



Efficient Timed Diagnosis Using Automata with Timed Domains

Patricia Bouyer¹, Samy Jaziri¹, and Nicolas Markey²(✉)

¹ LSV – CNRS & ENS Paris-Saclay, Univ. Paris-Saclay, Cachan, France

² IRISA – Univ. Rennes & CNRS & INRIA, Rennes, France

nmarkey@irisa.fr

Abstract. We consider the problems of efficiently diagnosing and predicting what did (or will) happen in a partially-observable one-clock timed automaton. We introduce *timed sets* as a formalism to keep track of the evolution of the reachable configurations over time, and use our previous work on automata over timed domains to build a candidate diagnoser for our timed automaton. We report on our implementation of this approach compared to the approach of [Tripakis, *Fault diagnosis for timed automata*, 2002].

1 Introduction

Formal Methods in Verification. Because of the wide range of applications of computer systems, and of their increasing complexity, the use of formal methods for checking their correct behaviours has become essential [10, 16]. Numerous approaches have been introduced and extensively studied over the last 40 years, and mature tools now exist and are used in practice. Most of these approaches rely on building mathematical models, such as automata and extensions thereof, in order to represent and reason about the behaviours of those systems; various algorithmic techniques are then applied in order to ensure correctness of those behaviours, such as *model checking* [11, 12], *deductive verification* [13, 17] or *testing* [25].

Fault Diagnosis. The techniques listed above mainly focus on assessing correctness of the set of all behaviours of the system, in an offline manner. This is usually very costly in terms of computation, and sometimes too strong a requirement. *Runtime verification* instead aims at checking properties of a running system [19]. *Fault diagnosis* is a prominent problem in runtime verification: it consists in (deciding the existence and) building a diagnoser, whose role is to monitor real executions of a (partially-observable) system, and decide *online* whether some property holds (e.g., whether some unobservable fault has occurred) [24, 26]. A diagnoser can usually be built (for finite-state models) by determinizing a model of the system, using the powerset construction; it will keep track of all

Work supported by ERC project EQualIS.

© Springer Nature Switzerland AG 2018

C. Colombo and M. Leucker (Eds.): RV 2018, LNCS 11237, pp. 205–221, 2018.

https://doi.org/10.1007/978-3-030-03769-7_12

possible states that can be reached after each (observable) step of the system, thereby computing whether a fault may or must have occurred. The related problem of *prediction*, a.k.a. prognosis, (that e.g. no fault may occur in the next five steps) [15], is also of particular interest in runtime verification, and can be solved using similar techniques.

Verifying Real-Time Systems. Real-time constraints often play an important role for modelling and specifying correctness of computer systems. Discrete models, such as finite-state automata, are not adequate to model such real-time constraints; timed automata [1], developed at the end of the 1980's, provide a convenient framework for both representing and efficiently reasoning about computer systems subject to real-time constraints. Efficient offline verification techniques for timed automata have been developed and implemented [2, 3]. Diagnosis of timed automata however has received less attention; this problem is made difficult by the fact that timed automata can in general not be determinized [14, 27]. This has been circumvented by either restricting to classes of determinizable timed automata [6], or by keeping track of all possible configurations of the automaton after a (finite) execution [26]. The latter approach is computationally very expensive, as one step consists in maintaining the set of all configurations that can be reached by following (arbitrarily long) sequences of unobservable transitions; this limits the applicability of the approach.

Our Contribution. In this paper, we (try to) make the approach of [26] more efficiently applicable (over the class of one-clock timed automata). Our improvements are based on two ingredients: first, we use *automata over timed domains* [7] as a model for representing the diagnoser. Automata over timed domains can be seen as an extension of timed automata with a (much) richer notion of time and clocks; these automata enjoy determinizability. The second ingredient is the notion of *timed sets*: timed sets are pairs (E, I) where E is any subset of \mathbb{R} , and I is an interval with upper bound $+\infty$; such a timed set represents a set of clock valuations evolving over time: the timed set $(E; I)$ after a delay d represents the set $(E + d) \cap I$. As we prove, timed sets can be used to finitely represent the evolution of the set of all reachable configurations after a finite execution.

In the end, our algorithm can compute a finite representation of the reachable configurations after a given execution, as well as all the configurations that can be reached from there after any delay. This can be used to very quickly update the set of current possible configurations (which would be expensive with the approach of [26]). Besides diagnosis, this can also be used to efficiently predict the occurrence of faults occurring after some delay (which is not possible in [26]). We implemented our technique in a prototype tool: as we report at the end of the paper, our approach requires heavy precomputation, but can then efficiently handle delay transitions.

Related Works. Model-based diagnosis has been extensively studied in the community of discrete-event systems [23, 24, 28]. This framework gave birth to a number of ramifications (e.g. active diagnosis [22], fault prediction [15], opacity [18]),

and was applied in many different settings besides discrete-event systems (e.g. Petri nets, distributed systems [4], stochastic systems [5,20], discrete-time systems [9], hybrid systems [21]).

Much fewer papers have focused on continuous-time diagnosis: Tripakis proposed an algorithm for deciding diagnosability [26]. Cassez developed a uniform approach for diagnosability of discrete and continuous time systems, as a reduction to Büchi-automata emptiness [8]. A construction of a diagnoser for timed systems is proposed in [26]: the classical approach of determinizing using the powerset construction does not extend to timed automata, because timed automata cannot in general be determinized [14,27]. Tripakis proposed the construction of a diagnoser as an online algorithm that keeps track of the possible states and zones the system can be in after each event (or after a sufficiently-long delay), which requires heavy computation at each step and is hardly usable in practice. Bouyer, Chevalier and D’Souza studied a restricted setting, only looking for diagnosers under the form of deterministic timed automata with limited resources [6].

2 Definitions

2.1 Intervals

In this paper, we heavily use intervals, and especially unbounded ones. For $r \in \mathbb{R}$, we define

$$\uparrow r = [r; +\infty) \quad \uparrow r = (r; +\infty) \quad \downarrow r = (-\infty; r) \quad \downarrow r = (-\infty; r].$$

We let $\widehat{\mathbb{R}}_{\geq 0} = \{\uparrow r, \uparrow r \mid r \in \mathbb{R}_{\geq 0}\}$ for the set of upward-closed intervals of $\mathbb{R}_{\geq 0}$; in the sequel, elements of $\widehat{\mathbb{R}}_{\geq 0}$ are denoted with \widehat{r} . Similarly, we let $\mathbb{R}_{\geq 0} = \{\downarrow r, \downarrow r \mid r \in \mathbb{R}_{\geq 0}\}$, and use notation \mathcal{r} for intervals in $\mathbb{R}_{\geq 0}$. The elements of $\widehat{\mathbb{R}}_{\geq 0}$ can be (totally) ordered using inclusion: we write $\widehat{r} \prec \widehat{r}'$ whenever $\widehat{r}' \subset \widehat{r}$ (so that $r < r'$ entails $\uparrow r \prec \uparrow r'$).

2.2 Timed Automata

Let Σ be a finite alphabet.

