint from $(D^M \cup I^M) \setminus (D^N \cup I^N)$ to $(D^N \cup I^N)$, i.e., $D^N \cup I^N \subseteq D^{N'} \cup I^{N'} \subseteq (D^M \cup I^M \cup D^N \cup I^N)$ and $|D^{N'} \cup I^{N'}| = |M| - 1$. By definition of **CAND**, $N' \in \mathbf{CAND}$. Moreover, since $N \not\sqsubseteq N'$ and N is sufficient, then N' is also sufficient (Proposition 2.8). \square

Based on the above observations, we build a set \mathcal{C}_s of indispensable candidates on s-seeds that need to be tested for sufficiency:

$$\mathcal{C}_s = \{N \in \mathcal{R}_{\mathcal{A}} \mid N \notin \mathcal{I} \wedge |N| = |M| - 1\} \quad (3.1)$$

The procedure **findSSeed**, shown in Algorithm 3, in each iteration picks a reduction $N \in \mathcal{C}_s$ and checks it for sufficiency (via the verifier). If N is sufficient, **findSSeed** returns N as the s-seed. Otherwise, when N is insufficient, the algorithm first *enlarges* N into a maximal insufficient reduction (MIR) E such that $N \sqsubseteq E$. By Proposition 2.8, every reduction N' such that $N' \sqsubseteq E$ is also insufficient, thus all these reductions are subsequently added to \mathcal{I} and hence removed from \mathcal{C}_s (note that this includes also N). If \mathcal{C}_s becomes empty, then there is no s-seed.

The purpose of *enlarging* N into E is to quickly prune the candidate set \mathcal{C}_s . We could just add all the insufficient reductions $\{N' \mid N' \sqsubseteq N\}$ to \mathcal{I} , but note that $|\{N' \mid N' \sqsubseteq E\}|$ is exponentially larger than $|\{N' \mid N' \sqsubseteq N\}|$ w.r.t. $|E| - |N|$. The enlargement of N into an MIR E is carried out via Algorithm 5 and it is described later on in Section 4. Note that Algorithm 5 exploits and updates the set \mathcal{S} of already known sufficient reductions.

Finally, let us note that we need to somehow efficiently represent and maintain the sets \mathcal{I} , \mathcal{S} and \mathcal{C}_s . In particular, we need to be able to add elements to these sets and obtain elements from these sets. The problem is that there can be up to exponentially many reductions w.r.t. $|\Psi(\Delta) \cup \Psi(\text{Inv})|$, and hence these sets can be also exponentially large and cannot be stored explicitly. In Section 5, we describe how we efficiently maintain these sets.




FIGURE 3. The situation before the first call of findSSeed.




FIGURE 4. The situation after the first call of findSSeed.

3.4. Example Execution. We illustrate an execution of Algorithm 1 on the TA \mathcal{A} defined in Example 2.3 (Fig. 1) with an initial location l_0 and a target unreachable set of locations $L_T = \{l_4\}$. For the sake of a graphical illustration, we restrict our analysis to possible removal of only 4 simple clock constraints: $c_1 = x \geq 9$, $c_2 = z \geq 9$, $c_3 = x \leq 14$ and $c_4 = u \leq 26$ that appear on edge e_1 , edge e_5 , location l_1 , and location l_3 , respectively (same as in Example 2.14). We will use a bitvector notation to denote the individual reductions, e.g., \mathcal{A}_{1011} represents the reduction $\mathcal{A}_{\langle D, I \rangle}$ where $D \cup I = \{c_1, c_3, c_4\}$.

The computation starts by setting N to \mathcal{A}_{1111} , $\mathcal{I} = \emptyset$ and $\mathcal{S} = \emptyset$. Subsequently, in the first iteration of Algorithm 1, N is shrunk into an MSR M . Assume that $M = \mathcal{A}_{0011}$, and that \mathcal{I} was enlarged to $\mathcal{I} = \{\mathcal{A}_{0001}, \mathcal{A}_{0010}, \mathcal{A}_{0000}\}$. After the shrinking, Algorithm 1 also enlarges the sets \mathcal{I} and \mathcal{S} by adding to them reductions that are smaller and larger than M w.r.t. \sqsubseteq and \sqsupseteq , respectively. We depict the situation at this moment in Figure 3. The power-set in the figure represents all possible reductions of \mathcal{A} (in the picture, we denote a reduction \mathcal{A}_B by the bitvector B). The reductions with dashed border are insufficient, and the reductions with solid border are sufficient. We use green and red background color to highlight the reductions in sets \mathcal{S} and \mathcal{I} , respectively. Moreover, we highlight in blue two reductions, \mathcal{A}_{0100} and \mathcal{A}_{1000} , that will form the set \mathcal{C}_s in the subsequent call of $\text{findSSeed}(M, \mathcal{I}, \mathcal{S})$.

During the execution of $\text{findSSeed}(M, \mathcal{I}, \mathcal{S})$, assume we first pick the candidate reduction $N = \mathcal{A}_{0100} \in \mathcal{C}_s$ and check it for sufficiency. It is insufficient, hence we enlarge it (via $\text{enlarge}(N, \mathcal{S})$) to an insufficient reduction E ; assume $E = \mathcal{A}_{1101}$. Subsequently, we add to \mathcal{I} every reduction N' such that $N' \sqsubseteq E$. The situation at this moment is depicted in Figure 4. At this point, \mathcal{C}_s is empty, i.e., we have the guarantee that there is no s-seed that would be smaller than M w.r.t. \sqsubseteq . Hence, findSSeed terminates, and Algorithm 1 then also terminates determining that the M from the first (and only) iteration is a minimum MSR.

Finally, let us note that there are different possible executions of our algorithm on the given example. In particular, in Algorithm 3, we choose a reduction N from the candidate set \mathcal{C}_s and the choice determines which sufficient reduction will be produced (if any). Similarly, in Algorithm 2, we pick constraints c in some order and this order determines which MSR will be produced. We observed that different reduction and constraint choices affect the performance of the overall algorithm, both in the runtime and the number of performed verifier calls. However, we postpone a development of a suitable heuristic for making good choices here for a future work.

Algorithm 4: Maximum MIR Extraction Scheme

```

1  $N \leftarrow \mathcal{A}_{\langle \emptyset, \emptyset \rangle}; \mathcal{I} \leftarrow \emptyset; \mathcal{S} \leftarrow \emptyset$ 
2 while  $N \neq \text{null}$  do
3    $M, \mathcal{S} \leftarrow \text{enlarge}(N, \mathcal{S})$  // Algorithm 5
4    $\mathcal{I} \leftarrow \mathcal{I} \cup \{M' \mid M' \sqsubseteq M\}$ 
5    $N, \mathcal{S} \leftarrow \text{findSeed}(M, \mathcal{I}, \mathcal{S})$  // Algorithm 6
6 return  $M$ 

```

Algorithm 5: $\text{enlarge}(\mathcal{A}_{\langle D, I \rangle}, \mathcal{S})$

```

1  $X \leftarrow \emptyset$ 
2 while  $(\Psi(\Delta) \cup \Psi(\text{Inv})) \setminus (D \cup I) \neq X$  do
3    $c \leftarrow \text{pick a constraint from } (\Psi(\Delta) \cup \Psi(\text{Inv})) \setminus (D \cup I \cup X)$ 
4   let  $R = \mathcal{A}_{\langle D \cup \{c\}, I \rangle}$  if  $c \in \Psi(\Delta)$  and  $R = \mathcal{A}_{\langle D, I \cup \{c\} \rangle}$  otherwise
5   if  $R \notin \mathcal{S}$  and  $R$  is not sufficient then
6      $\mathcal{A}_{\langle D, I \rangle} \leftarrow R$ 
7   else
8      $X \leftarrow X \cup \{c\}$ 
9      $\mathcal{S} \leftarrow \mathcal{S} \cup \{N \in \mathcal{R}_{\mathcal{A}} \mid N \sqsupseteq R\}$ 
10 return  $\mathcal{A}_{\langle D, I \rangle}, \mathcal{S}$ 

```

4. FINDING MAXIMAL INSUFFICIENT REDUCTIONS

In this section, we describe our approach for finding maximum maximal insufficient reductions (MIRs), and consequently also their complementary minimum minimal guarantees (MGs).

4.1. Base scheme for Computing a Maximum MIR. Our scheme for computing a maximum MIR is shown in Algorithm 4 and it works in a dual way to the scheme for computing a minimum MSR (Algorithm 1). We iteratively identify a sequence M_1, M_2, \dots, M_k of MIRs such that $|M_1| < |M_2| < \dots < |M_k|$ and the last MIR, M_k , is a maximum MIR. To find each MIR M_i in the sequence, we proceed in two steps. First, we identify an *i-seed*, i.e., an insufficient reduction N_i such that $|N_i| > |M_{i-1}|$. Second, we *enlarge* N_i into the MIR M_i (i.e., $N_i \sqsubseteq M_i$ and hence $|N_i| \leq |M_i|$). Once there is no more i-seed, it is guaranteed that the last identified MIR $M_{i-1} = M_k$ is a maximum MIR. The initial i-seed N_1 is the reduction $\mathcal{A}_{\langle \emptyset, \emptyset \rangle} = \mathcal{A}$ (we assume that L_T is indeed unreachable on the input TA \mathcal{A}).

Same as in case of Algorithm 1, this scheme also maintains the auxiliary sets \mathcal{I} and \mathcal{S} to store all identified insufficient and sufficient reductions, respectively.

4.2. Enlarging an I-Seed. The procedure *enlarge* is based on a concept of *conflicting simple clock constraints*.

Definition 4.1 (conflicting constraint). Given an insufficient reduction $\mathcal{A}_{\langle D, I \rangle}$, a simple clock constraint $c \in (\Psi(\Delta) \cup \Psi(\text{Inv})) \setminus (D \cup I)$ is *conflicting* for $\mathcal{A}_{\langle D, I \rangle}$ if the reduction $\mathcal{A}_{\langle D', I' \rangle}$ with $D' \cup I' = D \cup I \cup \{c\}$ is sufficient.

Algorithm 6: findSeed($M, \mathcal{I}, \mathcal{S}$)

```

1 while  $\{N \in \mathcal{R}_{\mathcal{A}} \mid N \notin \mathcal{S} \wedge |N| = |M| + 1\} \neq \emptyset$  do
2    $N \leftarrow$  pick from  $\{N \in \mathcal{R}_{\mathcal{A}} \mid N \notin \mathcal{S} \wedge |N| = |M| + 1\}$ 
3   if  $N$  is insufficient then return  $N, \mathcal{I}, \mathcal{S}$ 
4   else
5      $E, \mathcal{I} \leftarrow$  shrink( $N, \mathcal{I}$ ) // Algorithm 2
6      $\mathcal{S} \leftarrow \mathcal{S} \cup \{N' \in \mathcal{R}_{\mathcal{A}} \mid N' \sqsupseteq E\}$ 
7 return null,  $\mathcal{I}, \mathcal{S}$ 

```

Note that if a constraint c is conflicting for an insufficient reduction N then c is also conflicting for every insufficient reduction N' with $N \sqsubseteq N'$. Moreover, note that a reduction $\mathcal{A}_{\langle D, I \rangle}$ is an MIR iff every $c \in (\Psi(\Delta) \cup \Psi(Inv)) \setminus (D \cup I)$ is conflicting for $\mathcal{A}_{\langle D, I \rangle}$.

