Quantitative Robustness Analysis of Flat Timed Automata

Remi Jaubert, Pierre-Alain Reynier

Quantitative Robustness Analysis of Flat Timed Automata

Rémi Jaubert*, Pierre-Alain Reynier

LIF, Université Aix-Marseille & CNRS, France
{remi.jaubert,pierre-alain.reynier}@lif.univ-mrs.fr

Abstract. Whereas formal verification of timed systems has become a very active field of research, the idealized mathematical semantics of timed automata cannot be faithfully implemented. Recently, several works have studied a parametric semantics of timed automata related to implementability: if the specification is met for some positive value of the parameter, then there exists a correct implementation. In addition, the value of the parameter gives lower bounds on sufficient resources for the implementation. In this work, we present a symbolic algorithm for the computation of the parametric reachability set under this semantics for flat timed automata. As a consequence, we can compute the largest value of the parameter for a timed automaton to be safe.

1 Introduction

Verification of real-time systems. In the last thirty years, formal verification of reactive, critical, or embedded systems has become a very active field of research in computer science. It aims at checking that (the model of) a system satisfies (a formula expressing) its specifications. The importance of taking real-time constraints into account in verification has quickly been understood, and the model of timed automata [AD94] has become one of the most established models for real-time systems, with a well studied underlying theory, the development of mature model-checking tools (UPPAAL [BDL04], KRONOS [BDM+98], ...), and numerous success stories.

Implementation of real-time systems. Implementing mathematical models on physical machines is an important step for applying theoretical results on practical examples. This step is well-understood for many untimed models that have been studied (*e.g.*, finite automata, pushdown automata). In the timed setting, while timed automata are widely-accepted as a framework for modelling real-time aspects of systems, it is known that they cannot be faithfully implemented on finite-speed CPUs [CHR02]. Studying the "implementability" of timed automata is thus a challenging issue of obvious theoretical and practical interest.

A semantical approach. Timed automata are governed by a mathematical, idealized semantics, which does not fit with the digital, imprecise nature of the hardware on which they will possibly be implemented. An *implementation semantics* has been defined in [DDR05] in order to take the hardware into account: that semantics models a

* Funded by a doctoral grant of "Conseil Régional Provence-Alpes-Côte d'Azur".

digital CPU which, every δ_P time units (at most), reads the value of the digital clock (updated every δ_L time units), computes the values of the guards, and fires one of the available transitions. A timed automaton is then said to be *implementable* if there exist positive values for those parameters (δ_P and δ_L) for which, under this new semantics, the behaviours of the automaton satisfy its specification. In order to study it efficiently, this semantics is over-approximated by the *AASAP semantics*, which consists in "enlarging" the constraints on the clocks by some parameter δ. For instance, "$x \in [a, b]$" is transformed into "$x \in [a - \delta, b + \delta]$". Moreover, a formal link is drawn in [DDR05] between these two semantics: as soon as $\delta > 4\delta_P + 3\delta_L$, the AASAP semantics simulates the semantics of the implementation. As a consequence, implementability can be ensured by establishing the existence of some positive δ for which the AASAP semantics meets the specification.

Robustness problems. We call the above problem (existence of some positive δ) the *qualitative problem of robustness*. This problem was proven decidable for different kind of properties: the problem is PSPACE-complete for safety properties [Pur00,DDMR08] and LTL formula [BMR06]. It is EXPTIME-complete for a fragment of the timed logic MTL [BMR08]. In addition, for safety properties, it is proven in [Pur00,DDMR08] that if there exists a safe positive value of δ, then the system is also safe for a specific value of the form $1/2^{|\mathcal{A}|}$. While this allows to deduce a correct value for the parameter δ, computing the largest value of δ for which the AASAP semantics meets the specification was still an open problem. We are interested here in this last problem, which we call the *quantitative problem of robustness* for safety properties.

Our contributions. In this paper, we prove that the quantitative robustness problem for safety properties is decidable for flat timed automata (*i.e.* where each location belongs to at most one cycle). In addition, we show that the maximal safe value of δ is a rational number. To this end, we solve a more general problem: we prove that it is possible to compute the parametric reachability set for flat timed automata, and present a forward algorithm based on parametric zones (constraints on clocks). As a parametric forward analysis does not terminate for (flat) timed automata, we need some acceleration techniques. To solve the qualitative robustness problem, different algorithms have been proposed in [Pur00,DDMR08,DK06] which compute an enlarged reachability set corresponding to states reachable for any positive perturbation, and include an acceleration of cycles. The algorithm we propose can be understood as a parametric version of the symbolic algorithm proposed in [DK06] for flat timed automata. We then tackle two issues: the termination of our procedure and its correction. For the first aspect, as we are in a parametric setting, we need completely new arguments of termination (the number of parametric zones we compute cannot be bounded as it is the case for zones). Considering a graph representation of zones introduced in [CJ99a], we obtain proofs of termination depending only on the number of clocks, and not on the constants appearing in the automaton. Up to our knowledge, this constitutes an original approach in the context of timed automata. Regarding correctness, we identify under which conditions the enlarged reachability set coincides with the standard reachability set, and propose a modification of the algorithm to obtain the computation of the parametric reachability set (and not of the parametric enlarged reachability set).

Related work. Since its definition in [Pur00,DDR05], the approach based on the AASAP semantics has received much attention, and other kind of perturbations, like the drift of clocks, have been studied [DDMR08,ALM05,Dim07]. In the case of safety properties and under some natural assumptions, this perturbation is equivalent to constraint enlargement and relies on similar techniques, as proved in [DDMR08]. Also, several works have considered variants of the robustness problem. In [SF07,SFK08], the case of systems with bounded life-time or regular resynchronization of clocks is considered, while in [Dim07], a symbolic algorithm is proposed to handle strict constraints.

Many other notions of "robustness" have been proposed in the literature in order to relax the mathematical idealization of the semantics of timed automata [GHJ97,OW03,BBB+07]. Those approaches are different from ours, since they roughly consist in dropping "isolated" or "unlikely" executions, and are thus more related to language theoretical issues than to implementability issues.

Finally, our work is somewhat related to parametric timed automata. It is proven in [WT99] that emptiness is already undecidable for timed automata with three clocks and one parameter. In our setting, decidability results follow from strong restrictions on the use of the parameter. They correspond to the notion of upper parameter introduced in [HRSV02], but the problems we consider are different. In addition, to obtain termination, we introduce acceleration techniques based on [CJ99a]. Two recent works [BIL06,BIK10] also rely on [CJ99a] to propose acceleration techniques, but these concern parametric flat counter automata, and their parameter takes its values in natural numbers.

Organisation of the paper. In Section 2, we introduce standard definitions. We present in Section 3 the definition of the enlarged reachability set, and a modification of the algorithm of [DK06] for its computation. In Section 4, we first recall the graph representation of constraints, then present how we use it to obtain a new acceleration technique, and finally we present our parametric algorithm and its proof of termination and of correction.

2 Definitions

2.1 Timed Automata, Zones

Let $\mathcal{X} = \{x_1, \ldots, x_n\}$ be a finite set of *clock* variables. We extend it with a fictive clock x_0, whose value will always be 0, and denote $\overline{\mathcal{X}}$ the set $\mathcal{X} \cup \{x_0\}$. An *atomic (clock) constraint* on $\mathcal{X}$ is of the form $x - y \leqslant k$, where $x \neq y \in \overline{\mathcal{X}}$ and $k \in \mathbb{Q}$. Note that we only consider non-strict inequalities. This makes sense as we will later enlarge these constraints. We say that the constraint is *non-diagonal* if the comparison involves the clock x_0. We denote by $\mathcal{G}(\mathcal{X})$ (resp. $\mathcal{G}_{nd}(\mathcal{X})$) the set of *(clock) constraints* (resp. *non-diagonal constraints*) defined as conjunctions of atomic constraints (resp. non-diagonal atomic constraints).

A *(clock) valuation* v for $\mathcal{X}$ is an element of $\mathbb{R}^{\mathcal{X}}_{\geqslant 0}$. A valuation $v \in \mathbb{R}^{\mathcal{X}}_{\geqslant 0}$ is extended to $\mathbb{R}^{\overline{\mathcal{X}}}_{\geqslant 0}$ by $v(x_0) = 0$. If $v \in \mathbb{R}^{\mathcal{X}}_{\geqslant 0}$ and $t \in \mathbb{R}_{\geqslant 0}$, we write $v+t$ for the valuation assigning $v(x) + t$ to every clock $x \in \mathcal{X}$. If $r \subseteq \mathcal{X}$, $v[r \leftarrow 0]$ denotes the valuation assigning 0 to every clock in r and $v(x)$ to every clock in $\mathcal{X} \backslash r$. Whether a valuation $v \in \mathbb{R}^{\mathcal{X}}_{\geqslant 0}$ satisfies

a constraint $g \in \mathcal{G}(\mathcal{X})$, written $v \models g$, is defined inductively as follows: the conjunction is handled naturally, and $v \models x - y \leqslant k$ iff $v(x) - v(y) \leqslant k$ (recall that $v(x_0) = 0$). The set of valuations satisfying a constraint g is denoted $[\![g]\!]$.