Definition 1. *A one-clock timed automaton over Σ is a tuple $\mathcal{A} = (S, \{s_0\}, T, F)$, where S is a finite set of states, $s_0 \in S$ is the initial state, $T \subseteq S \times \widehat{\mathbb{R}}_{\geq 0} \times \mathbb{R}_{\geq 0} \times (\Sigma \uplus \{\epsilon\}) \times \{0, id\} \times S$ is the set of transitions, and $F \subseteq S$ is a set of final states. A configuration of \mathcal{A} is a pair $(s, v) \in S \times \mathbb{R}_{\geq 0}$. There is a transition from (s, v) to (s', v') if*

- either $s' = s$ and $v' \geq v$. In that case, we write $(s, v) \xrightarrow{d} (s, v')$, with $d = v' - v$, for such delay transitions (notice that we have no invariants);

- or there is a transition $e = (s, \widehat{l}, \underline{y}, a, r, s')$ s.t. $v \in \widehat{l} \cap \underline{y}$ and $v' = v$ if $r = id$, and $v' = 0$ otherwise. For those action transitions, we write $(s, v) \rightarrow_e (s', v')$. We assume that for each transition $e = (s, \widehat{l}, \underline{y}, a, r, s')$, it holds $\widehat{l} \cap \underline{y} \neq \emptyset$.

Fix a one-clock timed automaton \mathcal{A} ; we write T_{id} for the set of *non-resetting* transitions, i.e., having *id* as their fifth component, and T_0 for the complement set of *resetting* transitions.

For a transition $e = (s, \widehat{l}, \underline{y}, a, r, s')$, we write \widehat{e} and \underline{e} for \widehat{l} and \underline{y} , respectively. We write $\text{src}(e) = s$ and $\text{tgt}(e) = s'$, and $\text{lab}(e) = a \in \Sigma$. We extend these definitions to sequences of transitions $w = (e_i)_{0 \leq i < n}$ as $\text{src}(w) = \text{src}(e_0)$, $\text{tgt}(w) = \text{tgt}(e_{n-1})$, and $\text{lab}(w) = (\text{lab}(e_i))_{0 \leq i < n}$.

Let w be a sequence $(e_{2i+1})_{0 \leq 2i+1 < n}$ of transitions of T , and $d \in \mathbb{R}_{\geq 0}$. We write $(s, v) \xrightarrow{d}_w (s', v')$ if there exist finite sequences $(s_i, v_i)_{0 \leq i \leq n} \in (S \times \mathbb{R}_{\geq 0})^{n+1}$ and $(d_{2i})_{0 \leq 2i < n} \in \mathbb{R}_{\geq 0}^{\lfloor n/2 \rfloor}$ such that $\sum_{0 \leq 2i < n} d_{2i} = d$, and $(s_0, v_0) = (s, v)$ and $(s_n, v_n) = (s', v')$, and for all $0 \leq j < n$, $(s_j, v_j) \xrightarrow{d_j} (s_{j+1}, v_{j+1})$ if j is even and $(s_j, v_j) \rightarrow_{e_j} (s_{j+1}, v_{j+1})$ if j is odd. We write $(s, v) \rightarrow (s', v')$ when $(s, v) \xrightarrow{d}_w (s', v')$ for some $w \in T^*$ and some $d \in \mathbb{R}_{\geq 0}$.

For any $\lambda \in \Sigma^*$ and any $d \in \mathbb{R}_{\geq 0}$, we write $(s, v) \xrightarrow{d}_\lambda (s', v')$ whenever there exists a sequence of transitions w such that $\lambda = \text{lab}(w)$ and $(s, v) \xrightarrow{d}_w (s', v')$. Notice that¹ $(s, v) \xrightarrow{d}_\perp (s', v')$ (sometimes simply written $(s, v) \xrightarrow{d} (s', v')$) indicates a delay-transition (hence it must be $s = s'$). The untimed language $\mathcal{L}(\mathcal{A})$ of \mathcal{A} is the set of words $\lambda \in \Sigma^*$ such that $(s_0, 0) \xrightarrow{d}_\lambda (s', v')$ for some $s' \in F$ and $d \in \mathbb{R}_{\geq 0}$.

We borrow some of the formalism of [7], in order to define a kind of *powerset construction* for timed automata. For a one-clock timed automaton $\mathcal{A} = (S, \{s_0\}, T, F)$ on Σ , we write $\mathbb{M} = (\mathcal{P}(\mathbb{R}_{\geq 0}))^S$ for the set of *markings*, mapping states of \mathcal{A} to sets of valuations for the unique clock of \mathcal{A} . For a marking $m \in \mathbb{M}$, we write $\text{supp}(m) = \{s \in S \mid m(s) \neq \emptyset\}$. For any $l \in \Sigma$, we define the function $\mathbf{O}_l: \mathbb{M} \rightarrow \mathbb{M}$ by letting, for any $m \in \mathbb{M}$ and any $s' \in S$,

$$\mathbf{O}_l(m): s' \mapsto \{v' \in \mathbb{R}_{\geq 0} \mid \exists s \in S. \exists v \in m(s). (s, v) \rightarrow_l (s', v')\}.$$

Similarly, for any $d \in \mathbb{R}_{\geq 0}$, we let

$$\mathbf{O}_d(m): s' \mapsto \{v' \in \mathbb{R}_{\geq 0} \mid \exists s \in S. \exists v \in m(s). (s, v) \xrightarrow{d} (s', v')\}.$$

Notice that \mathbf{O}_d simply shifts all valuations by d .

Definition 2 ([7]). *The powerset automaton of a timed automaton $\mathcal{A} = (S, \{s_0\}, T, F)$ is a tuple $\mathbf{DA} = (\mathcal{P}(S), \{\{s_0\}\}, \mathbf{PT}, \mathbf{PF})$, where $\mathcal{P}(S)$ is the set of states, $\{s_0\}$ is the initial state, $\mathbf{PT} = \{(q, m, a, q') \in \mathcal{P}(S) \times \mathbb{M} \times \Sigma \times \mathcal{P}(S) \mid q = \text{supp}(m), q' = \text{supp}(\mathbf{O}_a(m))\}$ is the set of transitions, and $\mathbf{PF} = \{E \in \mathcal{P}(S) \mid E \cap F \neq \emptyset\}$ is a set of final states.*

¹ In this paper, we write \perp for the empty word (or empty sequences) over any alphabet.

Configurations of \mathbf{DA} are all pairs $(q, m) \in \mathcal{P}(S) \times \mathbb{M}$ for which $q = \text{supp}(m)$. There is a transition from a configuration (q, m) to a configuration (q', m') labelled with $l \in \Sigma \cup \mathbb{R}_{\geq 0}$ whenever $m' = \mathbf{O}_l(m)$. We extend this definition to sequences alternating delay- and action transitions, and write $(q, m) \xrightarrow{d}_w (q', m')$ when there is a path from (q, m) to (q', m') following the transitions of w in d time units. Similarly, we write $(q, m) \xrightarrow{d}_\sigma (q', m')$ if $(q, m) \xrightarrow{d}_w (q', m')$ and $\text{lab}(w) = \sigma$.

Following [7], the automaton \mathbf{DA} is deterministic, and it simulates \mathcal{A} in the sense that given two marking m and m' and a word σ of Σ^* , we have $(\text{supp}(m), m) \xrightarrow{d}_\sigma (\text{supp}(m'), m')$ if, and only if, $m'(s') = \{v' \in \mathbb{R}_{\geq 0} \mid \exists s \in S. \exists v \in m(s). (s, v) \xrightarrow{d}_\sigma (s', v')\}$ for all $s' \in S$.