The procedure `enlarge(N, \mathcal{S})` is shown in Algorithm 5. The algorithm iteratively maintains an insufficient reduction $\mathcal{A}_{\langle D, I \rangle}$ and a set X of constraints that are known to be conflicting for $\mathcal{A}_{\langle D, I \rangle}$. Initially, $\mathcal{A}_{\langle D, I \rangle} = N$ and $X = \emptyset$. In each iteration, the algorithm picks a simple clock constraint $c \in (\Psi(\Delta) \cup \Psi(Inv)) \setminus (D \cup I \cup X)$ and checks whether c is conflicting for $\mathcal{A}_{\langle D, I \rangle}$. If c is conflicting, then it is added to X . Otherwise, if c is not conflicting, then $\mathcal{A}_{\langle D, I \rangle}$ is extended either to $\mathcal{A}_{\langle D \cup \{c\}, I \rangle}$ or to $\mathcal{A}_{\langle D, I \cup \{c\} \rangle}$ (depending on if $c \in \Psi(\Delta)$ or $c \in \Psi(Inv)$). The algorithm maintains the invariant that every $c \in X$ is conflicting for $\mathcal{A}_{\langle D, I \rangle}$, and hence when $X = (\Psi(\Delta) \cup \Psi(Inv)) \setminus (D \cup I)$, it is guaranteed that $\mathcal{A}_{\langle D, I \rangle}$ is an MIR.

The check whether a constraint $c \in (\Psi(\Delta) \cup \Psi(Inv)) \setminus (D \cup I \cup X)$ is conflicting for $\mathcal{A}_{\langle D, I \rangle}$ is carried out by testing whether the reduction $\mathcal{A}_{\langle D', I' \rangle}$ with $D' \cup I' = D \cup I \cup \{c\}$ is sufficient. In particular, to save some invocations of the verifier, we first, in a lazy manner, check whether $\mathcal{A}_{\langle D', I' \rangle} \in \mathcal{S}$ (i.e., $\mathcal{A}_{\langle D', I' \rangle}$ is already known to be sufficient). If $\mathcal{A}_{\langle D', I' \rangle} \notin \mathcal{S}$, we check $\mathcal{A}_{\langle D', I' \rangle}$ for sufficiency via the verifier. Also, note that whenever we identify a sufficient reduction $\mathcal{A}_{\langle D', I' \rangle}$, we add every reduction X such that $X \sqsupseteq \mathcal{A}_{\langle D', I' \rangle}$ to \mathcal{S} (by Proposition 2.8, every such X is also sufficient). Finally, note that we do not add any insufficient reduction that is identified during the enlargement to the set \mathcal{I} . The reason is that all insufficient reductions that are identified during the enlargement are smaller (w.r.t. \sqsubseteq) than the resultant MIR, and we update \mathcal{I} based on the MIR after the enlargement (Algorithm 4, line 4).

4.3. Finding an I-Seed. The procedure `findSeed` works dually to the procedure `findSeed`. The input is the latest identified MIR M and the sets \mathcal{I} and \mathcal{S} of known insufficient and sufficient reductions. The output is an i-seed, i.e., an insufficient reduction N such that $|N| > |M|$, or `null` if there is no i-seed.

We exploit two basic observations while searching for N . First, observe that no reduction that is already known to be sufficient, i.e., belongs to \mathcal{S} , can be an i-seed. Second, observe that:

Observation 4.2. For every insufficient reduction N with $|N| > |M|$, there exists an insufficient N' such that $N' \sqsubseteq N$ and $|N'| = |M| + 1$.

Proof. Dually to the proof of Observation 3.5. □

Exploiting the above two observations, we build a set \mathcal{C}_i of indispensable candidates on i-seeds that need to be tested for sufficiency to either find an i-seed or to prove that there are no more i-seeds:

$$\mathcal{C}_i = \{N \in \mathcal{R}_{\mathcal{A}} \mid N \notin \mathcal{S} \wedge |N| = |M| + 1\} \quad (4.1)$$

The procedure `findSeed` (Algorithm 6) iteratively picks a reduction $N \in \mathcal{C}_i$ and checks it for sufficiency via the verifier. If N is found to be insufficient, it is returned as the i-seed. Otherwise, when N is sufficient, the algorithm *shrinks* N to an MSR E via Algorithm 2. By Proposition 2.8, every reduction N' such that $N' \sqsupseteq E$ is also sufficient; hence, we add all these reductions to \mathcal{S} (and thus implicitly remove them from \mathcal{C}_i). If \mathcal{C}_i becomes empty, then there is no i-seed.

5. REPRESENTATION OF \mathcal{I} , \mathcal{S} , \mathcal{C}_s , AND \mathcal{C}_i

Let us now describe how to efficiently represent and maintain the sets \mathcal{I} , \mathcal{S} , \mathcal{C}_s and \mathcal{C}_i that are used in our algorithms. Recall that we need to be able to add elements to these sets, obtain elements from these sets, and in case of \mathcal{C}_s and \mathcal{C}_i also perform emptiness checks. The problem is that the size of these sets can be exponential w.r.t. $|\Psi(\Delta) \cup \Psi(Inv)|$ (there are exponentially many reductions), and thus, it is practically intractable to maintain the sets explicitly. Instead, we use a symbolic representation.

Given a timed automaton \mathcal{A} with simple clock constraints $\Psi(\Delta) = \{(e_1, \varphi_1), \dots, (e_p, \varphi_p)\}$ and $\Psi(Inv) = \{(l_1, \varphi_1), \dots, (l_q, \varphi_q)\}$, we introduce two sets of Boolean variables $X = \{x_1, \dots, x_p\}$ and $Y = \{y_1, \dots, y_q\}$. Note that every valuation of the variables $X \cup Y$ one-to-one maps to the reduction $\mathcal{A}_{\langle D, I \rangle}$ such that $(e_i, \varphi_i) \in D$ iff x_i is assigned *True* and $(l_j, \varphi_j) \in I$ iff y_j is assigned *True*.

The sets \mathcal{I} and \mathcal{S} are used both in Algorithm 1 and Algorithm 4, and in both cases, they are gradually maintained during the whole computation of the algorithms. To represent \mathcal{I} , we build a Boolean formula \mathbb{I} such that a reduction N **does not** belong to \mathcal{I} iff N **does** correspond to a model of \mathbb{I} . Initially, $\mathcal{I} = \emptyset$, thus $\mathbb{I} = \text{True}$. To add an insufficient reduction $\mathcal{A}_{\langle D, I \rangle}$ and all reductions N , $N \sqsubseteq \mathcal{A}_{\langle D, I \rangle}$, to \mathcal{I} , we add to \mathbb{I} the clause $(\bigvee_{(e_i, \varphi_i) \in \Psi(\Delta) \setminus D} x_i) \vee (\bigvee_{(l_j, \varphi_j) \in \Psi(Inv) \setminus I} y_j)$. To test if a reduction N is in the set \mathcal{I} , we check if the valuation of $X \cup Y$ that corresponds to N is not a model of \mathbb{I} .

Similarly, to represent \mathcal{S} , we build a Boolean formula \mathbb{S} such that a reduction N **does not** belong to \mathcal{S} iff N **does** correspond to a model of \mathbb{S} . Initially, $\mathcal{S} = \emptyset$, thus $\mathbb{S} = \text{True}$. To add a sufficient reduction $\mathcal{A}_{\langle D, I \rangle}$ and all reductions N , $N \sqsupseteq \mathcal{A}_{\langle D, I \rangle}$, to \mathcal{S} , we add to \mathbb{S} the clause $(\bigvee_{(e_i, \varphi_i) \in D} \neg x_i) \vee (\bigvee_{(l_j, \varphi_j) \in I} \neg y_j)$.

The set \mathcal{C}_s is used only in Algorithm 1; namely in its subroutine `findSSeed`. We build the set \mathcal{C}_s repeatedly during each call of `findSSeed`($M, \mathcal{M}, \mathcal{I}, \mathcal{S}$) based on Equation (3.1) and we encode it via a Boolean formula \mathbb{C}_s such that every model of \mathbb{C}_s **does** correspond to a reduction $N \in \mathcal{C}_s$:

$$\mathbb{C}_s = \mathbb{I} \wedge \text{trues}(|M| - 1) \quad (5.1)$$

where $\text{trues}(|M| - 1)$ is a cardinality encoding forcing that exactly $|M| - 1$ variables from $X \cup Y$ are set to *True*. To check if $\mathcal{C}_s = \emptyset$ or to pick a reduction $N \in \mathcal{C}_s$, we ask a SAT solver for a model of \mathbb{C}_s . To remove an insufficient reduction from \mathcal{C}_s , we update the formula \mathbb{I} (and thus also \mathbb{C}_s) as described above.

Finally, the set \mathcal{C}_i is used in the subroutine `findSeed` of Algorithm 4. We build the set repeatedly during each call of `findSeed`($M, \mathcal{M}, \mathcal{I}, \mathcal{S}$) and to represent it, we maintain a

Boolean formula \mathbb{C}_i such that every model of \mathbb{C}_i **does** correspond to a reduction $N \in \mathcal{C}_i$:

$$\mathbb{C}_i = \mathbb{S} \wedge \mathbf{trues}(|M| + 1) \quad (5.2)$$

where $\mathbf{trues}(|M| + 1)$ is a cardinality encoding forcing that exactly $|M| + 1$ variables from $X \cup Y$ are set to *True*. To check if $\mathcal{C}_i = \emptyset$ or to pick a reduction $N \in \mathcal{C}_i$, we ask a SAT solver for a model of \mathbb{C}_i , and to remove a sufficient reduction from \mathcal{C}_i , we update the formula \mathbb{S} .

6. RELAXING MINIMAL SUFFICIENT REDUCTIONS

In Section 3, we considered a timed automaton $\mathcal{A} = (L, l_0, C, \Delta, Inv)$ and a set of its locations $L_T \subseteq L$, and we presented an efficient algorithm to find a sufficient reduction (see Definition 2.9), i.e., a set of simple clock constraints $D \subseteq \Psi(\Delta)$ (2.1) (over transitions) and $I \subseteq \Psi(Inv)$ (2.2) (over locations) such that L_T is reachable when constraints D and I are removed from \mathcal{A} . In other words, L_T is reachable on $\mathcal{A}_{\langle D, I \rangle}$. Here, instead of completely removing $D \cup I$, our goal is to find a relaxation valuation $\mathbf{r} : D \cup I \rightarrow \mathbb{N} \cup \{\infty\}$ such that L_T is reachable on $\mathcal{A}_{\langle D, I, \mathbf{r} \rangle}$. In addition, we intend to minimize the total change in the timing constants, i.e., $\sum_{\phi \in D \cup I} \mathbf{r}(\phi)$. We present two methods to find such a valuation. The first one solves an MILP using a witness path π'_{L_T} of $\mathcal{A}_{\langle D, I \rangle}$ that ends in L_T . The second one parametrizes each constraint from $D \cup I$ and solves a parameter synthesis problem on the resulting parametric timed automata. While the second method assumes all witness paths of $\mathcal{A}_{\langle D, I \rangle}$ and hence it is guaranteed to find the relaxation \mathbf{r} with minimal $\sum_{\phi \in D \cup I} \mathbf{r}(\phi)$ for the considered MSR, the first method is computationally more efficient.