A *zone* Z over $\mathcal{X}$ is a convex subset of $\mathbb{R}^{\mathcal{X}}_{\geqslant 0}$ which can be defined as the set of valuations satisfying a clock constraint, *i.e.* there exists $g \in \mathcal{G}(\mathcal{X})$ such that $Z = [\![g]\!]$. We note $\mathsf{Zones}(\mathcal{X})$ the set of zones on $\mathcal{X}$. The zone $\mathbb{R}^{\mathcal{X}}_{\geqslant 0}$ is denoted $\top$.

Definition 1 (Timed Automaton). *A* TA *is a tuple $\mathcal{A} = (L, \ell_0, \mathcal{X}, \Sigma, T)$ where L is a finite set of locations, $\ell_0 \in L$ is an initial location, $\mathcal{X}$ is a finite set of clocks, Σ is a finite set of actions, and $T \subseteq L \times \mathcal{G}_{nd}(\mathcal{X}) \times \Sigma \times 2^{\mathcal{X}} \times L$ is a finite set of transitions.*

We define the semantics of $\mathcal{A}$ as a timed transition system $[\![\mathcal{A}]\!] = \langle S, S_0, \Sigma, \rightarrow \rangle$. The set S of states of $[\![\mathcal{A}]\!]$ is $L \times \mathbb{R}^{\mathcal{X}}_{\geqslant 0}$ and $S_0 = \{(\ell_0, v_0) \mid v_0(x) = v_0(y),\ \forall x, y \in \mathcal{X}\}$. A transition in $[\![\mathcal{A}]\!]$ is composed either of a delay move $(\ell, v) \xrightarrow{d} (\ell, v + d)$, with $d \in \mathbb{R}_{\geqslant 0}$, or of a discrete move $(\ell, v) \xrightarrow{\sigma} (\ell', v')$ when there exists a transition $(\ell, g, \sigma, r, \ell') \in T$ with $v \models g$, and $v' = v[r \leftarrow 0]$. The graph $[\![\mathcal{A}]\!]$ is thus an infinite transition system. A *run* of $[\![\mathcal{A}]\!]$ is a finite or infinite sequence $(\ell_0, v_0) \xrightarrow{\sigma_1} (\ell_1, v_1) \xrightarrow{d_1} (\ell_1, v_1 + d_1) \xrightarrow{\sigma_2} (\ell_2, v_2) \ldots$ where for each $i \geqslant 1$, $d_i \in \mathbb{R}_{\geqslant 0}$, and $(\ell_0, v_0) \in S_0$. A state (ℓ, v) is *reachable* in $[\![\mathcal{A}]\!]$ iff there exists a run from an initial state $(\ell_0, v_0) \in S_0$ to (ℓ, v); the set of reachable states is denoted $\mathsf{Reach}(\mathcal{A})$.

Note that standard definitions of timed automata also allow invariants on locations which restrict time elapsing. For the sake of simplicity, we do not consider this technical addition here, however all our results hold in presence of invariants.

A *cycle* of $\mathcal{A}$ is a finite sequence of transitions corresponding to a cycle of the underlying finite state automaton. We say that a timed automaton is *flat* if each location belongs to at most one cycle. A *progress cycle* is a cycle where each clock is reset at least once. We say $\mathcal{A}$ is *progressive* if it only contains progress cycles.

Assumptions. As our results rely on previous works on robustness in TA [Pur00,DDMR08], we assume that our TA are progressive, and that all the clocks are always bounded by some constant M. In addition, as the algorithm we propose in based on [DK06], we also require our timed automata to be flat.

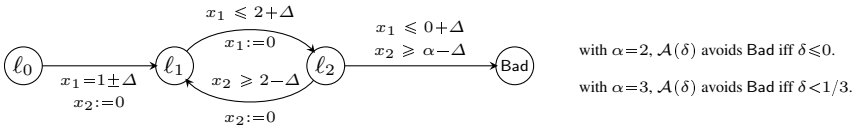

Fig. 1: A timed automaton $\mathcal{A}$, with its parametric semantics.

2.2 Parametric objects

We define the parametric semantics introduced in [Pur00] that enlarges the set of runs of timed automata. This semantics can be defined in terms of timed automata extended with one parameter, denoted Δ, with syntactic constraints on the use of this parameter.

We denote by $\mathcal{PG}(\mathcal{X})$ the set of *parametric (clock) constraints* generated by the grammar [1] $g ::= g \wedge g \mid x - y \leqslant k + b\Delta$, where $x \neq y \in \overline{\mathcal{X}}$, $k \in \mathbb{Q}$ and $b \in \mathbb{N}$. Given a parametric constraint g and $\delta \in \mathbb{Q}_{\geqslant 0}$, we denote by $g(\delta)$ the constraint obtained by evaluating the parameter Δ in δ. As the parameter helps in "relaxing" the clock constraint, we have that $\delta \leqslant \delta'$ implies $[\![g(\delta)]\!] \subseteq [\![g(\delta')]\!]$.

Definition 2 (Parametric Zone). *A* parametric zone $\mathcal{Z}$ *over* $\mathcal{X}$ *is a partial mapping from* $\mathbb{Q}_{\geqslant 0}$ *to zones over* $\mathcal{X}$*, which satisfies the following properties:* (i) *its domain* $\mathsf{dom}(\mathcal{Z})$ *is an interval with rational bounds, and* (ii) *it can be defined as the parametric satisfiability set of a parametric clock constraint,* i.e. *there exists* $g \in \mathcal{PG}(\mathcal{X})$ *such that for all* $\delta \in \mathsf{dom}(\mathcal{Z})$, $\mathcal{Z}(\delta) = [\![g(\delta)]\!]$. *We denote by* $\mathsf{PZones}(\mathcal{X})$ *the set of parametric zones on* $\mathcal{X}$. [2]

By default the considered domain for a parametric zone is $\mathbb{Q}_{\geqslant 0}$. Given a rational interval I, we denote $\mathcal{Z}_{|I}$ the parametric zone whose domain is restricted to I *i.e.*, $\mathsf{dom}(\mathcal{Z}_{|I}) = \mathsf{dom}(\mathcal{Z}) \cap I$, and which coincides with $\mathcal{Z}$ on $\mathsf{dom}(\mathcal{Z}_{|I})$. Given $\mathcal{Z}, \mathcal{Z}' \in \mathsf{PZones}(\mathcal{X})$, we define $\mathcal{Z} \subseteq \mathcal{Z}'$ if, and only if, we have $\mathsf{dom}(\mathcal{Z}) \subseteq \mathsf{dom}(\mathcal{Z}')$, and for any $\delta \in \mathsf{dom}(\mathcal{Z})$, $\mathcal{Z}(\delta) \subseteq \mathcal{Z}'(\delta)$. We say that a parametric zone $\mathcal{Z}$ is non-empty if there exists $\delta \in \mathsf{dom}(\mathcal{Z})$ such that $\mathcal{Z}(\delta) \neq \varnothing$. Let $\mathcal{Z}$ be a non-empty parametric zone. As the mapping represented by $\mathcal{Z}$ is monotone, we define $\delta_{\neg\varnothing}(\mathcal{Z}) = \inf\{\delta \geqslant 0 \mid \mathcal{Z}(\delta) \neq \varnothing\}$ the minimal value of the parameter for the zone it denotes to be nonempty. As $\mathcal{Z}$ only involves non-strict linear inequalities, $\delta_{\neg\varnothing}(\mathcal{Z})$ is a rational number and we have $\mathcal{Z}(\delta_{\neg\varnothing}(\mathcal{Z})) \neq \varnothing$ (provided that $\delta_{\neg\varnothing}(\mathcal{Z}) \in \mathsf{dom}(\mathcal{Z})$).

Definition 3 (Parametric Semantics [Pur00,DDMR08]). *Let* $\mathcal{A} = (L, \ell_0, \mathcal{X}, \Sigma, T)$ *be a TA. The* parametric semantics of $\mathcal{A}$ *consists in replacing each constraint* $g \in \mathcal{G}_{nd}(\mathcal{X})$ *appearing in some transition of* $\mathcal{A}$ *by the parametric constraint obtained by enlarging it with the parameter* Δ*. Formally, each atomic constraint of the form* $x - y \leqslant k$ *is replaced by the parametric constraint* $x - y \leqslant k + \Delta$.

Given $\delta \in \mathbb{Q}_{\geqslant 0}$, the instantiation of all constraints of $\mathcal{A}$ in δ leads to a timed automaton that we denote by $\mathcal{A}(\delta)$. The semantics used implies the following monotonicity property: $\delta \leqslant \delta' \Rightarrow \mathsf{Reach}(\mathcal{A}(\delta)) \subseteq \mathsf{Reach}(\mathcal{A}(\delta'))$. An example of timed automaton is shown in Figure 1.

2.3 Symbolic computations using (parametric) zones

A *symbolic state* is a pair $(\ell, Z) \in L \times \mathsf{Zones}(\mathcal{X})$. Consider a transition $t = (\ell, g, \sigma, r, \ell') \in T$ of a TA $\mathcal{A}$. We define the operator Post^t computing the symbolic successors over t starting from the zone Z, with $Z \in \mathsf{Zones}(\mathcal{X})$, by $\mathsf{Post}^t(Z) = \{v'' \in \mathbb{R}^{\mathcal{X}}_{\geqslant 0} \mid \exists v \in Z, \exists d \in \mathbb{R}_{\geqslant 0} : (\ell, v) \xrightarrow{\sigma} (\ell', v') \xrightarrow{d} (\ell', v' + d) \text{ and } v'' = v' + d\}$. It is well known that $\mathsf{Post}^t(Z)$ is still a zone. We define similarly the operator Pre^t for the set of predecessors by t. Given a sequence of transitions ϱ, we define the operators Post^ϱ and Pre^ϱ as the compositions of these operators for each transition of ϱ. We define the set of successors

[1] Compared with L/U TA introduced in [HRSV02], our parameter is "upper".