2.3 Timed Automata with Silent Transitions

The work reported in [7] only focuses on the case when there are no silent transitions. In that case, for any $d \in \mathbb{R}_{\geq 0}$, the operation \mathbf{O}_d can be computed easily, since it amounts to adding d to each item of the marking (in other terms, for any marking m , any state $s \in S$, and any $v \in \mathbb{R}_{\geq 0}$ such that $v + d \in \mathbb{R}_{\geq 0}$, we have $v \in m(s)$ if, and only if, $v + d \in \mathbf{O}_d(m)(s)$). This leads to an efficient expression of a (deterministic) powerset automaton simulating \mathcal{A} .

However fault diagnosis should deal with timed automata with unobservable transitions (unobservable transitions then correspond to internal transitions). So we now assume that Σ contains a special *silent letter* ϵ , whose occurrence is not *visible*. This requires changing the definition of lab : we now let

$$\begin{aligned} \text{lab}(\perp) &= \perp \\ \text{lab}(w \cdot e) &= \text{lab}(w) && \text{if } \text{lab}(e) = \epsilon \\ \text{lab}(w \cdot e) &= \text{lab}(w) \cdot \text{lab}(e) && \text{if } \text{lab}(e) \neq \epsilon \end{aligned}$$

Notice that $\text{lab}(w) \in (\Sigma \setminus \{\epsilon\})^*$. We may write \rightarrow_ϵ in place of \rightarrow_\perp , to stick to classical notations and make it clear that it allows silent transitions.

In that case, \mathbf{DA} still is a (deterministic) powerset automaton that simulates \mathcal{A} , and hence is still a *diagnoser*, but the function \mathbf{O}_d cannot be computed by just shifting valuations by d . In its raw form, the function \mathbf{O}_d can be obtained by the computation of the set of reachable configurations in a delay d by following silent transitions. This is analogous to the method proposed by Tripakis [26], and turns out to be very costly. However, a diagnoser must be able to simulate all possible actions of the diagnosed automaton quickly enough, so that it can be used at runtime. In this paper, we introduce a new data structure called *timed sets*, which we use to represent markings in the timed powerset automaton; as we explain, using timed sets we can compute \mathbf{O}_d more efficiently (at the expense of more precomputations).

3 (Regular) Timed Sets

3.1 Timed Sets

For a set E and a real d , we define the set $E + d = \{x + d \mid x \in E\}$. We introduce *timed sets* as a way to represent sets of clock valuations (and eventually markings), and their evolution over time.

Definition 3. An atomic timed set is a pair $(E; \hat{r})$ where $E \subseteq \mathbb{R}$ and $\hat{r} \in \widehat{\mathbb{R}}_{\geq 0}$. With an atomic timed set $F = (E; \hat{r})$, we associate a mapping $F: \mathbb{R}_{\geq 0} \rightarrow 2^{\mathbb{R}}$ defined as $F(d) = (E + d) \cap \hat{r}$.

We define the union of two timed sets F and F' , denoted as $F \sqcup F'$, as their pointwise union: $(F \sqcup F')(d) = F(d) \cup F'(d)$.

Definition 4. A timed set is a countable set $F = \{F_i \mid i \in I\}$ (intended to be a union, hence sometimes also denoted with $\bigsqcup_{i \in I} F_i$) of atomic timed sets. With such a timed set, we again associate a mapping $F: \mathbb{R}_{\geq 0} \rightarrow 2^{\mathbb{R}_{\geq 0}}$ defined as $F(d) = \bigcup_{i \in I} F_i(d)$. A timed set is finite when it is made of finitely many atomic timed sets. We write $\mathcal{T}(\mathbb{R})$ for the set of timed sets of \mathbb{R} .

Given two timed sets F and F' , we write $F \sqsubseteq F'$ whenever $F(d) \subseteq F'(d)$ for all $d \in \mathbb{R}_{\geq 0}$. This is a pre-order relation; it is not anti-symmetric as for instance $(\{1\}; \uparrow 0) \sqsubseteq (\{1\}; \uparrow 1)$ and $(\{1\}; \uparrow 1) \sqsubseteq (\{1\}; \uparrow 0)$. We write $F \equiv F'$ whenever $F \sqsubseteq F'$ and $F' \sqsubseteq F$.

Example 1. Figure 1 displays an example of an atomic timed set $F = (E; \uparrow 1)$, with $E = [-3; -1] \cup [0; 2]$. The picture displays the sets $F(0) = [1; 2]$ and $F(3) = [1; 2] \cup [3; 5]$. ◀

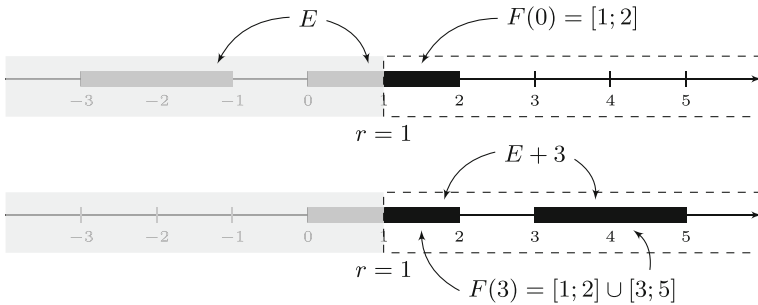


Fig. 1. Example of an atomic timed set $F = ([-3; -1] \cup [0; 2]; \uparrow 1)$.

3.2 Regular Timed Sets

In order to effectively store and manipulate timed sets, we need to identify a class of timed sets that is expressive enough but whose timed sets have a finite representation.

Definition 5. A regular union of intervals is a 4-tuple $E = (I, J, p, q)$ where I and J are finite unions of intervals of \mathbb{R} with rational (or infinite) bounds, $p \in \mathbb{Q}_{\geq 0}$ is the period, and $q \in \mathbb{N}$ is the offset. It is required that $J \subseteq (-p; 0]$ and $I \subseteq \uparrow(-q \cdot p)$.

The regular union of intervals $E = (I, J, p, q)$ represents the set (which we still write E) $I \cup \bigcup_{k=q}^{+\infty} J - k \cdot p$.

Regular unions of intervals enjoy the following properties:

Proposition 6. Let E and E' be regular unions of intervals, K be an interval, and $d \in \mathbb{Q}$. Then $E \cup E'$, \overline{E} , $E + d$ and $E - K$ are regular unions of intervals.

Definition 7. A regular timed set is a finite timed set $F = \{(E_i; \widehat{r}_i) \mid i \in I\}$ such that for all $i \in I$, the set E_i is a regular union of intervals.