6.1. MILP Based Relaxation. By the definition of a sufficient reduction, the set L_T is reachable on $\mathcal{A}_{\langle D, I \rangle}$. Consequently, when a verifier is used to check the reachability of L_T , it generates a finite witness run $\rho'_{L_T} = (l_0, \mathbf{0}) \rightarrow_{d_0} (l_1, v_1) \rightarrow_{d_1} \dots \rightarrow_{d_{n-1}} (l_n, v_n)$ of $\mathcal{A}_{\langle D, I \rangle}$ such that $l_n \in L_T$. Let $\pi'_{L_T} = l_0, e'_1, l_1, \dots, e'_{n-1}, l_n$ be the corresponding path on $\mathcal{A}_{\langle D, I \rangle}$, i.e., π'_{L_T} is realizable on $\mathcal{A}_{\langle D, I \rangle}$ due to the delay sequence d_0, d_1, \dots, d_{n-1} and the resulting run is ρ'_{L_T} . The corresponding path on the original TA \mathcal{A} is defined in (2.4):

$$\pi'_{L_T} = M(\pi_{L_T}), \text{ and } \pi_{L_T} = l_0, e_1, l_1, \dots, e_{n-1}, l_n, \quad (6.1)$$

While π'_{L_T} is realizable on $\mathcal{A}_{\langle D, I \rangle}$, π_{L_T} is not realizable on \mathcal{A} since L_T is not reachable on \mathcal{A} . We present an MILP based method to find a relaxation valuation $\mathbf{r} : D \cup I \rightarrow \mathbb{N} \cup \{\infty\}$ such that the path induced by π_{L_T} is realizable on $\mathcal{A}_{\langle D, I, \mathbf{r} \rangle}$.

For a given automaton path $\pi = l_0, e_1, l_1, \dots, e_{n-1}, l_n$ with $e_i = (l_{i-1}, \lambda_i, \phi_i, l_i)$ for each $i = 1, \dots, n-1$, we introduce real valued delay variables $\delta_0, \dots, \delta_{n-1}$ that represent the time spent in each location along the path except the last one (l_n). For a particular path, the value of a clock on a given constraint (invariant or guard) can be mapped to a sum of delay variables as each clock measures the time passed since its last reset:

$$\Gamma(x, \pi, i) = \delta_k + \delta_{k+1} + \dots + \delta_{i-1} \text{ where } k = \max(\{m \mid x \in \lambda_m, m < i\} \cup \{0\}) \quad (6.2)$$

The value of clock x equals to $\Gamma(x, \pi, i)$ on the i -th transition e_i along π . In (6.2), k is the index of the transition where x is last reset before e_i along π , and it is 0 if it is not reset. $\Gamma(0, \pi, i)$ is defined as 0 for notational convenience.

Next, we define an MILP (6.3) for the path π . By using the transformation (6.2), we map each clock constraint along the given path π to constraints over the sequence of

delay variables $\delta_0, \dots, \delta_{n-1}$ as shown in (6.4),(6.5),(6.6). In addition, we introduce integer valued constraint relaxation variables $\{p_{l,\varphi} \mid (l, \varphi) \in I\}$ and $\{p_{e,\varphi} \mid (e, \varphi) \in D\}$ for each simple constraint from $D \cup I$. In particular, for each transition e_i , the simple constraints $\varphi = x - y \sim c \in \mathcal{S}(\phi_i)$ of the guard ϕ_i of e_i are mapped to the new delay variables (6.4), where $p_{e_i,\varphi}$ is the integer valued relaxation variable if $(e_i, \varphi) \in D$, otherwise it is set to 0. On the other hand, for each location l_i , the simple clock constraints $\varphi = x - y \sim c \in \mathcal{S}(Inv(l_i))$ of the invariant $Inv(l_i)$ of l_i are mapped to arriving (6.5) and leaving (6.6) constraints over the delay variables. In (6.5) and (6.6), \mathbf{I} is a binary function mapping *true* to 1 and *false* to 0, and p_{l_i,φ_i} is the integer valued variable if $(l_i, \varphi_i) \in I$, otherwise it is set to 0 as in (6.4). Note that if the invariant is satisfied when arriving and leaving, then, due to the convexity of the constraints, it is satisfied at every time when \mathcal{A} is at the corresponding location along π .

$$\text{minimize } \sum_{(l,\varphi) \in I} p_{l,\varphi} + \sum_{(e,\varphi) \in D} p_{e,\varphi} \quad \text{subject to} \quad (6.3)$$

$$\Gamma(x, \pi, i) - \Gamma(y, \pi, i) \sim c + p_{e_i,\varphi} \quad (\text{guard})$$

for each $i = 1, \dots, n-1$, and $\varphi = x - y \sim c \in \mathcal{S}(\phi_i)$ (6.4)

$$\Gamma(x, \pi, i) \cdot \mathbf{I}(x \notin \lambda_i) - \Gamma(y, \pi, i) \cdot \mathbf{I}(y \notin \lambda_i) \sim c + p_{l_i,\varphi} \quad (\text{arriving, invariant})$$

for each $i = 1, \dots, n, \varphi = x - y \sim c \in \mathcal{S}(Inv(l_i))$ (6.5)

$$\Gamma(x, \pi, i+1) - \Gamma(y, \pi, i+1) \sim c + p_{l_i,\varphi} \quad (\text{leaving, invariant})$$

for each $i = 0, \dots, n-1, \varphi = x - y \sim c \in \mathcal{S}(Inv(l_i))$ (6.6)

$$p_{l,\varphi} \in \mathbb{Z}_+ \quad \text{for each } (l, \varphi) \in I \quad (6.7)$$

$$p_{e,\varphi} \in \mathbb{Z}_+ \quad \text{for each } (e, \varphi) \in D \quad (6.8)$$

$$\delta_i \geq 0 \quad \text{for each } i = 0, \dots, n-1 \quad (6.9)$$

Let $\{p_{l,\varphi}^* \mid (l, \varphi) \in I\}$, $\{p_{e,\varphi}^* \mid (e, \varphi) \in D\}$, and $\delta_0^*, \dots, \delta_{n-1}^*$ denote the solution of MILP (6.3). Define a relaxation valuation \mathbf{r} with respect to the solution as

$$\mathbf{r}(l, \varphi) = p_{l,\varphi}^* \text{ for each } (l, \varphi) \in I, \quad \mathbf{r}(e, \varphi) = p_{e,\varphi}^* \text{ for each } (e, \varphi) \in D. \quad (6.10)$$

Theorem 6.1. *Let $\mathcal{A} = (L, l_0, C, \Delta, Inv)$ be a timed automaton, $\pi = l_0, e_1, l_1, \dots, e_n, l_n$ be a finite path of \mathcal{A} , and $D \subseteq \Psi(\Delta)$, $I \subseteq \Psi(I)$ be guard and invariant constraint sets. If the MILP constructed from \mathcal{A} , π , D and I as defined in (6.3) is feasible, then l_n is reachable on $\mathcal{A}_{\langle D, I, \mathbf{r} \rangle}$ with \mathbf{r} as defined in (6.10).*

Proof. Denote the optimal solution of MILP (6.3) by $\{p_{l,\varphi}^* \mid (l, \varphi) \in I\}$, $\{p_{e,\varphi}^* \mid (e, \varphi) \in D\}$, and $\delta_0^*, \dots, \delta_{n-1}^*$. For simplicity of presentation set $p_{l,\varphi}^*$ to 0 for each $(l, \varphi) \in \Psi(Inv) \setminus I$ and set $p_{e,\varphi}^*$ to 0 for each $(e, \varphi) \in \Psi(\Delta) \setminus D$. Let $\mathcal{A}_{\langle D, I, \mathbf{r} \rangle} = (L, l_0, C, \Delta', Inv')$ and $T(\mathcal{A}_{\langle D, I, \mathbf{r} \rangle}) = (S, s_0, \Sigma, \rightarrow)$. Define clock value sequence v_0, v_1, \dots, v_n with respect to the path π with $e_i = (l_{i-1}, \lambda_i, \phi_i, l_i)$ and the delay sequence $\delta_0^*, \dots, \delta_{n-1}^*$ iteratively as $v_i = \mathbf{0}$ and $v_i = (v_{i-1} + \delta_{i-1}^*)[\lambda_i := 0]$ for each $i = 1, \dots, n$. Along the path π , v_i is consistent with $\Gamma(\cdot, \pi, i)$ (6.2) such that

$$(a) \ v_i(x) = \Gamma(x, \pi, i) \cdot \mathbf{I}(x \notin \lambda_i) \quad \text{and} \quad (b) \ v_i(x) + \delta_i^* = \Gamma(x, \pi, i+1) \quad (6.11)$$

For a simple constraint $\varphi = x - y \sim c + p_{l_i,\varphi}^* \in Inv'(l_i)$ (i.e. $x - y \sim c \in Inv(l_i)$ via Definition 2.5 and (6.10)), it holds that $v_i(x) - v_i(y) \sim c + p_{l_i,\varphi}^*$ via (6.5) and (6.11)-a. Then by (6.10) $v_i \models Inv'(l_i)$ and $(l_i, v_i) \in S$. Similarly, $v_i + \delta_i^* \models Inv'(l_i)$ via (6.6) and (6.11)-b. Hence, $(l_i, v_i + \delta_i^*) \in S$ and $(l_i, v_i) \xrightarrow{\delta_i^*} (l_i, v_i + \delta_i^*)$ (delay transition). Furthermore,

by (6.4), (6.10), (6.11)-b and Definition 2.5, we have $v_i + \delta_i^* \models R(\phi_i, D|_{e_i}, \mathbf{r}|_{e_i})$ and $(l_i, v_i + \delta_i^*) \xrightarrow{act} (l_{i+1}, v_{i+1})$ (discrete transition). As $s_0 = (l_0, \mathbf{0}) \in S$, and the derivation applies to each $i = 1, \dots, n$, we reach that $\rho = (l_0, v_0), \dots, (l_n, v_n) \in [[\mathcal{A}_{\langle D, I, \mathbf{r} \rangle}]]$, and l_n is reachable on $\mathcal{A}_{\langle D, I, \mathbf{r} \rangle}$. \square

A linear programming (LP) based approach was used in [BBBR07] to generate the optimal delay sequence for a given path of a weighted timed automata. In our case, the optimization problem is in MILP form since we find an integer valued relaxation valuation (\mathbf{r}) in addition to the delay variables.

Recall that we construct relaxation sets D and I via Algorithm 1, and define π_{L_T} (6.1) that reach L_T such that the corresponding path π'_{L_T} is realizable on $\mathcal{A}_{\langle D, I \rangle}$. Then, we define MILP (6.3) with respect to π_{L_T} , D and I , and define \mathbf{r} (6.10) according to the optimal solution. Note that this MILP is always feasible since π'_{L_T} is realizable on $\mathcal{A}_{\langle D, I \rangle}$. Finally, by Theorem 6.1, we conclude that L_T is reachable on $\mathcal{A}_{\langle D, I, \mathbf{r} \rangle}$.