[2] In the sequel, Z and Y denote a zone, while $\mathcal{Z}$ and $\mathcal{Y}$ denote a parametric zone.

from a symbolic state by $\mathsf{Succ}(\ell, Z) = \{(\ell', Z') \in L \times \mathsf{Zones}(\mathcal{X}) \mid \exists t = (\ell, g, \sigma, r, \ell') \in T \text{ s.t. } Z' = \mathsf{Post}^t(Z)\}$.

In order to perform parametric computations, we will use parametric zones. Our parametric constraints are less expressive [3] than those considered in [AAB00]. In particular, we can perform the operations of intersection, time elapsing, clock reset, inclusion checking... and extend operators Post^ϱ and Pre^ϱ to a parametric setting. We denote these extensions by PPost^ϱ and PPre^ϱ.We also define the operator $\mathsf{Succ}(\ell, \mathcal{Z})$, where $\mathcal{Z} \in \mathsf{PZones}(\mathcal{X})$, using the PPost operator.

3 The enlarged reachability set $\mathsf{Reach}^*(\mathcal{A})$

Definition of $\mathsf{Reach}^*(\mathcal{A})$. We are interested here in the *quantitative problem of robustness* for safety properties: given a set of states Bad to be avoided, compute the maximal value of δ for the system to be safe, *i.e.* the value $\delta_{\max} = \sup\{\delta \geqslant 0 \mid \mathsf{Reach}(\mathcal{A}(\delta)) \cap \mathsf{Bad} = \varnothing\}$ (recall the monotonicity of $\mathsf{Reach}(\mathcal{A}(\delta))$ *w.r.t.* δ). Note that the value $\delta_{\max}$ may be safe or not (see Examples in Appendix B.4).

In this paper, we propose an algorithm that computes a representation of the parametric reachability set of a flat timed automaton. It is then easy to derive the optimal value $\delta_{\max}$. A standard forward parametric analysis does not terminate in general for timed automata. Such a phenomenon is due to cycles: it can be the case that a state (ℓ, v) is reachable for any $\delta > 0$, but the length of paths allowing to reach (ℓ, v) in $\mathcal{A}(\delta)$ diverges when δ converges to 0.

Example 1. Consider the timed automaton represented on Figure 1. State (ℓ_2, v) with $v(x_1) = 0$ and $v(x_2) = 2$ is reachable in $[\![\mathcal{A}(\delta)]\!]$ for any $\delta > 0$. Let denote by t_1 (resp. t_2) the transition from ℓ_1 to ℓ_2 (resp. from ℓ_2 to ℓ_1), and let $\varrho = t_1 t_2$. In $[\![\mathcal{A}(\delta)]\!]$, this state is reachable only after $\lceil \frac{1}{2\delta} \rceil$ iterations of the cycle ϱ (see Figure 2).

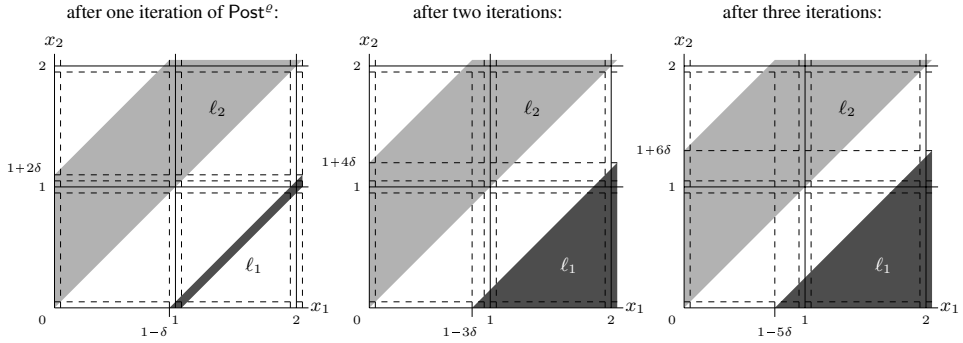

Fig. 2: Reachable states during the parametric forward analysis of $\mathcal{A}(\delta)$.

[3] Note that in our setting, one can define a data structure more specific than parametric DBMs considered in [AAB00]. Indeed, we do not need to split DBMs as the constraints only involve conjunctions. Moreover, we can perform basic operations (future, reset, intersection with an atomic constraint) in quadratic time, as for DBMs, see [Jau09].

This difficulty has first been identified by Puri in [Pur00] when studying the qualitative robustness problem, and solved by computing the enlarged reachability set defined as $\mathsf{Reach}^*(\mathcal{A}) \stackrel{def}{=} \bigcap_{\delta \in \mathbb{Q}_{>0}} \mathsf{Reach}(\mathcal{A}(\delta))$. It is the set of states of the automaton reachable by an arbitrarily small value of the parameter. While [Pur00] proposed an algorithm based on the region graph, we use an algorithm proposed in [DK06] which relies on zones, as it is better suited for a parametric setting. The drawback of [DK06] is that it requires the timed automaton to be flat.

Algorithm 1 Computation of $\mathsf{Reach}^*(\mathcal{A})$.

Require: a progressive flat timed automaton $\mathcal{A}$ with bounded clocks.
Ensure: the set $\mathsf{Reach}^*(\mathcal{A})$.

```
 1: Compute νY.Pre^ϱ(Y), νY.Post^ϱ(Y), for each cycle ϱ in A.
 2: Wait = {(ℓ_0, Z_0)} ;                                   // Initial states
 3: Passed = ∅ ;
 4: while Wait ≠ ∅ do
 5:   pop (ℓ, Z) from Wait ;
 6:   if ∀(ℓ, Z') ∈ Passed, Z ⊈ Z' then
 7:     if there exists a cycle ϱ around location ℓ then
 8:       if Z ∩ νY.Pre^ϱ(Y) ≠ ∅ then
 9:         Wait = Wait ∪ Succ(ℓ, νY.Post^ϱ(Y));
10:         Passed = Passed ∪ {(ℓ, νY.Post^ϱ(Y))};
11:     Wait = Wait ∪ Succ(ℓ, Z) ;
12:     Passed = Passed ∪ {(ℓ, Z)} ;
13: return Passed ;
```

A new procedure for the computation of Reach^*. We present Algorithm 1 which is a modification of the algorithm proposed in [DK06] to compute Reach^*. This modification allows us in the next section to prove the termination of a parametric version of this algorithm.

The original algorithm proposed in [DK06] (see Appendix A) relies on the notion of *stable zone* of a cycle ϱ. This zone represents states having infinitely many predecessors and successors by ϱ, and is defined as the intersection of two greatest fixpoints: $W_\varrho = \nu Y.\mathsf{Post}^\varrho(Y) \cap \nu Y.\mathsf{Pre}^\varrho(Y)$. Then, the algorithm is obtained by the following modifications of the standard forward analysis of the timed automaton: for each new symbolic state (ℓ, Z) considered, if there exists a cycle ϱ around location ℓ, and if Z intersects the stable zone W_ϱ, then the stable zone is marked as reachable. The correction of this algorithm relies on the following property of the stable zone: given two valuations $v, v' \in W_\varrho$, for any $\delta > 0$, there exists a path in $[\![\mathcal{A}(\delta)]\!]$ from state (ℓ, v) to state (ℓ, v') (while such a path may not exist in $[\![\mathcal{A}]\!]$). The addition of the stable zone can be viewed as the acceleration of cycle ϱ.

Our new algorithm is obtained as follows: (i) at line 8, we test the intersection of Z with $\nu Y.\mathsf{Pre}^\varrho(Y)$ instead of W_ϱ, and (ii) at line 9 and 10, instead of declaring W_ϱ as reachable, we declare $\nu Y.\mathsf{Post}^\varrho(Y)$ reachable. We state below that this modification is correct.

Theorem 1. *Algorithm 1 is sound and complete.*

Proof. We show that Algorithm 1 is equivalent to that of [DK06]. As W_ϱ is included in both greatest fixpoints, the completeness of the algorithm is trivial. To prove the soundness, let us consider the region graph construction (see for instance [AD94]). We do not recall this standard construction as it will only be used in this proof. As there are finitely many regions, it is easy to verify that if a region is included in $\nu Y.\mathsf{Pre}^\varrho(Y)$, it has infinitely many successors by ϱ and then one of them is included in W_ϱ. In other terms, the test of line 8 of intersection with $\nu Y.\mathsf{Pre}^\varrho(Y)$ instead of W_ϱ simply anticipates the acceleration of the cycle ϱ. Similarly, any region included in $\nu Y.\mathsf{Post}^\varrho(Y)$ is the successor of a region included in W_ϱ. Thus, our modification can be understood as a speed-up of the original algorithm of [DK06]. □

We also state the following Lemma whose proof follows from a similar reasoning:

Lemma 1. *Let ϱ be a cycle of a TA $\mathcal{A}$. Then we have:*
$\nu Y.\mathsf{Pre}^\varrho(Y) \neq \varnothing \Leftrightarrow \nu Y.\mathsf{Pre}^\varrho(Y) \cap \nu Y.\mathsf{Post}^\varrho(Y) \neq \varnothing \Leftrightarrow \nu Y.\mathsf{Post}^\varrho(Y) \neq \varnothing$

4 Parametric computation of $\mathsf{Reach}(\mathcal{A}(\delta))$

4.1 Representing constraints as a graph

In the sequel, we will use a representation of clock constraints as a weighted directed graph introduced in [CJ99a,CJ99b]. Due to lack of space, we recall here only succintly its definition. Intuitively, the value of a clock can be recovered from its date of reset and the current time. The vertices of the graph represent these values, with one duplicate for each fired transition. Constraints on clock values are expressed as weights on arcs.