4 Computing the Powerset Automaton

In this section, we fix a one-clock timed automaton $\mathcal{A} = (S, \{s_0\}, T, F)$ over alphabet Σ , assumed to contain a silent letter ϵ . We assume that some silent transitions are *faulty*, and we want to detect the occurrence of such faulty transitions based on the sequence of actions we can observe. Following [26], this can be reduced to a *state-estimation problem*, even if it means duplicating some states of the model in order to keep track of the occurrence of a faulty transition. In the end, we aim at computing (a finite representation of) the powerset automaton \mathbf{DA} , which amounts to computing the transition functions \mathbf{O}_l for any $l \in \Sigma$ and \mathbf{O}_d for any $d \in \mathbb{R}_{\geq 0}$. Computing $\mathbf{O}_l(m)$ for $l \in \Sigma$ is not very involved: for each state $s \in S$ and each transition e labelled with l with source s and target s' , it suffices to intersect $m(s)$ with the guard $\widehat{e} \cap e$, and add the resulting interval (or the singleton $\{0\}$ if e is a resetting transition) in $\mathbf{O}_l(m)(s')$.

From now on, we only focus on computing \mathbf{O}_d , for $d \in \mathbb{R}_{\geq 0}$. For this, it is sufficient to only consider silent transitions of \mathcal{A} : we let $U = U_0 \uplus U_{id}$ be the subset of T containing all transitions labelled ϵ , partitioned into those transitions that reset clock x (in U_0), and those that do not (in U_{id}). We write \mathcal{A}_ϵ for the restriction of \mathcal{A} to silent transitions, and only consider that automaton in the sequel. All transitions are silent in \mathcal{A}_ϵ , but for convenience, we assume that transitions are labelled with their name in \mathcal{A} , so that the *untimed language* of \mathcal{A}_ϵ is the set of sequences of consecutive (silent) transitions fireable from the initial configuration.

4.1 Linear Timed Markings and Their ϵ -Closure

We use markings to represent sets of configurations; in order to compute \mathbf{O}_d , we need to represent the evolution of markings over time. For this, we introduce *timed markings*. A timed marking is a mapping $M: S \rightarrow (\mathbb{R}_{\geq 0} \rightarrow 2^{\mathbb{R}_{\geq 0}})$. For any $s \in S$ and any $d \in \mathbb{R}_{\geq 0}$, $M(s)(d)$ is intended to represent all clock valuations that can be obtained in s after a delay of d time units. For any delay $d \in \mathbb{R}$, we may (abusively) write $M(d)$ for the marking represented by M after delay d (so that for any $s \in S$ and any $d \in \mathbb{R}_{\geq 0}$, both notations $M(d)(s)$ and $M(s)(d)$ represent the same subset of \mathbb{R}).

A special case of timed marking are those timed markings that can be defined using timed sets; timed markings of this kind will be called *linear timed markings* in the sequel. As we prove below, linear timed markings are expressive enough to represent how markings evolve over time in one-clock timed automata. Atomic (resp. finite, regular) timed markings are linear timed markings whose values are atomic (resp. finite, regular) timed sets (we may omit to mention linearity in these cases to alleviate notations). Union, inclusion and equivalence of (timed) markings are defined statewise.

With any marking m , we associate a linear timed marking, which we write \vec{m} (or sometimes m if no ambiguity arises), defined as $\vec{m}(s)(d) = \{v+d \mid v \in m(s)\}$. This timed marking is linear since it can be defined e.g. as $\vec{m}(s) = (m(s); \uparrow 0)$. This timed marking can be used to represent all clock valuations that can be reached from marking m after any delay $d \in \mathbb{R}_{\geq 0}$.

Given a marking m , a delay d and a sequence $w \in U^*$ of silent transitions of \mathcal{A}_ϵ , we define the marking $m \oplus_w d$ as follows:

$$m \oplus_w d: s' \mapsto \{v \in \mathbb{R}_{\geq 0} \mid \exists s \in S. \exists v \in m(s). (s, v) \xrightarrow{d}_w (s', v')\}$$

(remember that w here represents a sequence of silent transitions). This corresponds to all configurations reachable along w from a configuration in m with a delay of exactly d time units. By definition of the transition relation \rightarrow_w , for $m \oplus_w d$ to be non-empty, w must be a sequence of consecutive transitions. With this definition, for any sequence w of silent transitions and any marking m , we can define a timed marking $m^w: d \mapsto m \oplus_w d$. In particular, for the empty sequence \perp , the timed marking m^\perp is equivalent to the timed marking \vec{m} (hence it is linear).

For any $d \in \mathbb{R}_{\geq 0}$, we define $m \oplus_\epsilon d$ as $(m \oplus_\epsilon d)(s) = \bigcup_{w \in U^*} (m \oplus_w d)(s)$. The marking $m \oplus_\epsilon d$ represents the set of configurations that can be reached after a delay d through sequences of silent transitions. This gives rise to a timed marking, which we write m^ϵ . By definition of \mathbf{O}_d , we have $\mathbf{O}_d(m) = m^\epsilon(d)$ for any marking m and any delay d .

The definition is extended to timed markings as follows: for a timed marking M and a delay d , we let

$$M \oplus_w d: s' \mapsto \{v' \in \mathbb{R}_{\geq 0} \mid \exists s \in S. \exists d_0 \leq d. \exists v \in M(d_0)(s). (s, v) \xrightarrow{d-d_0}_w (s', v')\},$$

Again, this gives rise to a timed marking $M^w : d \mapsto M \oplus_w d$. Observe that for any linear timed marking, we have $M^\perp \equiv M$. Notice also that for any marking m , it holds $(\vec{m})^w \equiv m^w$. We let $M \oplus_\epsilon d$ be the union of all $M \oplus_w d$ when w ranges over U^* , and M^ϵ be the associated timed marking.

Definition 8. Let M be a timed marking. A timed marking N is an ϵ -closure of M if $N \equiv M^\epsilon$. The timed marking M is said ϵ -closed if it is an ϵ -closure of itself.

Our aim in this section is to compute (a finite representation of) an ϵ -closure of any given initial marking (defined using regular unions of intervals).

Example 2. Consider the (silent) timed automaton of Fig. 2. The initial configuration can be represented by the timed marking M with $M(s_0) = ([0; 0]; \uparrow 0)$ and $M(s_1) = (\emptyset; \uparrow 0)$, corresponding to the single configuration $\{(s_0, x = 0)\}$. This timed marking is not closed under delay- and silent-transitions, as for instance configuration $(s_1, x = 0)$ is reachable; however, this configuration cannot be reached after any delay: it is only reachable after delay 0, or after a delay larger than or equal to 2 time units. In the end, an ϵ -closed timed marking for this automaton is $M^\epsilon(s_0) = M(s_0)$, and $M^\epsilon(s_1) = ((-\infty; -2] \cup [0; 0]; \uparrow 0)$. \triangleleft

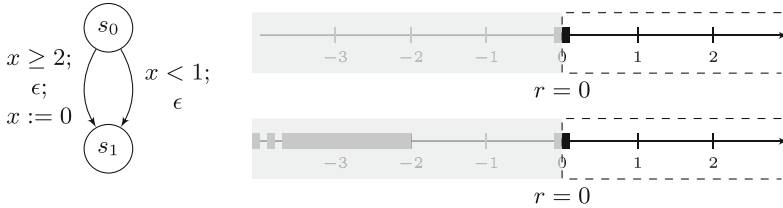


Fig. 2. A silent timed automaton

4.2 Computing ϵ -Closures

Let E and F be two subsets of \mathbb{R} . We define their gauge as the set $E \bowtie F = (E - F) \cap \mathbb{R}_{\leq 0}$. Equivalently, $E \bowtie F = \{d \in \mathbb{R}_{\leq 0} \mid F + d \cap E \neq \emptyset\}$.