6.2. Parameter Synthesis Based Relaxation. As our second approach, we parametrize each simple constraint in the considered MSR. In particular, for each $(v, \varphi = x - y \sim c) \in D \cup I$ (v is either a transition e or a location l), we introduce a positive valued parameter $p_{v, \varphi}$ and replace the corresponding constraint with $x - y \sim c + p_{v, \varphi}$. The resulting TA $\mathcal{A}^{D \cup I}$ is parametric with parameter set $P = \{p_{(v, \varphi)} \mid (v, \varphi) \in D \cup I\}$. $\mathcal{A}^{D \cup I}$ has $|D \cup I|$ parametric constraints and each parameter appears in a single constraint. Subsequently, we use a parameter synthesis tool that generates the set of all parameter valuations $\mathbb{P} \subseteq \mathbb{R}_+^{|P|}$ for $\mathcal{A}^{D \cup I}$ such that the target set L_T becomes reachable, i.e., for each $\mathbf{p} \in \mathbb{P}$, L_T is reachable on $\mathcal{A}^{D \cup I}(\mathbf{p})$, where $\mathcal{A}^{D \cup I}(\mathbf{p})$ is a non-parametric TA obtained from $\mathcal{A}^{D \cup I}$ and \mathbf{p} by replacing each parameter $p_{v, \varphi}$ with the corresponding valuation $\mathbf{p}(p_{v, \varphi})$. Then, we choose the integer valued parameter valuation $\mathbf{p}^* : P \rightarrow \mathbb{N}$ that minimizes the total change, i.e., $\mathbf{p}^* = \arg \min_{\mathbf{p} \in \mathbb{P} \cap \mathbb{N}} \sum_{(v, \varphi) \in D \cup I} \mathbf{p}(p_{v, \varphi})$. The parameter synthesis method ensures that L_T is reachable on $\mathcal{A}_{\langle D, I, \mathbf{r} \rangle}$, where \mathbf{r} is defined from \mathbf{p}^* as in (6.10).

6.3. Comparison of the MSR Relaxation Methods. The MILP based relaxation method minimizes the total change in the timing constants ($\sum_{\phi \in D \cup I} \mathbf{r}(\phi)$) for a particular path π_{L_T} . Thus, the resulting relaxation valuation (6.10) is not necessarily minimal for the considered MSR $D \cup I$. Whereas, the parameter synthesis based relaxation method is guaranteed to find the minimal valuation (as it considers all paths of $\mathcal{A}_{\langle D, I \rangle}$). However, it is computationally more expensive compared to the MILP approach due to the complexity of the parameter synthesis for timed automata.

Let us note that both our approaches work with a fixed minimum MSR $\mathcal{A}_{\langle D, I \rangle}$. However, observe that there might exist another minimum MSR $\mathcal{A}_{\langle D', I' \rangle}$ with $|D' \cup I'| = |D \cup I|$ that would lead to a smaller overall change of the constraints (i.e., smaller $\sum_{c \in D' \cup I'} \mathbf{r}(c)$). While our approach can be applied to a number of minimum MSRs, processing all of them can be practically intractable.

7. RELAXING MINIMAL GUARANTEES

In Section 4, we presented a method to find a minimal guarantee $D \cup I$ (MG), i.e., a minimal subset of the constraints that need to be left in the system to ensure that a target (unsafe) location is still not reachable. In particular, L_T is not reachable on $\mathcal{A}' = \mathcal{A}_{\langle \Psi(\Delta) \setminus D, \Psi(Inv) \setminus I \rangle}$ (see Definition 2.11). In this section, we attempt to relax the timing constraints in the resulting TA \mathcal{A}' , i.e., $D \cup I$, as much as possible while ensuring that L_T is still unreachable. Thus, we analyze how *robust* the resulting TA is against constraint perturbations with respect to the safety specification. We consider two settings for relaxing the constraints from the MG. First, as in the MSR case, we find the maximal total relaxation of the remaining clock constraints such that L_T is still unreachable. Second, we find a single relaxation value δ such that L_T is still unreachable when each constraint is relaxed by δ , that is referred as the robustness degree in literature [BMS13].

7.1. Maximizing the Total Change. As described in Section 6.2, we parametrize each simple constraint from the considered constraint set, i.e. in this case, it is the MG $D \cup I$ on $\mathcal{A}' = \mathcal{A}_{\langle \Psi(\Delta) \setminus D, \Psi(Inv) \setminus I \rangle}$. Note that each constraint in the resulting TA that is denoted by $\mathcal{A}^{D \cup I}$ is parametric and the parameter set is $P = \{p_{(v,\varphi)} \mid (v,\varphi) \in D \cup I\}$. Then, we use a parameter synthesis tool that generates the set of all parameter valuations $\mathbb{P} \subseteq \mathbb{R}_+^{|P|}$ for $\mathcal{A}^{D \cup I}$ such that the set L_T is still unreachable. Finally, we chose the integer valued parameter valuation $\mathbf{p}^* : P \rightarrow \mathbb{N}$ that maximize the total change, i.e., $\mathbf{p}^* = \arg \max_{\mathbf{p} \in \mathbb{P} \cap \mathbb{N}} \sum_{(v,\varphi) \in D \cup I} \mathbf{p}(p_{v,\varphi})$. Note that, the maximal total change is finite since $\mathbf{p}(p_{v,\varphi})$ is finite for each valuation $\mathbf{p} \in \mathbb{P}$ and constraint $(v,\varphi) \in D \cup I$ due to the minimality of the MG. The integer valued parameter valuation identifies the maximal total change in the constraint thresholds that can be applied to the TA $\mathcal{A}_{\langle \Psi(\Delta) \setminus D, \Psi(Inv) \setminus I \rangle}$ without violating the safety specification. In particular, for any relaxation valuation \mathbf{r} over $D \cup I$ with $\sum_{(v,\varphi) \in D \cup I} \mathbf{r}(v,\varphi) > \sum_{(v,\varphi) \in D \cup I} \mathbf{p}^*(p_{v,\varphi})$, the automaton $\mathcal{A}'_{\langle D, I, \mathbf{r} \rangle}$ violates the safety specification.

7.2. Finding the Robustness Degree. A timed automaton $\mathcal{A} = (L, l_0, C, \Delta, Inv)$ is said to δ -robustly satisfy a linear-time property, such as a safety property, if the TA $\mathcal{A}_{\langle \Psi(\Delta), \Psi(Inv), \mathbf{r}_\delta \rangle}$ obtained by relaxing each simple constraint of \mathcal{A} by δ satisfies the property [BMS13, WDMR08], where

$$\mathbf{r}_\delta(v, \varphi) = \delta \text{ for each } (v, \varphi) \in \Psi(\Delta) \cup \Psi(Inv).$$

A robustness value δ can be found via parametric analysis [BMS13, AS11]. Here, our goal is to find the maximal robustness value δ^* for the timed automaton \mathcal{A}' such that L_T is not reachable on $\mathcal{A}'_{\langle D, I, \mathbf{r}_{\delta^*} \rangle}$ (recall that $\mathcal{A}' = \mathcal{A}_{\langle \Psi(\Delta) \setminus D, \Psi(Inv) \setminus I \rangle}$). Let \mathbb{P} be the parameter valuation set defined as in Section 7.1. Then, L_T is not reachable on $\mathcal{A}'_{\langle D, I, \mathbf{r}_{\delta^*} \rangle}$ for each $\delta \in \mathbf{D}$ (\mathcal{A}' δ -robustly satisfies the safety specification), where

$$\mathbf{D} = \{\delta \mid \mathbf{p} \in \mathbb{P} \text{ and } \mathbf{p}(v, \varphi) \geq \delta \text{ for each } (v, \varphi) \in D \cup I\}$$

Alternatively, one can use the same parameter p for each simple constraint of \mathcal{A}' to obtain a parametric TA $\mathcal{A}'^{D \cup I}$ from \mathcal{A}' by replacing each simple constraint $(v, x - y \sim c) \in D \cup I$ with $(v, x - y \sim c + p)$. The resulting TA $\mathcal{A}'^{D \cup I}$ has a single parameter p and $|D \cup I|$ parametric constraints. Then, a parameter synthesis tool generates the set of all parameter

valuations for $\mathcal{A}^{D \cup I}$ such that the set L_T is still unreachable. Note that the set obtained in the second case is equal to \mathbf{D} .

In literature, the robustness analysis is studied considering the imperfect implementations of the \mathcal{A} , e.g. timing or measuring errors, thus real valued robustness is used. In this work, we analyze the properties of the timed automata model itself, i.e., constraints and the constraint thresholds. Hence, we focus on integer valued relaxations of the TA. For this reason, we define the optimal relaxation value as $\delta^* = \max \mathbf{D} \cap \mathbb{N}$.

8. RELATED WORK

8.1. Timed Automata. In the literature, the uncertainties about timing constants are handled by representing such constants as parameters in a parametric timed automaton (PTA), i.e., a TA where clock constants can be represented with parameters. Subsequently, a parameter synthesis method, such as [AFKS12, LRST09, BBČB18], is used to find suitable values of the parameters for which the resultant TA satisfies the specification. However, most of the parameter synthesis problems are undecidable [And19b]. While symbolic algorithms without termination guarantees exist for some subclasses [AS11, BBBČ16, JLR15, AKL⁺19], these algorithms are computationally very expensive compared to model checking (see [And19a]). Furthermore, it is not straightforward to integrate the minimization of the number of modified constraints in the parameter synthesis method for the reachability properties in an efficient way. For example, Imitator tool [AFKS12] generates all parameter valuations such that the reachability or the safety property holds when the synthesis algorithm terminates. One approach would be parametrizing each simple constraint of the TA, then finding the valuation minimizing the number of non-zero parameters returned by the tool for the reachability problem. However, due to the dependence of the computation time on the number of parameters, this approach would be impractical. Similarly, for the safety problem, each constraint can be parametrized and further analysis can be performed on the result returned by the synthesis tool in order to find the minimal set of constraints that need to be left in the TA to ensure safety. While assigning 0 to a parameter that bounds a clock from below (i.e. $p \leq x$) or infinity to a parameter that bounds a clock from above (i.e. $x < p$) are equivalent to removing these constraints, it is not straightforward to deduce the constraint removal decision for constraints that involve multiple clocks (i.e. $x - y \leq p$). Moreover, as mentioned for reachability, it would be impractical to solve the parameter synthesis problem when each constraint is parametrized.

Repair of a TA has been studied in recent works [KLW19, EYG21, AAGR19], where, similar to the reachability problem considered in this paper, the goal is to modify a given timed automaton such that the repaired TA satisfies the specification. In [AAGR19], it is assumed that some of the clock constraints are incorrect and the goal is to make the TA compliant with an oracle that decides if a trace of the TA belongs to a system or not. To repair the TA, the authors of [AAGR19] parametrize the initial TA and generate parameters by analyzing traces of the TA. They minimize the total change of the timing constraints, while we primarily minimize the number of changed constraints and then the total change. Furthermore, their approach cannot handle reachability properties. In [KLW19, EYG21], the goal is to repair the TA to avoid undesired behaviors, e.g., traces violating universal properties such as safety. In particular, in [KLW19], a single violating trace is analyzed by running an SMT solver on a linear arithmetic encoding of the trace. The generated

repair suggestions include introducing clock resets and changing the clock constraints (both constraint bounds and constraint operators). As these operations can significantly change the set of traces of the automaton, they check the equivalence of the original and the repaired models after applying the suggested repair. In [EYG21], new clocks and constraints over these new clocks are introduced to restrict the behavior of the automaton to eliminate the violating traces. Neither of these approaches can handle reachability properties. For safety properties, we consider a timed automaton satisfying the property, identify the constraints of the automaton that are effective in the satisfaction of the property and further analyze these constraints. On the other hand, both [KLW19] and [EYG21] aim at repairing a TA that violates the given property.