More formally, we introduce a new variable τ representing the total elapsed time. In addition, for each clock $x_i \in \overline{\mathcal{X}}$ we let variable X_i denote $X_i = \tau - x_i$. Note that for $x_i \in \mathcal{X}$, X_i thus represents last date of reset of clock x_i. For the special case of x_0, we have $X_0 = \tau$ (as x_0 always has value 0). We denote $\overrightarrow{V}$ the vector defined as $(\tau, X_1, \ldots, X_n)$. For a transition $t = (\ell, g, \sigma, r, \ell')$, we define the formula $T^t(\overrightarrow{V}, \overrightarrow{V'})$ which expresses the relationship between values of the variables before (represented by $\overrightarrow{V}$) and after the firing of the transition (represented by $\overrightarrow{V'} = (\tau', X'_1, \ldots, X'_n)$):

$$T^t(\overrightarrow{V}, \overrightarrow{V'}) := \bigwedge_{i=1}^{n} (X_i \leqslant \tau \wedge X'_i \leqslant \tau') \wedge \tau \leqslant \tau' \\ \wedge \bigwedge_{x_i \in r} \tau = X'_i \wedge \bigwedge_{x_i \notin r} X_i = X'_i \wedge \overline{g}$$

where $\overline{g}$ is the constraint g where for any i, clock x_i is replaced by $\tau - X_i$.

Let $\varrho = t_1 \ldots t_m$ be a sequence of transitions. For $j \in \{0, \ldots, m\}$, we denote by $\overrightarrow{V^j}$ the vector $(\tau^j, X_1^j, \ldots, X_n^j)$. Then we define formula $T^\varrho(\overrightarrow{V^0}, \overrightarrow{V^m})$ expressing the constraints between variables before and after the firing of the sequence ϱ as follows:

$$T^\varrho(\overrightarrow{V^0}, \overrightarrow{V^m}) = \exists \overrightarrow{V^1}, \ldots, \overrightarrow{V^{m-1}}. \bigwedge_{j=0}^{m-1} T^\varrho(\overrightarrow{V^j}, \overrightarrow{V^{j+1}})$$

Definition 4 (Graph $G_\varrho^\top$). *Let $\varrho = t_1 \dots t_m$ be a sequence of transitions. The weighted directed graph $G_\varrho^\top$ has a set of vertices $\mathcal{S} = \bigcup_{j=0}^{m} V^j$ (where V^j is the set associated with the vector $\overrightarrow{V^j}$). Given two vertices $v, v' \in \mathcal{S}$ and a weight $c \in \mathbb{Q}$, there is an arc from v to v' labelled by c if and only if constraint $v - v' \leqslant c$ appears in formula $T^\varrho(\overrightarrow{V^0}, \overrightarrow{V^m})$.*

For any path p, we write $w(p)$ the total weight of the path. Suppose now that there is no cycle of negative weight in graph $G_\varrho^\top$. Let P_{beg}^ϱ (resp. P_{end}^ϱ) denote the set of minimal weighted paths between vertices in V^0 (resp. in $V^{|\varrho|}$). We define the following mapping which interprets these shortest paths as clock constraints:

Let $\alpha = 0$. $\forall p \in P_{beg}^\varrho, C(p) = x_l - x_i \leqslant w(p)$ if p starts in X_i^α and ends in X_l^α.
Mapping C is defined similarly on P_{end}^ϱ, using $\alpha = |\varrho|$.

From Propositions 12 and 13 of [CJ99b], we have the following properties:

Proposition 1. *Let ϱ be a sequence of transitions. Then we have:*

- *there exists a cycle γ with $w(\gamma) < 0$ in $G_\varrho^\top \Leftrightarrow \mathsf{Post}^\varrho(\top) = \varnothing \Leftrightarrow \mathsf{Pre}^\varrho(\top) = \varnothing$*
- *if there is no cycle of negative weight, then:*
  $[\![\bigwedge_{p \in P_{beg}^\varrho} C(p)]\!] = \mathsf{Post}^\varrho(\top)$ *and* $[\![\bigwedge_{p \in P_{end}^\varrho} C(p)]\!] = \mathsf{Pre}^\varrho(\top)$

More generally, given a zone Z, we define the graph denoted G_ϱ^Z by adding the constraints of Z on the vertices in V^0. Mapping C applied on paths in P_{beg}^ϱ then defines the zone $\mathsf{Post}^\varrho(Z)$. Similarly, the zone $\mathsf{Pre}^\varrho(Z)$ can be represented by adding constraints of Z on vertices in $V^{|\varrho|}$.

It is easy to verify that this construction extends to a parametric setting: considering parametric constraints on arcs, we obtain a graph representation of the parametric computation of symbolic successors or predecessors. Note that a path p in this context will have a weight of the form $k + b\Delta$, where $b \in \mathbb{N}$ represents the number of atomic constraints of the TA used in p. In particular, while the value of a path depends on the value of Δ, its existence does not.

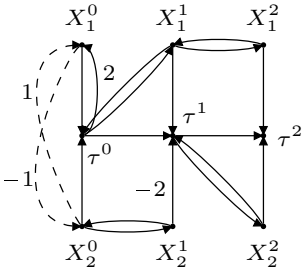

Example 2. Consider the sequence of transitions $\varrho = t_1 t_2$ in the TA of Figure 1 defined in Example 1. The graph depicted on the left-side figure with plain arcs represents $G_\varrho^\top$ (arcs without label have weight 0). For instance, the arc from vertex X_2^2 to vertex τ^2, labelled by -2, represents the lower bound for the clock x_2 in t_2 which means: $x_2 \geqslant 2$. Consider now the zone $Z = [\![x_1 - x_2 = 1]\!]$ (it corresponds to the set of reachable valuations after firing transition $\ell_0 \to \ell_1$), then additional dotted arcs allow to represent G_ϱ^Z.

Given a zone defined as the result of the firing of a sequence of transitions, this representation allows to recover how the constraints are obtained. Thus, the graph stores the complete history of the constraints.

In the sequel, we use this construction in the particular case of the iteration of a cycle ϱ, given as a sequence of transitions of a TA. Let Z_{init} be a zone. We consider two sequences of zones $(Z_k^\top)_{k\geqslant 0}$ (resp. $(Z_k^{init})_{k\geqslant 0}$) defined by $Z_0^\top = \top$ (resp. $Z_0^{init} = Z_{init}$) and $Z_{k+1}^* = \mathsf{Post}^\varrho(Z_k^*)$ (where * denotes either $^\top$ or ^{init}). Note that by monotonicity of Post^ϱ, the sequence $(Z_k^\top)_{k\geqslant 0}$ is decreasing and converges towards $Z_\infty^\top = \nu Y.\mathsf{Post}^\varrho(Y)$. According to previous definitions, $G_{\varrho^k}^\top$ (resp. $G_{\varrho^k}^{init}$) denotes the graph associated with the zone $Z_k^\top$ (resp. Z_k^{init}). As the cycle ϱ will be clear from the context, we will omit to mention it in the subscript, and use notations $G_k^\top$ and G_k^{init} respectively.

Moreover, we will only be interested in vertices at the frontier between the different copies of the graph of ϱ. Then, given a clock $x_i \in \overline{\mathcal{X}}$ and an index $j \leqslant k$, *vertex X_i^j now denotes the date of reset of clock x_i after the j-th execution of ϱ* (this notation is a shorthand for the notation $X_i^{j\times|\varrho|}$, as this last notation will never be used anymore).

Definition 5. *Let $N = |\overline{\mathcal{X}}|^2$. A* return path *is a pair $r = (p_1, p_2)$ of paths in the graph $G_N^\top$ such that there exist two clocks $x_u, x_v \in \overline{\mathcal{X}}$ and two indices $0 \leqslant i < j \leqslant N$ verifying:*

- *p_1 and p_2 are included in the subgraph associated with i-th to j-th copies of ϱ,*
- *p_1 is a shortest path from vertex X_u^j to vertex X_u^i,*
- *p_2 is a shortest path from vertex X_v^i to vertex X_v^j.*

The weight *of r is defined as $w(r) = w(p_1) + w(p_2)$. The set of return paths is finite and is denoted $\mathcal{R}$.*

4.2 Accelerating computations of greatest fixpoints

Let ϱ be a cycle. In this subsection, we only consider the operator Post^ϱ, but all our results also apply to the operator Pre^ϱ. We consider the decreasing sequence $(Z_k^\top)_{k\geqslant 0}$ converging towards $Z_\infty^\top = \nu Y.\mathsf{Post}^\varrho(Y) = \bigcap_{k\geqslant 0} Z_k^\top$. We prove the following lemma which provides a bound for termination only dependant on the number of clocks. Note that this result does not require the cycle ϱ to be progressive neither the clocks to be bounded.