- Lemma 9.**
- If $E \leq F$ (in particular, if $E \subseteq \mathbb{R}_{\leq 0}$ and $F \subseteq \mathbb{R}_{\geq 0}$), then $E \bowtie F = E - F$.
 - if $E > F$, then $E \bowtie F = \emptyset$.
 - If E and F are two intervals, then $E \bowtie F$ is an interval.
 - IF E is a regular union of intervals and J is an interval, then $E \bowtie F$ is a regular union of intervals.
 - If $F' \subseteq \mathbb{R}_{\geq 0}$, then $(E \bowtie F) \bowtie F' = (E \bowtie F) - F'$.

We now define a mapping $\epsilon : \mathcal{T}(\mathbb{R}) \times U^* \rightarrow \mathcal{T}(\mathbb{R})$, intended to represent the timed set that is reached by performing sequences of silent transitions from some given timed set. We first consider atomic timed sets, and the application of a single silent transition. The definition is based on the type of the transition:

$$\epsilon((E, \hat{r}), (s, \hat{e}, \underline{e}, \epsilon, \text{op}, s')) = \begin{cases} (\emptyset; \uparrow 0) & \text{if } \hat{r} \cap \underline{e} = \emptyset \\ (E \cap \underline{e}; \hat{p}) & \text{if } \hat{r} \cap \underline{e} \neq \emptyset \text{ and } \text{op} = id \\ (E \bowtie (\hat{p} \cap \underline{e}); \uparrow 0) & \text{if } \hat{r} \cap \underline{e} \neq \emptyset \text{ and } \text{op} = 0 \end{cases}$$

where $\hat{p} = \hat{r} \cap \hat{e}$. We extend ϵ to sequences of transitions inductively by letting $\epsilon((E; \hat{r}), \perp) = (E; \hat{r})$ and, for $w \in U^+$,

$$\epsilon((E; \hat{r}), w \cdot e) = \begin{cases} \epsilon(\epsilon((E; \hat{r}), w), e) & \text{if } \text{tgt}(w) = \text{src}(e) \\ (\emptyset; \uparrow 0) & \text{otherwise.} \end{cases}$$

Finally, we extend this definition to unions of atomic timed sets by letting $\epsilon(F_1 \sqcup F_2, w) = \epsilon(F_1, w) \sqcup \epsilon(F_2, w)$. We now prove that this indeed corresponds to applying silent- and delay-transitions from a given timed set.

Lemma 10. *Let F be a timed set and $w \in U^*$. Then for any $d \in \mathbb{R}_{\geq 0}$ and any $v \in \mathbb{R}_{\geq 0}$,*

$$v \in (\epsilon(F, w))(d) \Leftrightarrow \exists d_0 \in [0; d]. \exists v' \in F(d_0). (\text{src}(w), v') \xrightarrow{d-d_0}_w (\text{tgt}(w), v).$$

Proof. We carry the proof for the case where F is an atomic timed set. The extension to unions of atomic timed sets is straightforward. The proof for atomic timed sets is in two parts: we begin with proving the result for a single transition (the case where $w = \perp$ is easy), and then proceed by induction to prove the full result.

We begin with the case where w is a single transition $e = (s, \hat{e}, \underline{e}, \epsilon, \text{op}, s')$. In case F is empty, then also $\epsilon(F, e)$ is empty, and the result holds. We now assume that F is not empty, and consider three cases:

- if $\hat{r} \cap \underline{e} = \emptyset$, then $\epsilon(F, e) = (\emptyset; \uparrow 0)$. On the other hand, for any d_0 and any $v' \in F(d_0)$, it holds $v' \in \hat{r}$, so that $v' \notin \underline{e}$, and the transition cannot be taken from that valuation. Hence both sides of the equivalence evaluate to false, and the equivalence holds.
- now assume that $\hat{r} \cap \underline{e} \neq \emptyset$, and consider the case where e does not reset the clock. Then $v \in (\epsilon(F, w))(d)$ means that $v \in (E + d) \cap (\underline{e} + d) \cap \hat{p}$. If such a v exists, then $\hat{p} \cap \underline{e} \cap [v - d; v]$ is non-empty: indeed, this is trivial if either $v \in \underline{e}$, or $v - d \in \hat{p}$, or $d = 0$; otherwise, we have $\underline{e} \subseteq \downarrow v$ and $\hat{p} \subseteq \uparrow v - d$, so that $\hat{p} \cap \underline{e} \subseteq (v - d; v)$. Moreover, $\hat{p} \cap \underline{e} \neq \emptyset$ since $\hat{r} \cap \underline{e} \neq \emptyset$ and $\hat{e} \cap \underline{e} \neq \emptyset$. Then for any v' in that set $\hat{p} \cap \underline{e} \cap [v - d; v]$, letting $d_0 = v' - (v - d)$, we have $v' \in E + d_0$. In the end, $v' \in F(d_0)$, and $v' \in \hat{e} \cap \underline{e}$, so that $(\text{src}(e), v') \xrightarrow{d-d_0}_e (\text{tgt}(e), v)$. Conversely, if $d_0 \in [0; d]$ and $v' \in F(d_0)$ exist such that $(\text{src}(w), v') \xrightarrow{d-d_0}_w (\text{tgt}(w), v)$, then $v' \in E + d_0 \cap \hat{r}$, and for some $d_1 \leq d - d_0$, $v' + d_1 \in \hat{e} \cap \underline{e}$.

- Then letting $v = v' + d - d_0$, we have $v \in E + d$ and $v \in \hat{r} + (d - d_0) \subseteq \hat{r}$ and $v \in \hat{e} + (d - d_0 - d_1) \subseteq \hat{e}$ and $v \in \underline{e} + (d - d_0 - d_1) \subseteq \underline{e} + d$. This proves our result for this case.
- we finally consider the case where $\hat{r} \cap \underline{e} \neq \emptyset$ and e resets the clock. In this case, $v \in (\epsilon(F, w))(d)$ means that $v \in \uparrow 0$ and $v - d \in E \bowtie (\hat{p} \cap \underline{e})$, which rewrites as $0 \leq v \leq d$ and $(E + d - v) \cap (\hat{p} \cap \underline{e}) \neq \emptyset$. Let $d_0 = d - v$. The property above entails that $0 \leq d_0 \leq d$, and that there exists some $v' \in (E + d_0) \cap (\hat{p} \cap \underline{e})$, so that $0 \leq d_0 \leq d$, $v' \in F(d_0)$ and $(\text{src}(e), v') \xrightarrow{d-d_0}_e (\text{tgt}(e), v)$. Conversely, if those conditions hold, then for some $0 \leq d_1 \leq d - d_0$, we have $v' + d_1 \in \hat{e} \cap \underline{e}$, and $v = d - (d_0 + d_1)$ (remember that e resets the clock). Then $v' + d_1 \in E + (d_0 + d_1) \cap \hat{r} \cap \hat{e} \cap \underline{e}$, so that $-(d_0 + d_1) \in E \bowtie (\hat{p} \cap \underline{e})$, and finally $v \in \epsilon(F, w)(d)$.