The robustness of timed automata is studied considering non-ideal implementations of the model, i.e., imprecise clocks, measuring errors, etc. [BMS13, WDMR08]. A timed automaton is said to be robust against clock perturbations and drifts for safety specifications when a TA obtained by allowing the clocks to drift within the given limits and relaxing each constraint by a certain amount satisfies the specification. A complementary approach to robustness analysis is called shrinkability [SBM11, San13]: tighten (shrink) all of the constraints by a positive amount while guaranteeing that the resulting automaton is non-blocking and/or time abstract simulates the original one (thus preserves the safety and reachability properties). Consequently, the shrunk automaton is robust against constraint perturbations. Region automata construction and difference bound matrices are used for the computation of the robustness degree in [BMS13, WDMR08, SBM11, San13]. A parameter synthesis method is also utilized to find the robustness in [AFKS12]. In this work, a similar constraint relaxation approach is used for reachability and safety specifications. To satisfy reachability specifications, we relax the constraints from minimal sufficient reductions. For safety specifications, we first identify a set of constraints that are active in satisfying the safety specification (minimal guarantee, MG), and then perform robustness analysis only over these constraints. In order to relax the identified constraints, we present an MILP based approach and also employ parameter synthesis by parametrizing constraints from the identified sets.

8.2. Minimal Sets over a Monotone Predicate. Although the concepts of minimal sufficient reductions (MSRs) and minimal guarantees (MGs) are novel in the context of timed automata, similar concepts appear in other areas of computer science. For example, see minimal unsatisfiable subsets [dlBSW03], minimal correction subsets [MHJ⁺13], minimal inconsistent subsets [BBB⁺16, Ben17], or minimal inductive validity cores [GWG17]. All these concepts can be generalized as *minimal sets over monotone predicates (MSMPs)* [MJB13, MJM17]. The input is a reference set R and a monotone predicate $\mathbf{P} : \mathcal{P}(R) \rightarrow \{1, 0\}$, and the goal is to find minimal subsets of R that satisfy the predicate. In the case of MSRs, the reference set is the set of all simple constraints $\Psi(\Delta) \cup \Psi(Inv)$ and, for every $D \cup I \subseteq \Psi(\Delta) \cup \Psi(Inv)$, the predicate is defined as $\mathbf{P}(D \cup I) = 1$ iff $\mathcal{A}_{\langle D, I \rangle}$ is sufficient. Similarly, in the case of MGs, the reference set is the set of all simple constraints $\Psi(\Delta) \cup \Psi(Inv)$ and, for every $D \cup I \subseteq \Psi(\Delta) \cup \Psi(Inv)$, the predicate is defined as $\mathbf{P}(D \cup I) = 1$ iff $\mathcal{A}_{\langle \Psi(\Delta) \setminus D, \Psi(Inv) \setminus I \rangle}$ is insufficient.

Many algorithms for finding MSMPes were proposed (e.g., [IPLM15, LML⁺09, LPMM16, BK16, BBČB16, BČB18, BČ20a, MHJ⁺13, BČ20b, IMMV16, GWG17, BGWČ18]), including also several algorithms (e.g. [IPLM15, LML⁺09, IJM16]) for extracting minimum MSMPs. Most of the existing algorithms are *domain-specific*, i.e. tailored to a particular instance of

MSMP and extensively exploiting specific properties of the instances (such as we exploit reduction cores in case of MSRs). Hence, the domain-specific solutions cannot be directly used for finding MSRs and/or MGs. Several *domain-agnostic* MSMP identification algorithms (e.g. [BS05, SKFP12, LPMM16]) were also proposed, i.e., algorithms that can be used for any type of MSMPs. Due to their universality, domain-agnostic approaches are usually not as efficient as the domain-specific solutions. However, it is often the case that a domain-agnostic algorithm serves as a basis while building a domain-specific solution [BČ18, Ben21]. Some techniques we presented in this paper, including mainly the symbolic representation (Section 5) and the shrinking and growing procedures, are specializations of existing domain-agnostic solutions (see [LPMM16, Ben21]).

9. EXPERIMENTAL EVALUATION

We implemented the proposed reduction, guarantee and relaxation methods in a tool called Tamus. We use UPPAAL [BDL⁺06] for sufficiency checks and witness computation, Imitator [AFKS12] for parameter synthesis for PTA and CBC solver from Or-tools library [PF] for the MILP part. All experiments were run on a laptop with Intel i5 quad core processor at 2.5 GHz and 8 GB ram using a time limit of 20 minutes per benchmark. The tool and used benchmarks are available at <https://github.com/jar-ben/tamus>.

As discussed in Section 8, an alternative approach to solve the MSR problem (Problem 2.15) is to parameterize each simple clock constraint of the TA. Then, we can run a parameter synthesis tool on the parameterized TA to identify the set of all possible valuations of the parameters for which the TA satisfies the reachability property. Subsequently, we can choose the valuations that assign non-zero values (i.e., relax) to the minimum number of parameters, and out of these, we can choose the one with a minimum cumulative change of timing constants. In our experimental evaluation, we evaluate the state-of-the-art parameter synthesis tool Imitator [AFKS12] to run such analysis. Although Imitator is not tailored for our problem, it allows us to measure the relative scalability of our approach compared to a well-established synthesis technique. In addition, we employ Imitator to solve the parameter synthesis problems for finding the optimal relaxation for a given MSR (Section 6.2), to find the maximal total change for a given MG (Section 7.1) and to find the robustness degree for the MG (Section 7.2).

We used two collections of benchmarks to evaluate the proposed methods: one is obtained from the literature, and the other are crafted timed automata modeling a machine scheduling problem. In the following, we introduce these benchmarks and present the results of the experiments for reductions and guarantees.

9.1. Experimental Results on Machine Scheduling Automata. A scheduler automaton is composed of a set of paths starting in location l_0 and ending in location l_1 . Each path $\pi = l_0 e_k l_k e_{k+1} \dots l_{k+M-1} e_{k+M} l_1$ represents a particular scheduling scenario where an intermediate location, e.g. l_i for $i = k, \dots, k + M - 1$, belongs to a unique path (only one incoming and one outgoing transition). Thus, a TA that has p paths with M intermediate locations in each path has $M \cdot p + 2$ locations and $(M + 1) \cdot p$ transitions. Each intermediate location represents a machine operation, and periodic simple clock constraints are introduced to mimic the limitations on the corresponding durations. For example, assume that the total time to use machines represented by locations l_{k+i} and l_{k+i+1} is upper (or lower) bounded by c for $i = 0, 2, \dots, M - 2$. To capture such a constraint with a period of $t = 2$, a new clock x

| Model | | MSR Results | | | | MG Results | | | | | | | |
|--------------------------|----------|-------------|-------|-------|-------|------------|-------|-------|-------|------------|-------------------------|------------|-------------|
| Name | $ \Psi $ | d_R | v_R | t_R | c_R | d_G | v_G | t_G | c_G | t_G^{IT} | $\delta_{\mathbb{R}}^*$ | δ^* | t_G^{ITS} |
| $\mathcal{A}_{(3,1,12)}$ | 11 | 2 | 33 | 0.18 | 6 | 5 | 88 | 0.55 | 2 | 0.006 | 0.59 | 0 | 0.003 |
| $\mathcal{A}_{(3,2,12)}$ | 17 | 1 | 13 | 0.13 | 13 | 9 | 175 | 1.47 | 14 | 0.016 | 0.59 | 0 | 0.005 |
| $\mathcal{A}_{(3,1,18)}$ | 16 | 3 | 61 | 0.40 | 9 | 5 | 279 | 1.95 | 2 | 0.006 | 0.59 | 0 | 0.005 |
| $\mathcal{A}_{(3,2,18)}$ | 24 | 1 | 498 | 4.68 | 6 | 10 | 519 | 5.33 | 7 | 0.031 | 0.59 | 0 | 0.007 |
| $\mathcal{A}_{(3,1,24)}$ | 21 | 4 | 96 | 0.73 | 12 | 5 | 985 | 8.15 | 2 | 0.005 | 0.59 | 0 | 0.002 |
| $\mathcal{A}_{(3,2,24)}$ | 32 | 1 | 51 | 0.65 | 16 | 12 | 1291 | 16.76 | 0 | 0.051 | 0.59 | 0 | 0.008 |
| $\mathcal{A}_{(3,1,30)}$ | 26 | 5 | 140 | 1.24 | 15 | 5 | 1829 | 17.49 | 2 | 0.007 | 0.59 | 0 | 0.003 |
| $\mathcal{A}_{(3,2,30)}$ | 40 | 1 | 63 | 0.96 | 9 | 13 | 2901 | 45.54 | 0 | 0.176 | 0.59 | 0 | 0.009 |
| $\mathcal{A}_{(5,1,12)}$ | 16 | 3 | 90 | 0.55 | 10 | 4 | 214 | 1.39 | 1 | 0.004 | 0.49 | 0 | 0.002 |
| $\mathcal{A}_{(5,2,12)}$ | 24 | 1 | 192 | 1.54 | 13 | 6 | 166 | 1.46 | 14 | 0.007 | 1.33 | 1 | 0.003 |
| $\mathcal{A}_{(5,1,18)}$ | 23 | 4 | 149 | 1.04 | 16 | 4 | 399 | 3.02 | 1 | 0.004 | 0.49 | 0 | 0.004 |
| $\mathcal{A}_{(5,2,18)}$ | 35 | 1 | 25 | 0.35 | 6 | 7 | 796 | 9.13 | 3 | 0.010 | 0.24 | 0 | 0.008 |
| $\mathcal{A}_{(5,1,24)}$ | 31 | 6 | 327 | 2.67 | 24 | 4 | 1050 | 9.34 | 1 | 0.004 | 0.49 | 0 | 0.004 |
| $\mathcal{A}_{(5,2,24)}$ | 47 | 2 | 373 | 4.86 | 31 | 8 | 2708 | 40.46 | 13 | 0.021 | 0.49 | 0 | 0.015 |
| $\mathcal{A}_{(5,1,30)}$ | 39 | 7 | 571 | 5.54 | 29 | 4 | 1864 | 19.57 | 1 | 0.007 | 0.49 | 0 | 0.006 |
| $\mathcal{A}_{(5,2,30)}$ | 59 | 2 | 624 | 9.45 | 17 | 9 | 1556 | 26.45 | 6 | 0.028 | 0.49 | 0 | 0.010 |
| $\mathcal{A}_{(7,1,12)}$ | 19 | 3 | 119 | 0.74 | 11 | 3 | 153 | 0.97 | 0 | 0.004 | 0.33 | 0 | 0.003 |
| $\mathcal{A}_{(7,2,12)}$ | 28 | 1 | 70 | 0.62 | 13 | 5 | 247 | 2.16 | 10 | 0.005 | 0.33 | 0 | 0.003 |
| $\mathcal{A}_{(7,1,18)}$ | 28 | 5 | 314 | 2.33 | 25 | 3 | 402 | 3.19 | 0 | 0.002 | 0.33 | 0 | 0.002 |
| $\mathcal{A}_{(7,2,18)}$ | 42 | 1 | 175 | 2.00 | 6 | 6 | 528 | 6.45 | 3 | 0.010 | 0.33 | 0 | 0.009 |
| $\mathcal{A}_{(7,1,24)}$ | 38 | 7 | 615 | 5.29 | 39 | 3 | 1442 | 14.28 | 0 | 0.002 | 0.33 | 0 | 0.002 |
| $\mathcal{A}_{(7,2,24)}$ | 57 | 2 | 944 | 12.67 | 21 | 6 | 863 | 12.86 | 11 | 0.012 | 0.33 | 0 | 0.011 |
| $\mathcal{A}_{(7,1,30)}$ | 48 | 10 | 1559 | 16.75 | 47 | 3 | 2302 | 26.82 | 0 | 0.008 | 0.33 | 0 | 0.003 |
| $\mathcal{A}_{(7,2,30)}$ | 72 | 2 | 675 | 11.25 | 14 | 7 | 1295 | 23.43 | 4 | 0.021 | 0.33 | 0 | 0.015 |