Lemma 2. *Let $N = |\overline{\mathcal{X}}|^2$, and $k \geqslant N$. If $Z_{k+1}^\top \subsetneq Z_k^\top$, then we have $Z_\infty^\top = \varnothing$.*

Proof. First, we prove that $Z_{k+1}^\top \subsetneq Z_k^\top$ implies that there exists $r \in \mathcal{R}$ used in some shortest path of $Z_{k+1}^\top$ witness of the disequality. Indeed, as $Z_{k+1}^\top \subsetneq Z_k^\top$, there exists a bound $b = "x_p - x_q \leqslant \cdot"$ with $0 \leqslant p \neq q \leqslant n$, whose constraint is strictly smaller in $Z_{k+1}^\top$ than in $Z_k^\top$. In $Z_{k+1}^\top$, the constraint on b is obtained as a shortest path between vertices X_p^{k+1} and X_q^{k+1} in the graph $G_{k+1}^\top$. Let c be such a path. By definition of $G_k^\top$ and $G_{k+1}^\top$, the path c must use arcs in $G_1^\top$ (otherwise c would also exist in $G_k^\top$). The graph $G_{k+1}^\top$ is the concatenation of $k+1$ copies of the graph of ϱ. For each occurrence of ϱ, c goes through a pair of vertices when it enters/leaves it. Finally, as $k+1 > N = |\overline{\mathcal{X}}|^2$, there exists a pair that occurs twice, we denote these two clocks x_u and x_v. Thus c contains a return path $r \in \mathcal{R}$ (see Figure 3 representing the graph $G_{k+1}^\top$ and the return path r in the shortest path c).

Second, as $Z^{\top}_{k+1} \subsetneq Z^{\top}_{k}$, we have $w(r) < 0$. By contradiction, if $w(r) > 0$ then c would not be a shortest path and if $w(r) = 0$ then c would also exist in $G^{\top}_{k}$.

Finally, the existence of a return path $r \in \mathcal{R}$ such that $w(r) < 0$ implies that $Z^{\top}_{\infty} = \varnothing\ (= \nu Y.\mathsf{Post}^{\varrho}(Y))$. When k grows, one can build new paths by repeating this return path. As its weight is negative, the weights of the paths we build diverge towards $-\infty$. In particular, the constraint of the zone $Z^{\top}_{\infty}$ on the clock difference $x_p - x_q$ cannot be finite (as it is the limit of a sequence diverging towards $-\infty$), and thus we obtain $Z^{\top}_{\infty} = \varnothing$. □

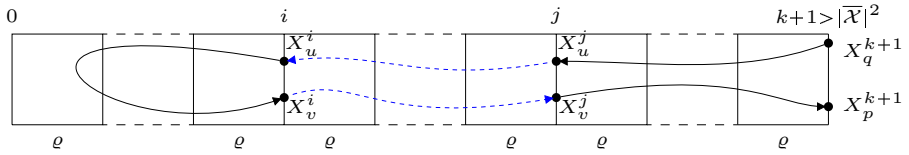

Fig. 3: Pumping lemma : a path from X^{k+1}_q to X^{k+1}_p using arcs in $G^{\top}_1$ exhibits a return path between pairs of vertices (X^i_u, X^i_v) and (X^j_u, X^j_v).

We can now compute, in the parametric setting, the greatest fixpoint of PPost^{ϱ} for every cycle ϱ of the automaton. We first evaluate the parametric zones $\mathcal{Z} = \mathsf{PPost}^{\varrho^N}(\top)$ and $\mathcal{Z}' = \mathsf{PPost}^{\varrho}(\mathcal{Z})$. Then, we determine the minimal value $\delta_0 = \min\{\delta \geqslant 0 \mid \mathcal{Z}(\delta) = \mathcal{Z}'(\delta)\}$. This definition is correct as $\mathcal{Z}' \subseteq \mathcal{Z}$ and there exists δ for which the greatest fixpoint is not empty. Finally the greatest fixpoint can be represented by $\mathcal{Z}_{|[\delta_0;+\infty[}$ as Lemma 2 ensures that the fixpoint is empty for all $\delta < \delta_0$.

4.3 Parametric Forward analysis with acceleration

We present Algorithm 2 for the parametric computation of $\mathsf{Reach}(\mathcal{A}(\delta))$. It can be understood as an adaptation in a parametric setting of Algorithm 1. First, at line 1 we perform parametric computation of greatest fixpoints using the procedure proposed in Section 4.2. Second, the test of intersection between the current zone and the greatest fixpoint of Pre^{ϱ} is realized in a parametric setting by the computation at line 8 of $\delta_{\min} = \delta_{\neg\varnothing}(\mathcal{Z} \cap \nu\mathcal{Y}.\mathsf{PPre}^{\varrho}(\mathcal{Y}))$. Finally, we split the domain of the current parametric zone into intervals I_1 and I_2. In interval I_1, no acceleration is done for cycles and thus the set $\mathsf{Reach}(\mathcal{A}(\delta))$ is computed. Acceleration techniques are used only for interval I_2, and for these values the algorithm computes the set $\mathsf{Reach}^*(\mathcal{A}(\delta))$. We prove below that in this case, the equality $\mathsf{Reach}(\mathcal{A}(\delta)) = \mathsf{Reach}^*(\mathcal{A}(\delta))$ holds. Note that the test at line 9 allows to handle differently the particular case of value $\delta_{\min}$ which does not always require to apply acceleration.

Theorem 2. *Algorithm 2 terminates and is correct.*

In the sequel, we denote $N = |\overline{\mathcal{X}}|^2$ and $\delta^{\varrho}_{\neg\varnothing} = \delta_{\neg\varnothing}(\nu\mathcal{Y}.\mathsf{PPre}^{\varrho}(\mathcal{Y})) = \delta_{\neg\varnothing}(\nu\mathcal{Y}.\mathsf{PPost}^{\varrho}(\mathcal{Y}))$ (by Lemma 1). Before turning to the proof, we state the following Lemma whose proof

Algorithm 2 Parametric Computation of the Reachability Set.

Require: a progressive flat timed automaton $\mathcal{A}$ with bounded clocks.
Ensure: the set $\mathsf{Reach}(\mathcal{A}(\delta))$ for all $\delta \in \mathbb{R}_{\geqslant 0}$.

Compute $\nu\mathcal{Y}.\mathsf{PPre}^{\varrho}(\mathcal{Y})$ and $\nu\mathcal{Y}.\mathsf{PPost}^{\varrho}(\mathcal{Y})$ for each cycle ϱ of $\mathcal{A}$.
Wait $= \{(\ell_0, \mathcal{Z}_0)\}$; *// Initial States*
Passed $= \varnothing$;
while Wait $\neq \varnothing$ **do**
  pop $(\ell, \mathcal{Z})$ from Wait ;
  if $\forall(\ell, \mathcal{Z}') \in$ Passed, $\mathcal{Z} \nsubseteq \mathcal{Z}'$ **then**
    if there exists a cycle ϱ around location ℓ **then**
      $\delta_{\min} = \delta_{\neg\varnothing}(\mathcal{Z} \cap \nu\mathcal{Y}.\mathsf{PPre}^{\varrho}(\mathcal{Y}))$;
      if $\delta_{\min} = \delta_{\neg\varnothing}(\nu\mathcal{Y}.\mathsf{PPre}^{\varrho}(\mathcal{Y}))$ **then**
        $I_1 = [0; \delta_{\min}]$; $I_2 =]\delta_{\min}; +\infty[$;
      else
        $I_1 = [0; \delta_{\min}[$; $I_2 = [\delta_{\min}; +\infty[$;
      Wait $=$ Wait $\cup\ \mathsf{Succ}(\ell, \mathcal{Z}_{|I_1}) \cup \mathsf{Succ}(\ell, \mathcal{Z}_{|I_2}) \cup \mathsf{Succ}(\ell, \nu\mathcal{Y}.\mathsf{PPost}^{\varrho}(\mathcal{Y})_{|I_2})$;
      Passed $=$ Passed $\cup\ (\ell, \mathcal{Z}_{|I_1}) \cup (\ell, \mathcal{Z}_{|I_2}) \cup (\ell, \nu\mathcal{Y}.\mathsf{PPost}^{\varrho}(\mathcal{Y})_{|I_2})$;
    else
      Wait $=$ Wait $\cup\ \mathsf{Succ}(\ell, \mathcal{Z})$;
      Passed $=$ Passed $\cup\ (\ell, \mathcal{Z})$;
return Passed ;

is given in Appendix B.2. Intuitively, it establishes that when all return paths have a positive weight, then either (i) the starting zone has finitely many successors and then it converges to the empty set after at most N steps, or (ii) it has infinitely many successors and then it converges towards $\nu Y.\mathsf{Post}^{\varrho}(Y)$. In this last case, the enlarged reachability set corresponds to the standard reachability set. Its proof relies on pumping techniques presented in Section 4.2. To illustrate property (ii), let consider the timed automaton of Figure 1, for which the enlarged reachability set strictly contains the standard reachability set. One can verify that there exists a return path associated with $\varrho = t_1 t_2$ which has weight 0.