We now extend this result to sequences of transitions. The case where $w = \perp$ is straightforward. Now assume that the result holds for some word w , and consider a word $w \cdot e$. In case $\text{tgt}(w) \neq \text{src}(e)$, the result is trivial.

The case of single transitions has been handled just above. We thus consider the case of $w \cdot e$ with $w \in U^+$. First assume that $v' \in (\epsilon(F, w \cdot e))(d)$, and let $F' = \epsilon(F, w)$. Then $v' \in (\epsilon(F', e))(d)$, thus there exist $0 \leq d_0 \leq d$ and $v \in F'(d_0)$ s.t. $(\text{src}(e), v) \xrightarrow{d-d_0}_e (\text{tgt}(e), v')$. Since $v \in F'(d_0)$, there must exist $0 \leq d_1 \leq d_0$ and $v'' \in F(d_1)$ such that $(\text{src}(w), v'') \xrightarrow{d_0-d_1}_w (\text{tgt}(w), v)$. We thus have found $0 \leq d_1 \leq d$ such that $(\text{src}(w), v'') \xrightarrow{d-d_1}_{w \cdot e} (\text{tgt}(e), v')$.

Conversely, if $(\text{src}(w), v'') \xrightarrow{d-d_1}_{w \cdot e} (\text{tgt}(e), v')$ for some $0 \leq d_1 \leq d$ and $v'' \in F(d_1)$, then we have $(\text{src}(w), v'') \xrightarrow{d_0-d_1}_w (\text{tgt}(w), v) \xrightarrow{d-d_0}_e (\text{tgt}(e), v')$ for some $d_0 \in [d_1; d]$ and some v . We prove that $v \in (\epsilon(F, w))(d_0)$: indeed, we have $d_1 \in [0; d_0]$, and $v'' \in F(d_1)$ such that $(\text{src}(w), v'') \xrightarrow{d_0-d_1}_w (\text{tgt}(w), v)$, which by induction hypothesis entails $v \in (\epsilon(F, w))(d_0)$. Thus we have $d_0 \in [0; d]$ and $v \in F'(d_0)$, where $F' = \epsilon(F, w)$, such that $(\text{tgt}(w), v) \xrightarrow{d-d_0}_e (\text{tgt}(e), v')$, which means $v' \in \epsilon(F', e)(d)$, and concludes the proof. \square

Thanks to this semantic characterization of $\epsilon(F, w)$, we get:

Corollary 11. *For any sequence w of transitions, and any two equivalent timed sets F and F' , the timed sets $\epsilon(F, w)$ and $\epsilon(F', w)$ are equivalent.*

Finally, we extend ϵ to linear timed markings in the expected way: given a linear timed marking M , and a sequence w of transitions, we let

$$\begin{aligned} \epsilon(M, w) : s \mapsto \epsilon(M(\text{src}(w)), w) & \text{ if } s = \text{tgt}(w), \\ s \mapsto (\emptyset; \uparrow 0) & \text{ otherwise.} \end{aligned}$$

Again, we have $\epsilon(M_1 \sqcup M_2, w) \equiv \epsilon(M_1, w) \sqcup \epsilon(M_2, w)$ for any $w \in U^*$. Then:

Lemma 12. *$\epsilon(M, w) \equiv M^w$ for all $w \in U^*$ and all linear timed marking M .*

Letting $\epsilon(M, \mathcal{L}) = \bigsqcup_{w \in \mathcal{L}} \epsilon(M, w)$ for any subset \mathcal{L} of U^* , and $\epsilon(M) = \epsilon(M, U^*)$, we immediately get:

Theorem 13. *For any linear timed marking M , it holds $\epsilon(M) \equiv M^\epsilon$.*

It follows that the closure of any marking can be represented as a linear timed marking. However, this linear timed marking is currently defined as an infinite union over all sequences of consecutive silent transitions. We make the computation more effective in the next section.

4.3 Finite Representation of the Closure

In this section, we prove that we can effectively compute a finite representation of the closure of any regular timed marking. More precisely, we show how to compute such a closure as a regular timed marking.

Finiteness. Indeed, let M be (a finite representation of) a regular timed marking: then M can be written as the finite union of atomic regular timed markings $M_{(s,E,\hat{r})}$, defined as $M_{(s,E,\hat{r})}(s) = (E; \hat{r})$ and $M_{(s,E,\hat{r})}(s') = (\emptyset; \uparrow 0)$ for all $s' \neq s$. In the end, any regular timed marking M can be written as the finite union $\bigsqcup_{i \in I} M_i$ of atomic regular timed markings. Thus we have $\epsilon(M) \equiv \bigsqcup_{i \in I} \epsilon(M_i)$, and it suffices to compute ϵ for atomic regular timed markings $M_{(s,E,\hat{r})}$. We prove that those closures can be represented as regular timed markings.

We write $\epsilon_1((E; \hat{r}), w)$ and $\epsilon_2((E; \hat{r}), w)$ for the first and second components of $\epsilon((E; \hat{r}), w)$. Notice that $\epsilon_2((E; \hat{r}), w)$ does not depend on E (so that we may denote it with $\epsilon_2(\hat{r}, w)$ in the sequel). In particular,

- $\epsilon_2((E; \hat{r}), w) = \uparrow 0$ if $w \in U^* \times U_0$ is a sequence of consecutive transitions ending with a resetting transition;
- $\epsilon_2((E; \hat{r}), w) = \hat{r} \cap \bigcap_{i < k} \hat{e}_i$ if $w = e_1 \dots e_k \in U_{id}^*$ is a sequence of consecutive non-resetting transitions.

Letting $\mathbb{J}_{\hat{r}} = \{\uparrow 0, \hat{r}\} \cup \{\hat{e} \mid e \in U_{id}\}$, it follows that $\epsilon_2((E; \hat{r}), w) \in \mathbb{J}_{\hat{r}}$ for any $(E; \hat{r})$ and any w . Thus $\epsilon(M)$ can be written as a finite union of atomic timed markings.

Regularity. To prove regularity, we first introduce some more formalism:

- we let $\widehat{G}_{id} = \{\hat{e} \mid e \in U_{id}\}$ and $\underline{G}_{id} = \{e \mid e \in U_{id}\}$. We thus have $\mathbb{J}_{\hat{r}} = \{\uparrow 0, \hat{r}\} \cup \widehat{G}_{id}$;
- for $\hat{r} \in \mathbb{R}_{\geq 0}$ and $e \in U$, we write $\Phi(\hat{r}, e)$ for the interval $\hat{r} \cap \hat{e} \cap e$;
- we define a mapping $\mathcal{J}_{\hat{r}}: U^* \rightarrow \mathbb{N}^{\mathbb{J}_{\hat{r}} \times U_0}$ that counts the number of occurrences of certain timing constraints at resetting transitions along a path: precisely, it is defined inductively as follows (where \uplus represents addition of an element to a multiset):

$$\begin{aligned} \mathcal{J}_{\hat{r}}(\perp) &= \{0\}^{\mathbb{J}_{\hat{r}} \times U_0} \\ \mathcal{J}_{\hat{r}}(w \cdot e) &= \mathcal{J}_{\hat{r}}(w) \uplus \{(\epsilon_2(\hat{r}, w), e)\} && \text{if } e \in U_0 \\ \mathcal{J}_{\hat{r}}(w \cdot e) &= \mathcal{J}_{\hat{r}}(w) && \text{if } e \in U_{id}. \end{aligned}$$