TABLE 1. Results for the scheduler TA, where $|\Psi| = |\Psi(\Delta) \cup \Psi(I)|$ is the total number of constraints, $d_R/d_G = |D \cup I|$ is the minimum MSR/MG size, v_R/v_G is the number of reachability checks during minimum MSR/MG computation, t_R/t_G is the computation time in seconds for minimum MSR/MG computation (including the reachability checks), c_R is the optimal cost of (6.3), c_G is the maximal total change $\sum_{(v,\varphi) \in D \cup I} \mathbf{P}^*(p_{v,\varphi})$ for the MG, $\delta_{\mathbb{R}}^*$ is the real valued maximal robustness value, and δ^* is the integer valued optimal relaxation value. t_G^{IT} and t_G^{ITS} are the Imitator computation times for maximizing the total change and finding the robustness degree, respectively.

is introduced and it is reset and checked on every t^{th} transition along the path, i.e., for every $m \in \{i \cdot t + k \mid i \cdot t \leq M - 1\}$, let $e_m = (l_m, \lambda_m, \phi_m, l_{m+1})$, add x to λ_m , set $\phi_m := \phi_m \wedge x \leq c$ ($x \geq c$ for lower bound). A periodic constraint is denoted by (t, c, \sim) , where t is its period, c is the timing constant, and $\sim \in \{<, \leq, >, \geq\}$. A set of such constraints are defined for each path to capture possible restrictions. In addition, a bound T on the total execution time is captured with the constraint $x \leq T$ on transition e_{k+M} over a clock x that is not reset on any transition. A realizable path to l_1 represents a feasible scheduling scenario. We have

generated 24 test cases. A test case $\mathcal{A}_{(c,p,M)}$ represents a timed automaton with $c \in \{3, 5, 7\}$ clocks, and $p \in \{1, 2\}$ paths with $M \in \{12, 18, 24, 30\}$ intermediate locations in each path. $R_{c,i}$ is the set of periodic restrictions defined for the i^{th} path of an automaton with c clocks:

$$\begin{aligned} R_{3,1} &= \{(2, 11, \geq), (3, 15, \leq)\} & R_{3,2} &= \{(4, 17, \geq), (5, 20, \leq)\} \\ R_{5,1} &= R_{3,1} \cup \{(4, 21, \geq), (5, 25, \leq)\} & R_{5,2} &= R_{3,2} \cup \{(8, 33, \geq), (9, 36, \leq)\} \\ R_{7,1} &= R_{5,1} \cup \{(6, 31, \geq), (7, 35, \leq)\} & R_{7,2} &= R_{5,2} \cup \{(12, 49, \geq), (12, 52, \leq)\} \end{aligned}$$

Note that $\mathcal{A}_{(c,2,M)}$ emerges from $\mathcal{A}_{(c,1,M)}$ by adding a path with restrictions $R_{c,2}$.

MSR analysis. A path to l_1 describes a scheduling scenario for a scheduler automaton ($\mathcal{A}_{(c,p,M)}$). However, location l_1 is unreachable for each of the introduced automata. Thus, our goal is to find a realizable path to l_1 by performing a minimum amount of change. In order to achieve this, we define the target set as $L_T = \{l_1\}$ and run the developed MSR methods. The results obtained on the scheduler automata are summarized in Table 1. Tamus solved all models and the longest computation time was 16.75 seconds. As expected, the computation time t_R is depends on the number $|\Psi|$ of simple clock constraints in the model.

When each simple constraint is parametrized, Imitator solved $\mathcal{A}_{(3,1,12)}$, $\mathcal{A}_{(3,2,12)}$, $\mathcal{A}_{(3,1,18)}$, and $\mathcal{A}_{(5,1,12)}$ within 0.09, 0.5, 62, and 71 seconds, respectively, and timed-out for the other models. In addition, we run Imitator with a flag “witness” that terminates the computation when a satisfying valuation is found. The use of this flag reduced the computation time for the aforementioned cases, and it allowed to solve two more models: $\mathcal{A}_{(3,2,18)}$ and $\mathcal{A}_{(5,2,12)}$. However, using this flag, Imitator often did not provide a solution that minimizes the number of relaxed simple clock constraints.

MG analysis. We also run the developed methods to find MGs, the corresponding maximal total changes and robustness values over the scheduler automata models with $L_T = \{l_1\}$. The results are reported in Table 1. Tamus was also able to generate MG results for all models and $\mathcal{A}_{(5,2,24)}$ took the longest with 40.46 seconds. In the MG case, when any of the $|\Psi| - d_G + 1$ simple constraints are removed from the original scheduler automata $\mathcal{A}_{(c,p,M)}$, L_T becomes reachable. In addition, we run Imitator to find the maximal total change (Section 7.1) and the robustness degree (Section 7.2) for the identified MG $D \cup I$. Specifically we first ran Tamus on TA $\mathcal{A}_{(c,p,M)}$ and then removed every constraint that is not in $D \cup I$ (i.e., obtained $\mathcal{A}_{(c,p,M), <\Psi(\Delta) \setminus D, \Psi(Inv) \setminus I>}$) and parameterized every constraint that is in $D \cup I$. A different parameter is used for every constraint to find the maximal total change and the same parameter is used for every constraint to find the robustness degree. Since d_G is much smaller than $|\Psi|$, Imitator generated the results for every model within 0.2 seconds for both parameter synthesis approaches. Both integer valued and real valued results for the parameters are reported in Table 1. Location l_1 becomes reachable when each simple constraint in $\mathcal{A}_{(c,p,M), <\Psi(\Delta) \setminus D, \Psi(Inv) \setminus I>}$ is relaxed by $\delta^* + 1$.

9.2. Experimental Results on Benchmarks from Literature. We collected 10 example models from the literature that include models with a safety specification that requires avoiding a set of locations L_T , and models with a reachability specification with a set of target locations L_T . In both cases, the original models satisfy the given specification. Eight of the examples are networks of TAs, and while a network of TAs can be represented as a single product TA and hence our methods can handle it, Tamus currently supports only

| Model | Source | Spec. | $ \Psi $ | $ \Psi^u $ | m |
|--------------|--------------------------------|--------|----------|------------|-----|
| accel1000 | [AHW18][HAF15] | reach. | 7690 | 13 | 3 |
| CAS | [ALN13] | reach. | 18 | 18 | 9 |
| coffee | [AKL ⁺ 19] | reach. | 10 | 10 | 3 |
| Jobshop4 | [AM01] | reach. | 64 | 48 | 5 |
| Pipeline3-3 | [KP10] | reach. | 41 | 41 | 12 |
| RCP | [CAS01] | reach. | 42 | 42 | 11 |
| SIMOP3 | [ACDS ⁺ 09] | reach. | 80 | 80 | 40 |
| Fischer | [HRSV01] | safety | 24 | 16 | 0 |
| JLR13-3tasks | [JLR13][ALNS15] | safety | 42 | 36 | 0 |
| WFAS | [BBLS15][FAWD ⁺ 14] | safety | 32 | 24 | 0 |

TABLE 2. Properties of the benchmarks, where $|\Psi|$ is as defined in Table 1, $|\Psi^u|$ is the number of constraints considered in the analysis and m is the number of mutated constraints.

MSR and MG computations for networks of TA, but not MILP relaxation. The properties of these models are summarized in Table 2.

MSR analysis. For the safety specifications, we define L_T as the target set and apply our methods for MSRs. Here, we find the minimal number of timing constants that should be changed to reach L_T , i.e., to violate the original safety specification. On the other hand, for reachability specifications, inspired by mutation testing [ALN13], we change a number of constraints on the original model so that L_T becomes unreachable. The number of mutated constraints are shown in Table 2.

The MSR results are shown in Table 3. Tamus computed a minimum MSR for all the models and also provided the MILP relaxation for the non-network models. Note that the bottleneck of our approach is the MSR computation and especially the verifier calls; the MILP part always took only few milliseconds (including models from Table 1), thus we believe that it would be also the case for the networks of TAs. The base variant of Imitator that computes the set of all satisfying parameter valuations solved only 4 of the 10 models. When run with the early termination flag, Imitator solved 3 more models, however, as discussed above, the provided solutions might not be optimal.

We have also evaluated Imitator for parameter synthesis based relaxation (Section 6.2). In particular, we first run Tamus to compute a minimum MSR $\mathcal{A}_{\langle D, I \rangle}$, then parameterized the constraints $D \cup I$ in the original TA \mathcal{A} , and run Imitator on the parameterized TA. In this case, Imitator solved 9 out of 10 models. Moreover, we have the guarantee that we found the optimal solution: the minimum MSR ensures that we relax the minimum number of simple clock constraints, and Imitator finds all satisfying parameterizations of the constraints hence also the one with minimum cumulative change of timing constants.

MG analysis. In order to apply the developed methods for the guarantee sets, we use L_T as the set of unsafe locations for both safety and reachability specifications. For the safety specifications, the minimal MG is the minimal set of constraints that are effective in avoiding the unsafe behaviors as intended. On the other hand, for the reachability specifications, as

| Model | d_R | v_R | t_R | c_R | c_I | t_R^I | t_R^{IT} | t_R^{Iw} | t_R^{ITw} |
|--------------|-------|-------|-------|-------|-------|---------|------------|------------|-------------|
| accel1000 | 2 | 22 | 1.50 | - | 5850 | 184.5 | 1.73 | 1.47 | 0.88 |
| CAS | 2 | 46 | 0.36 | 16 | 16 | 0.80 | 0.10 | 0.06 | 0.02 |
| coffee | 2 | 18 | 0.09 | 14 | 14 | 0.02 | 0.008 | 0.01 | 0.007 |
| Jobshop4 | 5 | 272 | 2.62 | - | 5 | to | 305.1 | to | 301.1 |
| Pipeline3-3 | 1 | 42 | 0.43 | - | 2 | to | 0.04 | to | 0.03 |
| RCP | 1 | 99 | 1.69 | - | 253 | to | 0.04 | 34.38 | 0.03 |
| SIMOP3 | 6 | 833 | 14.24 | - | 3755 | to | 3.98 | to | 0.37 |
| Fischer | 1 | 14 | 0.09 | - | - | to | to | 0.17 | 0.03 |
| JLR13-3tasks | 1 | 40 | 0.51 | - | 20 | to | 3.24 | 0.31 | 0.10 |
| WFAS | 1 | 10 | 0.09 | - | 6 | 14.69 | 0.02 | 14.66 | 0.03 |

TABLE 3. Experimental MSR analysis results for the benchmarks, where d_R , v_R , t_R and c_R are as defined in Table 1. c_I is the minimal total change $\sum_{(v,\varphi) \in D \cup I} \mathbf{P}^*(p_{v,\varphi})$ for the MSR as explained in Section 6.2. t_R^I , t_R^{IT} , t_R^{Iw} and t_R^{ITw} are the Imitator computation times, where w indicates that the early termination flag (“witness”) is used, otherwise the largest set of parameters is searched, and T indicates that only the constraints from the MSR identified by Tamus are parametrized, otherwise all constraints from Ψ^u are parametrized. *to* shows that the timeout limit is reached (20 min.).

| Model | d_G | v_G | t_G | c_G | t_G^{IT} | $\delta_{\mathbb{R}}^*$ | δ^* | t_G^{ITS} |
|--------------|-------|-------|-------|-------|------------|-------------------------|------------|-------------|
| accel1000 | 1 | 23 | 1.84 | 1557 | 1.3 | 1557.99 | 1557 | 1.2 |
| CAS | 2 | 39 | 0.30 | 11 | 0.004 | 5.99 | 5 | 0.003 |
| coffee | 2 | 25 | 0.12 | 7 | 0.005 | 3.99 | 3 | 0.002 |
| Jobshop4 | 2 | 281 | 2.94 | 0 | 4.44 | 0.49 | 0 | 3.47 |
| Pipeline3-3 | 2 | 82 | 0.67 | 1 | 0.04 | 0.99 | 0 | 0.03 |
| RCP | 8 | 255 | 4.46 | 463 | 0.32 | 8.99 | 8 | 0.03 |
| SIMOP3 | 2 | 214 | 3.81 | 417 | 0.003 | 208.99 | 208 | 0.002 |
| Fischer | 4 | 46 | 0.27 | 2 | 0.05 | 0.50 | 0 | 0.03 |
| JLR13-3tasks | 10 | 101 | 1.26 | - | to | 7.99 | 7 | 14.12 |
| WFAS | 5 | 66 | 0.47 | 0 | 0.06 | 0.00 | 0 | 0.018 |

TABLE 4. Experimental MG analysis results for the benchmarks, where d_G , v_G , t_G , c_G , $\delta_{\mathbb{R}}^*$, δ^* , t_G^{IT} and t_G^{ITS} are as defined in Table 1.

in the MSR analysis, we mutate a number of constraints on the original model so that L_T becomes unreachable and then apply our MG methods.