Lemma 3. *Let ϱ be such that for any return path $r \in \mathcal{R}$, we have $w(r) > 0$. Then we have:*

(i) *If $Z_{init} \cap \nu Y.\mathsf{Pre}^{\varrho}(Y) = \varnothing$, then $Z_N^{init} = \varnothing$.*
(ii) *If $Z_{init} \cap \nu Y.\mathsf{Pre}^{\varrho}(Y) \neq \varnothing$, then $Z_{\infty}^{init} = Z_{\infty}^{\top} (= \nu Y.\mathsf{Post}^{\varrho}(Y))$.*

Unlike Lemma 2, we use the progess cycle assumption to prove this lemma (see the proof of Lemma 4 in Appendix B.1).

Recall that the TA we consider are flat. As a consequence, in the following proofs of termination and correctness, we will only consider a simple cycle ϱ.

Termination. Consider a parametric symbolic state $(\ell, \mathcal{Z})$ and a cycle ϱ starting in ℓ. We have to prove that all the elements added to the Wait list have a finite number of successors. This is trivial for the successors of $(\ell, \nu\mathcal{Y}.\mathsf{PPost}^{\varrho}(\mathcal{Y})_{|I_2})$ as $\nu\mathcal{Y}.\mathsf{PPost}^{\varrho}(\mathcal{Y})_{|I_2}$ is by definition a fixpoint of PPost^{ϱ}. We now focus on the successors of $(\ell, \mathcal{Z}_{|I_1})$ and $(\ell, \mathcal{Z}_{|I_2})$. Note that we have $\delta_{\min} \geqslant \delta^{\varrho}_{\neg\varnothing}$.

– **Case of** $(\ell, \mathcal{Z}_{|I_2})$**:** We prove property $(*)$ $\mathsf{PPost}^{\varrho^N}(\mathcal{Z}_{|I_2}) \subseteq \nu\mathcal{Y}.\mathsf{PPost}^{\varrho}(\mathcal{Y})_{|I_2}$. Then the computation is stopped by the test of line 6 as the greatest fixpoint has been added to the Passed list. To prove $(*)$, we prove it holds for any $\delta \in I_2$. Fix some $\delta \in I_2$ and define $Z_{init} = \mathcal{Z}_{|I_2}(\delta)$. We consider the two sequences $(Z_i^*)_{i\geqslant 0}$ *w.r.t.* cycle ϱ enlarged by δ. Note that as $\delta \geqslant \delta_{\min} \geqslant \delta^{\varrho}_{\neg\varnothing}$, we have $\nu\mathcal{Y}.\mathsf{PPost}^{\varrho}(\mathcal{Y})(\delta) \neq \varnothing$. By Lemma 2, this entails $Z_N^{\top} = \nu\mathcal{Y}.\mathsf{PPost}^{\varrho}(\mathcal{Y})(\delta)$. By monotonicity of Post^{ϱ}, $Z_N^{init} \subseteq Z_N^{\top}$ holds. This yields the result.
– **Case of** $(\ell, \mathcal{Z}_{|I_1})$**:** We distinguish two cases whether $\delta_{\min} > \delta^{\varrho}_{\neg\varnothing}$ or not.

  If $\delta_{\min} > \delta^{\varrho}_{\neg\varnothing}$**:** for any $\delta \in [\delta^{\varrho}_{\neg\varnothing}, \delta_{\min}[$, Lemma 3.$(i)$ can be applied on cycle ϱ enlarged by δ. This implies that for any $\delta \in [\delta^{\varrho}_{\neg\varnothing}, \delta_{\min}[$, we have $\mathsf{PPost}^{\varrho^N}(\mathcal{Z}_{|I_1})(\delta) = \varnothing$. Then this property also holds for any $\delta \in I_1$, by monotonicity of $\mathcal{Z}$ and PPost^{ϱ}.

  If $\delta_{\min} = \delta^{\varrho}_{\neg\varnothing}$**:** the complete proof of this last case is more technical and is completely described in Appendix B.3. We only present here a sketch of proof. First note that for any fixed value of $\delta < \delta_{\min}$, as the zone does not intersect the greatest fixpoint of Pre^{ϱ}, the zone has finitely many successors. However, this argument cannot be lifted to a parametric setting as this number diverges when δ converges towards $\delta_{\min}$. By definition of $\delta^{\varrho}_{\neg\varnothing}$, some return paths, which we call *optimal*, have a weight equal to 0 in $\delta^{\varrho}_{\neg\varnothing}$ (and are thus strictly negative on $[0, \delta^{\varrho}_{\neg\varnothing}[$). Our proof consists in first showing that there exists some integer k for which after k steps, all shortest paths go through optimal return paths. Then, considering q as the least common multiple of lengths of optimal return paths, we can prove the following inclusion $\mathsf{PPost}^{\varrho^{k+q}}(\mathcal{Z}_{|I_1}) \subseteq \mathsf{PPost}^{\varrho^k}(\mathcal{Z}_{|I_1})$. The algorithm stops by test of line 6.

Correctness. As explained before, the algorithm is a standard forward analysis which may add some additional behaviours, according to test of line 8. We distinguish three cases:

1. **For** $\delta \in [0, \delta_{\min}[$ **:** For these values, the algorithm simply performs a forward analysis. As a consequence, the correctness is trivial.
2. **For** $\delta \in]\delta_{\min}, +\infty[$**:** For all these values, the addition occurs, and then the algorithm is equivalent to Algorithm 1. By correction of Algorithm 1, this implies that it computes the set $\mathsf{Reach}^*(\mathcal{A}(\delta))$. We will prove that for all these values, we have the equality $\mathsf{Reach}(\mathcal{A}(\delta)) = \mathsf{Reach}^*(\mathcal{A}(\delta))$. Therefore we need to prove that what has been added to obtain $\mathsf{Reach}^*(\mathcal{A}(\delta))$ was already in $\mathsf{Reach}(\mathcal{A}(\delta))$. Note that the only addition is the greatest fixpoint of Post^{ϱ}. The property is then a direct consequence of Lemma 3.(ii) as it states that the greatest fixpoint is reachable from the initial states. It is easy to verify that Lemma 3.(ii) can indeed be applied.
3. **For** $\delta = \delta_{\min}$**:** There are two cases, whether $\delta_{\min} = \delta^{\varrho}_{\neg\varnothing}$ or not. If the equality holds, then $\delta_{\min} \in I_1$ and the reasoning developed at point 1. also applies. If $\delta_{\min} > \delta^{\varrho}_{\neg\varnothing}$ holds, then $\delta_{\min} \in I_2$ and we can apply reasoning of point 2. as Lemma 3.(ii) also applies because we have $\delta_{\min} > \delta^{\varrho}_{\neg\varnothing}$.

4.4 Quantitative safety

Once the reachable state space of the automaton is computed by Algorithm 2, it is easy to compute the maximal value of the parameter such that the system avoids some set of bad states. Simply compute the value $\delta_{\neg\varnothing}$ on each parametric zone associated with a bad location and keep the lower one: $\delta_{\max} = \min\{\delta_{\neg\varnothing}(\mathcal{Z}) \mid \exists \ell \in Bad \text{ such that } (\ell, \mathcal{Z}) \in \text{Passed}\}$. We thus obtain:

Theorem 3. *The quantitative robustness problem for safety properties is decidable for flat progressive timed automata with bounded clocks. In addition, the value $\delta_{\max}$ is a rational number.*

5 Conclusion

In this paper, we considered the quantitative robustness problem for safety properties, which aims at computing the largest value of the parameter Δ under which the TA is safe. We proposed a symbolic forward algorithm for the computation of the parametric reachability set for flat timed automata. We proved its termination by means of original arguments using a representation of zones by graphs. As a consequence, it allows us to compute the largest safe value of the parameter, and prove it is a rational number.

There are several extensions we want to investigate. First, we are implementing the algorithm using a data structure specific to the parametric zones used in our setting. Second, we want to study the complexity of our algorithm. The difficulty is due to the argument of termination in the last case which leads to a large value and may be improved.

We also aim at enlarging the class of TA for which we can solve the quantitative robustness problem. For instance, if the parameter is not always introduced on guards with coefficient 1, but with other coefficients in $\mathbb{N}_{>0}$, we believe that our algorithm can also be applied. A challenging topic concerns the hypothesis of flatness: we plan to investigate a parametric extension of the algorithm introduced in [Dim07] which can be seen as an extension of that of [DK06] to non-flat TA.

Finally, we believe that it should be possible to solve the quantitative robustness problem for flat TA for other specifications like for instance LTL properties.

References

AAB00. A. Annichini, E. Asarin, and A. Bouajjani. Symbolic techniques for parametric reasoning about counter and clock systems. In *Proc. CAV'00*, vol. 1855 of *LNCS*, pp. 419–434. Springer, 2000.

AD94. R. Alur and D. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.

ALM05. R. Alur, S. La Torre, and P. Madhusudan. Perturbed timed automata. In *Proc. HSCC'05*, vol. 3414 of *LNCS*, pp. 70–85. Springer, 2005.

BBB+07. C. Baier, N. Bertrand, P. Bouyer, Th. Brihaye, and M. Größer. Probabilistic and topological semantics for timed automata. In *Proc. FSTTCS'07*, vol. 4855 of *LNCS*, pp. 179–191. Springer, 2007.