By induction on w , we prove:

Lemma 14. *Let $(E; \hat{r})$ be a timed set with $E \subseteq \mathbb{R}_{\leq 0}$, and $w \in U^*$. Then*

$$\epsilon_1((E; \hat{r}), w) = E - \sum_{J=(\hat{g}, e) \in \mathbb{J}_{\hat{r}}} \mathcal{J}_{\hat{r}}(w)(J) \times \Phi(\hat{g}, e) \subseteq \mathbb{R}_{\leq 0}.$$

Now, we fix an atomic regular timed marking $M_{(s, E, \hat{r})}$. For any state s' of \mathcal{A}_ϵ , we let $\mathcal{L}(s, s')$ be the set of all sequences of consecutive transitions from s to s' in \mathcal{A}_ϵ . Then

$$(M_{(s, E, \hat{r})})^\epsilon = \bigsqcup_{s' \in S} \epsilon(M_{(s, E, \hat{r})}, \mathcal{L}(s, s')).$$

Hence we need to prove that $\epsilon_1(M_{(s, E, \hat{r})}, \mathcal{L}(s, s'))$ is regular.

For any set \mathcal{L} of sequences of consecutive transitions, and for any \hat{r} and \hat{r}' in $\widehat{\mathbb{R}}_{\geq 0}$, we let $\mathcal{L}_{\hat{r}}^{\hat{r}'} = \{w \in \mathcal{L} \mid \hat{r}' = \epsilon_2(\hat{r}, w)\}$. One easily observes that for any $\hat{r} \in \widehat{\mathbb{R}}_{\geq 0}$ and any \mathcal{L} , it holds $\mathcal{L} = \bigcup_{\hat{r}' \in \mathbb{J}_{\hat{r}}} \mathcal{L}_{\hat{r}}^{\hat{r}'}$, so that

$$\epsilon_1(M_{(s, E, \hat{r})}, \mathcal{L}(s, s')) = \bigcup_{\hat{r}' \in \mathbb{J}_{\hat{r}}} \epsilon_1(M_{(s, E, \hat{r})}, [\mathcal{L}(s, s')]_{\hat{r}}^{\hat{r}'}).$$

The following property entails that this set is a regular union of intervals:

Lemma 15. *Let E be a regular union of intervals, \hat{r} and \hat{r}' be two elements of $\widehat{\mathbb{R}}_{\geq 0}$, and $\mathcal{L} \subseteq \mathcal{L}(\mathcal{A}_\epsilon)$ be a regular language. Then $\epsilon_1(M_{(s, E, \hat{r})}, \mathcal{L}_{\hat{r}}^{\hat{r}'})$ is a regular union of intervals.*

5 Experimentations

In order to evaluate the possible improvement of our approach compared to the diagnoser proposed in [26], we implemented and compared the performances of both approaches. Sources can be downloaded at <http://www.lsv.fr/~jaziri/DOTA.zip>.

5.1 Comparison of the Approaches

In the approach of [26], the set of possible current configurations is stored as a marking. If an action l occurs after some delay d , the diagnoser computes the set of all possible configurations reached after delay d (possibly following silent transitions), and applies from the resulting markings the set of all available transitions labelled l . This amounts to computing the functions \mathbf{O}_d and \mathbf{O}_l at each observation. There is also a timeout, which makes the diagnoser update the marking (with \mathbf{O}_d) regularly if no action is observed. The computation of \mathbf{O}_d is heavily used, and has to be performed very efficiently so that the diagnoser can be used at runtime.

In our approach, we use *timed markings* to store sets of possible configurations. Given a timed marking, when an action l is observed after some delay d , we can easily compute the set of configurations reachable after delay d , and have to apply \mathbf{O}_l and recompute the ϵ -closure. Following [7], \mathbf{O}_l can be performed as a series of set operations on intervals. The ϵ -closure can be performed as a series of subtractions between an interval and regular unions of intervals (see Lemma 14). Those regular unions can be precomputed; while this may require exponential time and space to compute and store, this makes the simulation of a delay transition very efficient.

5.2 Implementation

In our experimentations, in order to only evaluate the benefits of the precomputation and of the use of ϵ -closures in our approach compared to that of [26], we use the same data structure for both diagnosers. In particular, both diagnosers are implemented as automata over timed domains [7], where the timed domain is the set of timed markings. The only difference lies in the functions computing the action- and the delay transitions. As a consequence, both implementations benefit from the data structure we chose for representing timed intervals, which allows us to compute basic operations in linear time. Also, both structures use the same reachability graphs for either computing the sets of reachable configurations or the Parikh images.

Our implementation is written in Python3. One-clock timed automata and both diagnosers are instances of an abstract class of automata over timed markings; timed markings are implemented in a library `TILib`. Simulations of those automata are performed using an object called `ATDRunner`, which takes an automaton over timed markings and simulates its transitions according to the actions it observes on a given *input channel*. It may also write what it does on an *output channel*. A channel is basically a way of communicating with `ATDRunners`.

In order to diagnose a given one-clock timed automaton, stored in an object `OTAutomata`, we first generate a diagnoser object, either a `DiagOTA` or a `TripakisDOTA`, depending of which version we want to use. Then we launch two threads: one is an `ATDRunner` simulating the timed automaton, listening to some channel object `Input`, and writing every non-silent action it performs on some other channel object `Comm`. The other one is another `ATDRunner` simulating the diagnoser and listening to the `Comm` channel.

In a `DiagOTA` object, which corresponds to our approach, we have already precomputed the relevant timed intervals; action transitions are then made by operations over timed markings, and delay transitions are encoded by increasing a *padding information* on the timed markings, which is applied when performing the next action transition. In such a simulation, we can thus keep track of which states may have been reached, but also predict which states may be reached in the future and the exact time before we can reach them.

In a `TripakisDOTA` object, which corresponds to the approach of [26], action and delay transitions are simulated by computing all configurations reachable

through that action or delay, also allowing arbitrarily many silent transitions. This does not allow for prediction.

5.3 Results

Figure 3 reports on the performances of both implementations on a small set of (randomly generated) examples. Those examples are distributed with our prototype. In Fig. 3, we give the important characteristics of each automaton (number of states and of silent transitions), the amount of precomputation time used by our approach, and the average time (over 400 random runs) used in the two approaches to simulate action- and delay transitions.

	Example2	Example3	Example4	Example6	Example8
#State/#Silent Trans	3/6	4/6	4/7	7/10	7/5
Precomputation Time	173.25s	0.38s	791.06s	11.01s	4.96s
Actions DiagOTA	0.014s	0.019s	0.029s	0.17s	0.15s
Actions TripakisDOTA	0.020s	0.078s	0.049s	0.26s	0.042
Ratio (actions)	0.73	0.25	0.59	0.64	3.71
Delays DiagOTA	0.000012s	0.0000011s	0.000011s	0.000011s	0.000012s
Delays TripakisDOTA	0.032s	0.057s	0.049s	0.30s	0.033s
Ratio (delays)	0.0004	0.0002	0.0002	0.00003	0.0004

Fig. 3. Bench for 5 examples over 400 runs with 10 to 20 actions

As could be expected, our approach outperforms the approach of [26] on delay transitions by several orders of magnitude in all cases. The performances of both approaches are comparable when simulating action transitions.