The MG results are shown in Table 4. Tamus computed a minimum MG for all the models. The maximal total change and the robustness degrees are generated using Imitator. Imitator found robustness degrees for all of the models, whereas it created maximal total change results for 9 out of 10 models. The running time for Imitator for both parameter synthesis approaches is below 15 seconds for all the results it produced.

10. CONCLUSION

We proposed the novel concept of a minimum MSR for a TA, i.e., a minimal set of simple clock constraints that need to be relaxed to satisfy a given specification. Moreover, we developed efficient techniques to find a minimum MSR, and presented MILP and parameter synthesis methods how to further tune the constraints in the MSR. We also introduced the concept of a maximum MIR, i.e., a maximal set of simple clock constraints that can be removed from the TA without violating the specification. Dually, one can represent an MIR via its complementary MG, i.e., a minimal set of simple clock constraints that need to be left in the TA to ensure that the specification is not violated. Moreover, we proposed parameter synthesis based approaches that can further relax the constraints in the MG while still keeping the specification satisfied.

Our empirical analysis showed that our tool, Tamus, can generate minimum MSRs and minimum MGs within seconds even for large systems. For the task of MSR relaxation, we have shown that the MILP method is faster than the parameter synthesis approach (MSR + Imitator). However, the MILP approach minimizes the cumulative change of the constraints from a minimum MSR by considering a single witness path. If the goal is to find a minimal relaxation globally, i.e., w.r.t. all witness paths for the MSR, we recommend using the combined version of MSR and Imitator, i.e., first run Tamus to find a minimum MSR, parametrize each constraint from the MSR and run Imitator to find all satisfying parameter valuations, including the global optimum.

ACKNOWLEDGMENT

This research was supported in part by ERDF “CyberSecurity, CyberCrime and Critical Information Infrastructures Center of Excellence” (No. CZ.02.1.01/0.0/0.0/16_019/0000822) and in part by the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 798482.

REFERENCES

- [AAGR19] Étienne André, Paolo Arcaini, Angelo Gargantini, and Marco Radavelli. Repairing timed automata clock guards through abstraction and testing. In Dirk Beyer and Chantal Keller, editors, *Tests and Proofs*, pages 129–146, Cham, 2019. Springer International Publishing. doi:10.1007/978-3-030-31157-5_9.
- [ACDS⁺09] Étienne André, Thomas Chatain, Olivier De Smet, Laurent Fribourg, and Silvain Ruel. Synthèse de contraintes temporisées pour une architecture d’automatisation en réseau. *Journal Européen des Systèmes Automatisés*, 43, November 2009. doi:10.3166/jesa.43.1049-1064.
- [AD94] Rajeev Alur and David L Dill. A theory of timed automata. *Theoretical computer science*, 126(2):183–235, 1994. doi:10.1016/0304-3975(94)90010-8.
- [AFKS12] Étienne André, Laurent Fribourg, Ulrich Kühne, and Romain Soulat. Imitator 2.5: A tool for analyzing robustness in scheduling problems. In Dimitra Giannakopoulou and Dominique Méry, editors, *Formal Methods*, pages 33–36, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. doi:10.1007/978-3-642-32759-9_6.
- [AFMS19] Étienne André, Laurent Fribourg, Jean-Marc Mota, and Romain Soulat. Verification of an industrial asynchronous leader election algorithm using abstractions and parametric model checking. In Constantin Enea and Ruzica Piskac, editors, *Verification, Model Checking, and Abstract Interpretation*, pages 409–424, Cham, 2019. Springer International Publishing. doi:10.1007/978-3-030-11245-5_19.

- [AHW18] Étienne André, Ichiro Hasuo, and Masaki Waga. Offline timed pattern matching under uncertainty. In *International Conference on Engineering of Complex Computer Systems*, pages 10–20. IEEE Computer Society, 2018. doi:10.1109/ICECCS2018.2018.00010.
- [AKL⁺19] Étienne André, Michal Knapik, Didier Lime, Wojciech Penczek, and Laure Petrucci. Parametric verification: An introduction. *Trans. Petri Nets Other Model. Concurr.*, 14:64–100, 2019. doi:10.1007/978-3-662-60651-3_3.
- [ALN13] Bernhard K. Aichernig, Florian Lorber, and Dejan Ničković. Time for mutants — model-based mutation testing with timed automata. In Margus Veanes and Luca Viganò, editors, *Tests and Proofs*, pages 20–38, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. doi:10.1007/978-3-642-38916-0_2.
- [ALNS15] Étienne André, Giuseppe Lipari, Hoang Gia Nguyen, and Youcheng Sun. Reachability preservation based parameter synthesis for timed automata. In Klaus Havelund, Gerard Holzmann, and Rajeev Joshi, editors, *NASA Formal Methods*, pages 50–65, Cham, 2015. Springer International Publishing. doi:10.1007/978-3-319-17524-9_5.
- [Alu99] Rajeev Alur. Timed automata. In *International Conference on Computer Aided Verification*, pages 8–22. Springer, 1999. doi:10.1007/3-540-48683-6_3.
- [AM01] Yasmina Abdeddaïm and Oded Maler. Job-shop scheduling using timed automata. In Gérard Berry, Hubert Comon, and Alain Finkel, editors, *International Conference on Computer Aided Verification*, pages 478–492, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg. doi:10.1007/3-540-44585-4_46.
- [And19a] Étienne André. A benchmark library for parametric timed model checking. In Cyrille Artho and Peter Csaba Ölveczky, editors, *Formal Techniques for Safety-Critical Systems*, pages 75–83, Cham, 2019. Springer International Publishing. doi:10.1007/978-3-030-12988-0_5.
- [And19b] Étienne André. What’s decidable about parametric timed automata? *Int. J. Softw. Tools Technol. Transf.*, 21(2):203–219, April 2019. doi:10.1007/s10009-017-0467-0.
- [AS11] Étienne André and Romain Soulat. Synthesis of timing parameters satisfying safety properties. In Giorgio Delzanno and Igor Potapov, editors, *Reachability Problems - 5th International Workshop, RP 2011, Genoa, Italy, September 28-30, 2011. Proceedings*, volume 6945 of *LNCS*, pages 31–44. Springer, 2011. doi:10.1007/978-3-642-24288-5_5.
- [BBB⁺16] Jiří Barnat, Petr Bauch, Nikola Beneš, Luboš Brim, Jan Beran, and Tomáš Kratochvíla. Analysing sanity of requirements for avionics systems. *Formal Aspects of Computing 28*, pages 1–19, 2016. doi:10.1007/s00165-015-0348-9.
- [BBBČ16] Peter Bezděk, Nikola Beneš, Jiří Barnat, and Ivana Černá. LTL parameter synthesis of parametric timed automata. In Rocco De Nicola and Eva Kühn, editors, *Software Engineering and Formal Methods*, pages 172–187, Cham, 2016. Springer International Publishing. doi:10.1007/978-3-319-41591-8_12.
- [BBBR07] Patricia Bouyer, Thomas Brihaye, Véronique Bruyère, and Jean-François Raskin. On the optimal reachability problem of weighted timed automata. *Formal Methods in System Design*, 31:135–175, 2007. doi:10.1007/s10703-007-0035-4.
- [BBČB16] Jaroslav Bendík, Nikola Beneš, Ivana Černá, and Jiří Barnat. Tunable online MUS/MSS enumeration. In *36th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, volume 65 of *LIPICs*, pages 50:1–50:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPICs.FSTTCS.2016.50.
- [BBČB18] Peter Bezděk, Nikola Beneš, Ivana Černá, and Jiří Barnat. On clock-aware LTL parameter synthesis of timed automata. *J. Log. Algebraic Methods Program.*, 99:114–142, 2018. doi:10.1016/j.jlamp.2018.05.004.
- [BBL15] Nikola Beneš, Peter Bezděk, Kim Guldstrand Larsen, and Jiri Srba. Language emptiness of continuous-time parametric timed automata. In *International Colloquium on Automata, Languages, and Programming*, volume 9135 of *LNCS*, pages 69–81. Springer, 2015. doi:10.1007/978-3-662-47666-6_6.
- [BČ18] Jaroslav Bendík and Ivana Černá. Evaluation of domain agnostic approaches for enumeration of minimal unsatisfiable subsets. In *LPAR*, volume 57 of *EPiC Series in Computing*, pages 131–142. EasyChair, 2018. doi:10.29007/sxzb.