BDL04. G. Behrmann, A. David, and K. G. Larsen. A tutorial on UPPAAL. In *Proc. SFM-04:RT*, vol. 3185 of *LNCS*, pp. 200–236. Springer, 2004.

BDM+98. M. Bozga, C. Daws, O. Maler, A. Olivero, S. Tripakis, and S. Yovine. Kronos: a model-checking tool for real-time systems. In *Proc. CAV'98*, vol. 1427 of *LNCS*, pp. 546–550. Springer, 1998.

BIK10. M. Bozga, R. Iosif, and F. Konecný. Fast acceleration of ultimately periodic relations. In *Proc. CAV'10*, vol. 6174 of *LNCS*, pp. 227–242. Springer, 2010.

BIL06. M. Bozga, R. Iosif, and Y. Lakhnech. Flat parametric counter automata. In *Proc. ICALP'06*, vol. 4052 of *LNCS*, pp. 577–588. Springer, 2006.

BMR06. P. Bouyer, N. Markey, and P.-A. Reynier. Robust model-checking of linear-time properties in timed automata. In *Proc. LATIN'06*, vol. 3887 of *LNCS*, pp. 238–249. Springer, 2006.

BMR08. P. Bouyer, N. Markey, and P.-A. Reynier. Robust analysis of timed automata via channel machines. In *Proc. FoSSaCS'08*, vol. 4962 of *LNCS*, pp. 157–171. Springer, 2008.

CHR02. F. Cassez, Th. A. Henzinger, and J.-F. Raskin. A comparison of control problems for timed and hybrid systems. In *Proc. HSCC'02*, vol. 2289 of *LNCS*, pp. 134–148. Springer, 2002.

CJ99a. H. Comon and Y. Jurski. Timed automata and the theory of real numbers. In *Proc. CONCUR'99*, vol. 1664 of *LNCS*, pp. 242–257. Springer, 1999.

CJ99b. H. Comon and Y. Jurski. Timed automata and the theory of real numbers. Research Report LSV-99-6, Laboratoire Spécification et Vérification, ENS Cachan, France, July 1999. 44 pages.

DDMR08. M. De Wulf, L. Doyen, N. Markey, and J.-F. Raskin. Robust safety of timed automata. *Formal Methods in System Design*, 33(1-3):45–84, 2008.

DDR05. M. De Wulf, L. Doyen, and J.-F. Raskin. Almost ASAP semantics: from timed models to timed implementations. *Formal Aspects of Computing*, 17(3):319–341, 2005.

Dim07. C. Dima. Dynamical properties of timed automata revisited. In *Proc. FORMATS'07*, vol. 4763 of *LNCS*, pp. 130–146. Springer, 2007.

DK06. C. Daws and P. Kordy. Symbolic robustness analysis of timed automata. In *Proc. FORMATS'06*, vol. 4202 of *LNCS*, pp. 143–155. Springer, 2006.

GHJ97. V. Gupta, Th. A. Henzinger, and R. Jagadeesan. Robust timed automata. In *Proc. HART'97*, vol. 1201 of *LNCS*, pp. 331–345. Springer, 1997.

HRSV02. T. Hune, J. Romijn, M. Stoelinga, and F. Vaandrager. Linear parametric model checking of timed automata. *Journal of Logic and Algebraic Programming*, 2002.

Jau09. R. Jaubert. Aspects quantitatifs dans la réalisation de contrôleurs temps-réels robustes. Mémoire de Master Recherche, Master Informatique Fondamentale, Marseille, 2009.

OW03. J. Ouaknine and J. B. Worrell. Revisiting digitization, robustness and decidability for timed automata. In *Proc. LICS'03*. IEEE Computer Society Press, 2003.

Pur00. A. Puri. Dynamical properties of timed automata. *Discrete Event Dynamic Systems*, 10(1-2):87–113, 2000.

SF07. M. Swaminathan and M. Fränzle. A symbolic decision procedure for robust safety of timed systems. In *Proc. TIME'07*, p. 192. IEEE Computer Society Press, 2007.

SFK08. M. Swaminathan, M. Fränzle, and J.-P. Katoen. The surprising robustness of (closed) timed automata against clock-drift. In *Proc. TCS'08*, vol. 273 of *IFIP*, pp. 537–553. Springer, 2008.

WT99. H. Wong-Toi. Analysis of slope-parametric rectangular automata. In *Proc. Hybrid Systems'97*, vol. 1567 of *LNCS*, pp. 390–413. Springer, 1999.

A Algorithm of [DK06] for $\mathsf{Reach}^*(\mathcal{A})$

We present the algorithm proposed in [DK06] for the computation of the set $\mathsf{Reach}^*(\mathcal{A})$.

Algorithm 3 Algorithm of [DK06] for the computation of $\mathsf{Reach}^*(\mathcal{A})$

Require: a progressive flat timed automaton $\mathcal{A}$ with bounded clocks.
Ensure: the set $\mathsf{Reach}^*(\mathcal{A})$.

Compute $W_\varrho = \nu Y.\mathsf{Pre}^\varrho(Y) \cap \nu Y.\mathsf{Post}^\varrho(Y)$, for each cycle ϱ in $\mathcal{A}$.
Wait $= \{(\ell_0, Z_0)\}$; *// Initial States*
Passed $= \varnothing$;
while Wait $\neq \varnothing$ **do**
  pop (ℓ, Z) from Wait ;
  if $\forall(\ell, Z') \in$ Passed, $Z \not\subseteq Z'$ **then**
    if there exists a cycle ϱ around location ℓ **then**
      if $Z \cap W_\varrho \neq \varnothing$ **then**
        Wait $=$ Wait $\cup\ \mathsf{Succ}(\ell, W_\varrho)$;
        Passed $=$ Passed $\cup\ (\ell, W_\varrho)$;
    Wait $=$ Wait $\cup\ \mathsf{Succ}(\ell, Z)$;
    Passed $=$ Passed $\cup\ (\ell, Z)$;
return Passed ;

B Complements on Section 4

B.1 Preliminaries

In the following proofs, we will need to consider the weighted directed graph associated with a sequence of transitions in a parametric setting. The weight of an arc is then a parametric constraint (only arcs representing transitions of the TA are enlarged with the parameter). Given a path p in such a graph and a value $\delta \geqslant 0$, we denote by $w_\delta(p)$ the weight obtained when evaluating the parameter Δ in the value δ. Given a return path $r \in \mathcal{R}$, the *length of* r, defined as $j - i$ (with respect to Definition 5) , is denoted $|r|$.

There exists a standard data structure for representing zones which is called Difference Bound Matrix (DBM for short). We will not introduce its definition but assume the reader is familiar with it. Given a bound $b =$ "$x - y \leqslant \cdot$" and a non-empty zone Z, we denote by $Z[b]$ the value of the DBM in normal form associated with Z (which is either $+\infty$ or a relative number as all constraints are non-strict).

Finally, we prove the following lemma:

Lemma 4. *Let ϱ be a progress cycle. We consider the sequences $(Z_k^{init})_{k\geqslant 0}$ and $(Z_k^\top)_{k\geqslant 0}$. Let $k > N$ and $b =$ "$x - y \leqslant \cdot$" be a bound. Then we have $Z_k^{init}[b] < +\infty$ if and only if $Z_k^\top[b] < +\infty$.*

Proof. The right to left implication is trivial as $Z_k^{init} \subseteq Z_k^{\top}$. For the direct implication, the result follows from an examination of the form of the shortest paths. Assume that there exists a shortest path p associated with b in the graph G_k^{init}. Then if p does not exist in $G_k^{\top}$, this implies that p goes through arcs encoding Z_{init}. As ϱ is a progress cycle, we can substitute to these constraints another path p' in the graph $G_k^{\top}$ by going via clock τ. □

B.2 Proof of Lemma 3

Lemma 3. *Let ϱ be such that for any return path $r \in \mathcal{R}$, we have $w(r) > 0$. Then we have:*

(i) *If $Z_{init} \cap \nu Y.\mathsf{Pre}^{\varrho}(Y) = \varnothing$, then $Z_N^{init} = \varnothing$.*
(ii) *If $Z_{init} \cap \nu Y.\mathsf{Pre}^{\varrho}(Y) \neq \varnothing$, then $Z_{\infty}^{init} = Z_{\infty}^{\top} (= \nu Y.\mathsf{Post}^{\varrho}(Y))$.*

Proof. We consider successively these two properties:

(i) We prove the following property:

$$Z_N^{init} \neq \varnothing \Rightarrow Z_{\infty}^{init} \neq \varnothing \tag{1}$$

This concludes the proof as Z_{init} only has a finite number of non empty successors by Post^{ϱ} (because we assume $Z_{init} \cap \nu Y.\mathsf{Pre}^{\varrho}(Y) = \varnothing$) and thus we must have $Z_{\infty}^{init} = \varnothing$. To prove (1), we show that for any $i \geqslant N$, and any bound b, we have $Z_i^{init}[b] \geqslant \min_{j=0}^{N} Z_j^{init}[b]$. Let i and b, and consider a shortest path p associated with b in Z_i^{init}. Either p crosses less than N different copies of ϱ, and then p also exists in Z_N^{init}, yielding the result. Otherwise, as its length is larger than N, as it is done in the proof of Lemma 2, we can prove it contains some return path. We can iterate this reasoning until the resulting length is less or equal than N. Finally, we obtain a decomposition of it, exhibiting some return paths $r_i \in \mathcal{R}$ and a shorter path p' crossing less than N copies of ϱ. By hypothesis, we have $w(r_i) > 0$ for any i. In particular, we obtain $w(p) = \sum_i w(r_i) + w(p') \geqslant w(p')$. Let denote by $j \leqslant N$ the number of copies of ϱ crossed by p'. Then we have $Z_i^{init}[b] = w(p) \geqslant w(p') \geqslant Z_j^{init}[b]$, proving the property.