The precomputation phase of our approach is intrinsically very expensive. In our examples, it takes from less than a second to more than 13 min, and it remains to be understood which factors make this precomputation phase more or less difficult. We may also refine our implementation of the computation of Parikh images, which is heavily used in the precomputation phase.

6 Conclusion and Future Works

In this paper, we presented a novel approach to fault diagnosis for one-clock timed automata; it builds on a kind of powerset construction for automata over timed domains, using our new formalism of *timed sets* to represent the evolution of the set of reachable configurations of the automaton. Our prototype implementation shows the feasibility of our approach on small examples.

There remains space for improvements in many directions: first, our implementation can probably be made more efficient on the precomputation phase,

and at least we need to better understand why some very small examples are so hard to handle.

A natural continuation of this work is an extension to n -clock timed automata. This is not immediate, as it requires a kind of *timed zone*, and an adaptation of our operator \bowtie . Another possible direction of research could target priced timed automata, with the aim of monitoring the cost of the execution in the worst case.

References

1. Alur, R., Dill, D.L.: A theory of timed automata. *Theor. Comput. Sci.* **126**(2), 183–235 (1994)
2. Behrmann, G., et al.: Uppaal 4.0. In: Proceedings of the 3rd International Conference on Quantitative Evaluation of Systems (QEST 2006), pp. 125–126. IEEE Computer Society Press, September 2006
3. Bengtsson, J., Yi, W.: Timed automata: semantics, algorithms and tools. In: Desel, J., Reisig, W., Rozenberg, G. (eds.) ACPN 2003. LNCS, vol. 3098, pp. 87–124. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-27755-2_3
4. Benveniste, A., Fabre, É., Haar, S., Jard, C.: Diagnosis of asynchronous discrete event systems: a net-unfolding approach. *IEEE Trans. Autom. Control.* **48**(5), 714–727 (2003)
5. Bertrand, N., Haddad, S., Lefauchaux, E.: Foundation of diagnosis and predictability in probabilistic systems. In: Raman, V., Suresh, S.P. (eds.) Proceedings of the 34th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2014). Leibniz International Proceedings in Informatics, vol. 29, pp. 417–429. Leibniz-Zentrum für Informatik, December 2014
6. Bouyer, P., Chevalier, F., D’Souza, D.: Fault diagnosis using timed automata. In: Sassone, V. (ed.) FoSSaCS 2005. LNCS, vol. 3441, pp. 219–233. Springer, Heidelberg (2005). https://doi.org/10.1007/978-3-540-31982-5_14
7. Bouyer, P., Jaziri, S., Markey, N.: On the determinization of timed systems. In: Abate, A., Geeraerts, G. (eds.) FORMATS 2017. LNCS, vol. 10419, pp. 25–41. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-65765-3_2
8. Cassez, F.: A note on fault diagnosis algorithms. In: Proceedings of the 48th IEEE Conference on Decision and Control (CDC 2009), pp. 6941–6946. IEEE Computer Society Press, December 2009
9. Chen, Y.-L., Provan, G.: Modeling and diagnosis of timed discrete event systems - a factory automation example. In: Proceedings of the 1997 American Control Conference (ACC 1997), pp. 31–36. IEEE Computer Society Press, June 1997
10. Clarke, E.M., Allen Emerson, E., Sifakis, J.: Model checking: algorithmic verification and debugging. *Commun. ACM* **52**(11), 74–84 (2009)
11. Clarke, E.M., Grumberg, O., Peled, D.A.: Model Checking. MIT Press, Cambridge (2000)
12. Clarke, E.M., Henzinger, T.A., Veith, H., Bloem, R.: Handbook of Model Checking. Springer, Cham (2018). <https://doi.org/10.1007/978-3-319-10575-8>
13. Filliâtre, J.-C.: Deductive software verification. *Int. J. Softw. Tools Technol. Transf.* **13**(5), 397–403 (2011)
14. Finkel, O.: Undecidable problems about timed automata. In: Asarin, E., Bouyer, P. (eds.) FORMATS 2006. LNCS, vol. 4202, pp. 187–199. Springer, Heidelberg (2006). https://doi.org/10.1007/11867340_14

15. Genc, S., Lafortune, S.: Predictability of event occurrences in partially-observed discrete-event systems. *Automatica* **45**(2), 301–311 (2009)
16. Henzinger, T.A., Sifakis, J.: The embedded systems design challenge. In: Misra, J., Nipkow, T., Sekerinski, E. (eds.) *FM 2006*. LNCS, vol. 4085, pp. 1–15. Springer, Heidelberg (2006). https://doi.org/10.1007/11813040_1
17. Hoare, C.A.R.: An axiomatic basis for computer programming. *Commun. ACM* **12**(10), 576–580 (1969)
18. Lafortune, S., Lin, F., Hadjicostis, C.N.: On the history of diagnosability and opacity in discrete event systems. *Annu. Rev. Control.* **45**, 257–266 (2018)
19. Leucker, M., Schallart, C.: A brief account of runtime verification. *J. Log. Algebr. Program.* **78**(5), 293–303 (2009)
20. Lunze, J., Schröder, J.: State observation and diagnosis of discrete-event systems described by stochastic automata. *Discret. Event Dyn. Syst.* **11**(4), 319–369 (2001)
21. Narasimhan, S., Biswas, G.: Model-based diagnosis of hybrid systems. *IEEE Trans. Syst., Man, Cybern. Part A Syst. Hum.* **37**(3), 348–361 (2007)
22. Sampath, M., Lafortune, S., Teneketzis, D.: Active diagnosis of discrete-event systems. *IEEE Trans. Autom. Control.* **43**(7), 908–929 (1998)
23. Sampath, M., Sengupta, R., Lafortune, S., Sinnamohideen, K., Teneketzis, D.: Diagnosability of discrete-event systems. *IEEE Trans. Autom. Control.* **40**(9), 1555–1575 (1995)
24. Sampath, M., Sengupta, R., Lafortune, S., Sinnamohideen, K., Teneketzis, D.: Failure diagnosis using discrete-event models. *IEEE Trans. Comput.* **35**(1), 105–124 (1996)
25. Tretmans, J.: Conformance testing with labelled transition systems: Implementation relations and test generation. *Comput. Netw. ISDN Syst.* **29**(1), 49–79 (1996)
26. Tripakis, S.: Description and schedulability analysis of the software architecture of an automated vehicle control system. In: Sangiovanni-Vincentelli, A., Sifakis, J. (eds.) *EMSOFT 2002*. LNCS, vol. 2491, pp. 123–137. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-45828-X_10
27. Tripakis, S.: Folk theorems on the determinization and minimization of timed automata. *Inf. Process. Lett.* **99**(6), 222–226 (2006)
28. Zaytoon, J., Lafortune, S.: Overview of fault diagnosis methods for discrete event systems. *Annu. Rev. Control.* **37**(2), 308–320 (2013)