- [BČ20a] Jaroslav Bendík and Ivana Černá. Replication-guided enumeration of minimal unsatisfiable subsets. In *International Conference on Principles and Practice of Constraint Programming*, volume 12333 of *LNCS*, pages 37–54. Springer, 2020. doi:10.1007/978-3-030-58475-7_3.
- [BČ20b] Jaroslav Bendík and Ivana Černá. Rotation based MSS/MCS enumeration. In *LPAR*, volume 73 of *EPiC Series in Computing*, pages 120–137. EasyChair, 2020. doi:10.29007/8btb.
- [BČB18] Jaroslav Bendík, Ivana Černá, and Nikola Beneš. Recursive online enumeration of all minimal unsatisfiable subsets. In *International symposium on automated technology for verification and analysis*, volume 11138 of *LNCS*, pages 143–159. Springer, 2018. doi:10.1007/978-3-030-01090-4_9.
- [BDL⁺06] Gerd Behrmann, Alexandre David, Kim G. Larsen, John Hakansson, Paul Petterson, Wang Yi, and Martijn Hendriks. Uppaal 4.0. In *Proceedings of the 3rd International Conference on the Quantitative Evaluation of Systems*, QEST '06, pages 125–126, Washington, DC, USA, 2006. IEEE Computer Society. doi:10.1109/QEST.2006.59.
- [Ben17] Jaroslav Bendík. Consistency checking in requirements analysis. In *Proceedings of the 26th ACM SIGSOFT international symposium on software testing and analysis*, pages 408–411. ACM, 2017. doi:10.1145/3092703.3098239.
- [Ben21] Jaroslav Bendík. *Minimal Sets over a Monotone Predicate: Enumeration and Counting*. PhD thesis, Masaryk University, 2021. URL: <https://is.muni.cz/th/y4v8m/dissertationRevised.pdf>.
- [BGWČ18] Jaroslav Bendík, Elaheh Ghassabani, Michael W. Whalen, and Ivana Černá. Online enumeration of all minimal inductive validity cores. In *International Conference on Software Engineering and Formal Methods*, volume 10886 of *LNCS*, pages 189–204. Springer, 2018. doi:10.1007/978-3-319-92970-5_12.
- [BK16] Fahiem Bacchus and George Katsirelos. Finding a collection of muses incrementally. In *International Conference on AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*, volume 9676 of *LNCS*, pages 35–44. Springer, 2016. doi:10.1007/978-3-319-33954-2_3.
- [BMS13] Patricia Bouyer, Nicolas Markey, and Ocan Sankur. Robustness in timed automata. In Parosh Aziz Abdulla and Igor Potapov, editors, *Reachability Problems - 7th International Workshop, RP 2013, Uppsala, Sweden, September 24-26, 2013 Proceedings*, volume 8169 of *LNCS*, pages 1–18. Springer, 2013. doi:10.1007/978-3-642-41036-9_1.
- [BS05] James Bailey and Peter J. Stuckey. Discovery of minimal unsatisfiable subsets of constraints using hitting set dualization. In *International Workshop on Practical Aspects of Declarative Languages*, pages 174–186. Springer, 2005. doi:10.1007/978-3-540-30557-6_14.
- [BSGČ21] Jaroslav Bendík, Ahmet Sencan, Ebru Aydin Gol, and Ivana Černá. Timed automata relaxation for reachability. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, volume 12651 of *LNCS*, pages 291–310. Springer, 2021. doi:10.1007/978-3-030-72016-2_16.
- [CAS01] Aurore Collomb-Annichini and Mihaela Sighireanu. Parameterized reachability analysis of the IEEE 1394 root contention protocol using trex. 08 2001.
- [DILS09] Alexandre David, Jacob Illum, Kim G Larsen, and Arne Skou. Model-based framework for schedulability analysis using UPPAAL 4.1. In *Model-based design for embedded systems*, pages 117–144. 2009.
- [dIBSW03] Maria García de la Banda, Peter J. Stuckey, and Jeremy Wazny. Finding all minimal unsatisfiable subsets. In *Proceedings of the 5th ACM SIGPLAN international conference on Principles and practice of declarative programming*, pages 32–43. ACM, 2003. doi:10.1145/888251.888256.
- [EYG21] Mert Ergurtuna, Beyazit Yalcinkaya, and Ebru Aydin Gol. An automated system repair framework with signal temporal logic. *Acta Informatica*, pages 1–1, 2021. doi:10.1007/s00236-021-00403-z.
- [FAWD⁺14] Sergio Feo-Arenis, Bernd Westphal, Daniel Dietsch, Marco Muñoz, and Ahmad Siyar Andisha. The wireless fire alarm system: Ensuring conformance to industrial standards through formal verification. In Cliff Jones, Pekka Pihlajasaari, and Jun Sun, editors, *Formal Methods*, pages 658–672, Cham, 2014. Springer International Publishing. doi:10.1007/978-3-319-06410-9_44.
- [Feh99] A. Fehnker. Scheduling a steel plant with timed automata. In *Proceedings Sixth International Conference on Real-Time Computing Systems and Applications*, pages 280–286, 1999. doi:10.1109/RTCSA.1999.811256.

- [GBST14] Kahina Gani, Marinette Bouet, Michel Schneider, and Farouk Toumani. Formal modeling and analysis of home care plans. In Xavier Franch, Aditya K. Ghose, Grace A. Lewis, and Sami Bhiri, editors, *Service-Oriented Computing - 12th International Conference*, volume 8831 of *LNCS*, pages 494–501. Springer, 2014. doi:10.1007/978-3-662-45391-9_41.
- [GGD⁺07] Nan Guan, Zonghua Gu, Qingxu Deng, Shuaihong Gao, and Ge Yu. Exact schedulability analysis for static-priority global multiprocessor scheduling using model-checking. In *IFIP International Workshop on Software Technologies for Embedded and Ubiquitous Systems*, pages 263–272, 2007. doi:10.1007/978-3-540-75664-4_26.
- [GWG17] Elaheh Ghassabani, Michael W. Whalen, and Andrew Gacek. Efficient generation of all minimal inductive validity cores. In *Formal Methods in Computer Aided Design*, pages 31–38. IEEE, 2017. doi:10.23919/FMCD.2017.8102238.
- [HAF15] Bardh Hoxha, Houssam Abbas, and Georgios Fainekos. Benchmarks for temporal logic requirements for automotive systems. In Goran Frehse and Matthias Althoff, editors, *ARCH14-15. 1st and 2nd International Workshop on Applied verification for Continuous and Hybrid Systems*, volume 34 of *EPiC Series in Computing*, pages 25–30. EasyChair, 2015. doi:10.29007/xwrs.
- [HPW01] T. A. Henzinger, J. Preussig, and H. Wong-Toi. Some lessons from the hytech experience. In *Proceedings of the 40th IEEE Conference on Decision and Control (Cat. No.01CH37228)*, volume 3, pages 2887–2892 vol.3, 2001.
- [HRSV01] Thomas Hune, Judi Romijn, Mariëlle Stoelinga, and Frits Vaandrager. Linear parametric model checking of timed automata. In Tiziana Margaria and Wang Yi, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, pages 189–203, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg. doi:10.1007/3-540-45319-9_14.
- [IJM16] Alexey Ignatiev, Mikoláš Janota, and João Marques-Silva. Quantified maximum satisfiability. *Constraints An Int. J.*, 21(2):277–302, 2016. doi:10.1007/s10601-015-9195-9.
- [IMMV16] Alexander Ivrii, Sharad Malik, Kuldeep S. Meel, and Moshe Y. Vardi. On computing minimal independent support and its applications to sampling and counting. *Constraints An Int. J.*, 21(1):41–58, 2016. doi:10.1007/s10601-015-9204-z.
- [IPLM15] Alexey Ignatiev, Alessandro Previti, Mark H. Liffiton, and João Marques-Silva. Smallest MUS extraction with minimal hitting set dualization. In *International Conference on Principles and Practice of Constraint Programming*, volume 9255 of *LNCS*, pages 173–182. Springer, 2015. doi:10.1007/978-3-319-23219-5_13.
- [JLR13] Aleksandra Jovanović, Didier Lime, and Olivier H. Roux. Integer parameter synthesis for timed automata. In Nir Piterman and Scott A. Smolka, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, pages 401–415, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. doi:10.1007/978-3-642-36742-7_28.
- [JLR15] A. Jovanovic, D. Lime, and O. H. Roux. Integer parameter synthesis for real-time systems. *IEEE Transactions on Software Engineering*, 41(5):445–461, 2015. doi:10.1109/TSE.2014.2357445.
- [JPAM14] Zhihao Jiang, Miroslav Pajic, Rajeev Alur, and Rahul Mangharam. Closed-loop verification of medical devices with model abstraction and refinement. *Int. J. Softw. Tools Technol. Transf.*, 16(2):191–213, April 2014. doi:10.1007/s10009-013-0289-7.
- [KLW19] Martin Kölbl, Stefan Leue, and Thomas Wies. Clock bound repair for timed systems. In Isil Dillig and Serdar Tasiran, editors, *International Conference on Computer Aided Verification*, pages 79–96, Cham, 2019. Springer International Publishing. doi:10.1007/978-3-030-25540-4_5.
- [KMPP15] Marta Kwiatkowska, Alexandru Mereacre, Nicola Paoletti, and Andrea Patanè. Synthesising robust and optimal parameters for cardiac pacemakers using symbolic and evolutionary computation techniques. In Alessandro Abate and David Šafránek, editors, *Hybrid Systems Biology*, pages 119–140, Cham, 2015. Springer International Publishing. doi:10.1007/978-3-319-26916-0_7.
- [KP10] Michal Knapik and Wojciech Penczek. Bounded model checking for parametric timed automata. *Trans. Petri Nets Other Model. Concurr.*, 5:141–159, 2010.
- [LML⁺09] Mark H. Liffiton, Maher N. Mneimneh, Inês Lynce, Zaher S. Andraus, João Marques-Silva, and Karem A. Sakallah. A branch and bound algorithm for extracting smallest minimal unsatisfiable subformulas. *Constraints An Int. J.*, 14(4):415–442, 2009. doi:10.1007/s10601-008-9058-8.
- [LPMM16] Mark H. Liffiton, Alessandro Previti, Ammar Malik, and João Marques-Silva. Fast, flexible MUS enumeration. *Constraints*, 21(2):223–250, 2016. doi:10.1007/s10601-015-9183-0.

- [LRST09] Didier Lime, Olivier H. Roux, Charlotte Seidner, and Louis-Marie Traonouez. Romeo: A parametric model-checker for petri nets with stopwatches. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, volume 5505 of *LNCS*, pages 54–57. Springer, 2009. doi:10.1007/978-3-642-00768-2_6.
- [LY93] Kim G Larsen and Wang Yi. Time abstracted bisimulation: Implicit specifications and decidability. In *International Conference on Mathematical Foundations of Programming Semantics*, pages 160–176. Springer, 1993. doi:10.1006/inco.1997.2623.
- [MHJ⁺13] João Marques-Silva, Federico Heras, Mikolás Janota, Alessandro Previti, and Anton Belov. On computing minimal correction subsets. In *Twenty-Third International Joint Conference on Artificial Intelligence*, pages 615–622. IJCAI/AAAI, 2013.
- [MJB13] João Marques-Silva, Mikolás Janota, and Anton Belov. Minimal sets over monotone predicates in boolean formulae. In *International Conference on Computer Aided Verification*, volume 8044 of *LNCS*, pages 592–607. Springer, 2013. doi:10.1007/978-3-642-39799-8_39.
- [MJM17] João Marques-Silva, Mikolás Janota, and Carlos Mencía. Minimal sets on propositional formulae. problems and reductions. *Artif. Intell.*, 252:22–50, 2017. doi:10.1016/j.artint.2017.07.005.
- [PF] Laurent Perron and Vincent Furnon. Or-tools. URL: <https://developers.google.com/optimization/>.
- [San13] Ocan Sankur. Shrinktech: A tool for the robustness analysis of timed automata. In Natasha Sharygina and Helmut Veith, editors, *International Conference on Computer Aided Verification*, volume 8044 of *LNCS*, pages 1006–1012. Springer, 2013. doi:10.1007/978-3-642-39799-8_72.
- [SBM11] Ocan Sankur, Patricia Bouyer, and Nicolas Markey. Shrinking timed automata. In Supratik Chakraborty and Amit Kumar, editors, *Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, volume 13 of *LIPICs*, pages 90–102. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2011. doi:10.4230/LIPICs.FSTTCS.2011.90.
- [SKFP12] Roni Tzvi Stern, Meir Kalech, Alexander Feldman, and Gregory M. Provan. Exploring the duality in conflict-directed model-based diagnosis. In *Twenty-Sixth AAAI Conference on Artificial Intelligence*. AAAI Press, 2012.
- [Spe28] Emanuel Sperner. Ein satz über untermengen einer endlichen menge. *Mathematische Zeitschrift*, 27(1):544–548, 1928.
- [Wan04] Farn Wang. Formal verification of timed systems: a survey and perspective. *Proceedings of the IEEE*, 92(8):1283–1305, Aug 2004. doi:10.1109/JPROC.2004.831210.
- [WDMR08] Martin De Wulf, Laurent Doyen, Nicolas Markey, and Jean-François Raskin. Robust safety of timed automata. *Formal Methods Syst. Des.*, 33(1-3):45–84, 2008. doi:10.1007/s10703-008-0056-7.