(ii) We need to show that $Z_k^{init} \underset{k \to \infty}{\longrightarrow} Z_{\infty}^{\top}$. However, the sequence $(Z_k^{init})_k$ is not necessarily increasing. We will prove this result for each bound $b = "x - y \leqslant \cdot"$. Let us fix a bound b, and an integer $k \geqslant N$.
By Lemma 4, if coefficient $Z_k^{init}[b]$ is infinite, then the result holds.
Consider now a finite coefficient $Z_k^{init}[b] < +\infty$. If a shortest path p associated with b in graph G_k^{init} does not enter the arcs representing the initial zone Z^{init}, then this shortest path also exists in graph $G_k^{\top}$ and thus we obtain $Z_k^{init}[b] = Z_k^{\top}[b]$. By contradiction, assume thus that all shortest paths do enter these arcs. As a consequence, these paths are "long", as they cross k copies of ϱ. We will prove that in this case the values of the coefficient converge towards $Z_{\infty}^{\top}[b]$ when k diverges. Define $\eta = \min\{\frac{w(r)}{|r|} \mid r \in \mathcal{R}\}$. It represents the minimal weight that is accumulated through one iteration of ϱ using a return path. By hypotheses,

all return paths have a strictly postive weight and thus we have $\eta > 0$. Using a reasonning similar to that of point (i), we can prove the following inequality:

$$Z_k^{init}[b] \geqslant \left(\min_{j=0}^{N} Z_j^{init}[b]\right) + (k-N) \times \eta \underset{k\to\infty}{\longrightarrow} +\infty \tag{2}$$

By Lemma 4, we have that $Z_k^{\top}[b]$ is finite. In addition, by Lemma 2 and as $k > N$, we know that $Z_k^{\top} = Z_N^{\top} = Z_{\infty}^{\top}$. As a consequence, property (2) does not hold for all k and thus for some finite value of k, we obtain $Z_k^{init}[b] = Z_k^{\top}[b]$. This can be proven for any bound b, yielding the result. Note that the smaller is η, the slower is the convergence. This is precisely the setting of Figure 2.

□

B.3 Complements on Termination of Algorithm 2

We give some complements on the proof of termination of Algorithm 2. Indeed, we only sketched the proof for the following case:
Case of $(\ell, \mathcal{Z}_{|I_1})$, and $\delta_{\min} = \delta^{\varrho}_{\neg\varnothing}$: Recall that we consider an initial parametric zone $\mathcal{Z}$ and a cycle ϱ. By Lemma 2.(ii), we know that for any $\delta \geqslant \delta^{\varrho}_{\neg\varnothing}$, and any return path $r \in \mathcal{R}$, we have $w_\delta(r) \geqslant 0$ (because the existence of a return path of negative weight implies that greatest fixpoints of Post and Pre are empty). We define the set of *optimal return paths* as follows: let $\mathcal{R}_{opt} = \{r \in \mathcal{R} \mid w_{\delta^{\varrho}_{\neg\varnothing}}(r) = 0\}$. Then for any $r \in \mathcal{R}\backslash\mathcal{R}_{opt}$, we have $w_{\delta^{\varrho}_{\neg\varnothing}}(r) > 0$. Intuitively, once a shortest path associated with a bound b goes through an optimal return path, the value of the bound b cannot diverge towards $+\infty$. We will show that eventually, all shortest paths go through an optimal return path. Let $\delta_2 = \max\{\delta \geqslant 0 \mid \exists r \in \mathcal{R}\backslash\mathcal{R}_{opt}\ w_\delta(r) = 0\}$. As $\mathcal{R}$ is finite and as the parametric weight of a return path is an affine function, one obtains that δ_2 is rational. We now define $\delta_3 = (\delta_2 + \delta^{\varrho}_{\neg\varnothing})/2$, which is thus also a rational number, and divide interval I_1 into $I_1' = [0, \delta_3]$ and $I_1'' =]\delta_3, \delta^{\varrho}_{\neg\varnothing}[$. As $\delta_3 < \delta^{\varrho}_{\neg\varnothing}$, we have $\mathcal{Z}(\delta_3) \cap \nu\mathcal{Y}.\mathsf{PPre}^{\varrho}(\mathcal{Y})(\delta_3) = \varnothing$. This implies that there exists an integer k such that $\mathsf{PPost}^{\varrho^k}(\mathcal{Z})(\delta_3) = \varnothing$ (the value of k can be estimated for instance via the region graph construction applicated to the TA $\mathcal{A}(\delta_3)$ – this is possible because δ_3 is a rational number). By monotonicity of $\mathcal{Z}$ and PPost^{ϱ}, we obtain the same property for any $\delta \in I_1'$. We now consider the interval I_1''. First, there exists a postive rational number $\eta > 0$ such that the following property holds:

$$\forall \delta \in I_1'', \forall r \in \mathcal{R}\backslash\mathcal{R}_{opt}, \frac{w_\delta(r)}{|r|} \geqslant \eta$$

Intuitively, this means that any non-optimal return path r will have weight at least $\eta \times |r|$. The existence of η follows from the definition of δ_2 and the fact that $\forall \delta \in I_1'', |\delta - \delta_2| \geqslant \frac{1}{2}|\delta^{\varrho}_{\neg\varnothing} - \delta_2|$. We defined parametric zones $\mathcal{Z}_N^{init} = \mathsf{PPost}^{\varrho^N}(\mathcal{Z}_{|I_1''})$ and $\mathcal{Z}_N^{\top} = \mathsf{PPost}^{\varrho^N}(\top_{|I_1''})$. Consider now a bound $b =$ "$x - y \leqslant \cdot$". We can define $d_b = \max\{|\mathcal{Z}_N^{init}(\delta)[b] - \mathcal{Z}_N^{\top}(\delta)[b]| \mid \delta \in I_1''\}$. This value is finite by Lemma 4. Then, $\forall i > i_b = N + N \times \frac{d_b}{\eta}$, each shortest path associated with bound b must contain an

optimal return path $r \in \mathcal{R}_{opt}$. Otherwise we would obtain that $\mathcal{Z}_i^{init}(\delta)[b] > \mathcal{Z}_i^{\top}(\delta)[b]$, which is a contradiction. Define now $k' = \max\{i_b \mid b\} \cup \{k\}$, and let q be the least common multiple of the set $\{|r| \mid r \in \mathcal{R}_{opt}\}$. Then we obtain $\mathcal{Z}_{k'+q}^{init}(\delta) \subseteq \mathcal{Z}_{k'}^{init}(\delta)$ for any $\delta \in I_1''$. This concludes this last case because, as $k' \geqslant k$, we also have for any $\delta \in I_1'$, $\mathcal{Z}_{k'+q}^{init}(\delta) = \varnothing$.

B.4 Examples for safety of $\mathcal{A}(\delta_{\max})$

We present here some additional examples to illustrate the fact that the TA $\mathcal{A}(\delta_{\max})$ may be safe or not. These examples are slight variations of the TA depicted on Figure 1.

On Figures 4 and 5, in dark blue (resp. dark red) is depicted the reachable set $\mathsf{Reach}(\mathcal{A}(0))$ in location ℓ_1 (resp. ℓ_2). Light colors represent the sets $\mathsf{Reach}^*(\mathcal{A}(0))$. Note that for the TA of Figure 4, as $\delta_{\min} = \delta^{\varrho}_{\neg\varnothing}$, we have that $0 \in I_1$ (in Algorithm 2). We thus obtain that 0 is a safe value for Δ.

On Figure 6, in dark blue (resp. dark red) is depicted the reachable set $\mathsf{Reach}(\mathcal{A}(\frac{1}{2}))$ in location ℓ_1 (resp. ℓ_2). Light colors represent the sets $\mathsf{Reach}^*(\mathcal{A}(\frac{1}{2}))$. As in Figure 4, we obtain $\delta_{\min} = \delta^{\varrho}_{\neg\varnothing}$ and thus $\delta_{\min} \in I_1$. Note that here we have $\delta_{\neg\varnothing} > 0$.

On Figure 7, in light blue (resp. light red) is depicted the reachable set $\mathsf{Reach}(\mathcal{A}(\frac{3}{4})) = \mathsf{Reach}^*(\mathcal{A}(\frac{3}{4}))$ in location ℓ_1 (resp. ℓ_2). Dark colors represents here the stable zone for $\delta = \frac{3}{4}$. Note that here we have $\delta_{\min} \neq \delta^{\varrho}_{\neg\varnothing}$ and thus we obtain $\delta_{\min} \in I_2$.

On Figures 5 and 7, dotted lines represent the limits of the reachable state space for the value $\delta_{\max}$. Here we can see when the reachable space for location ℓ_2 is able to enter the Bad location. On Figure 7, this additional enlargement from $\delta_{\min}$ is equal to $\frac{1}{12}$ and thus we have $\delta_{\max} = \delta_{\min} + \frac{1}{12}$.

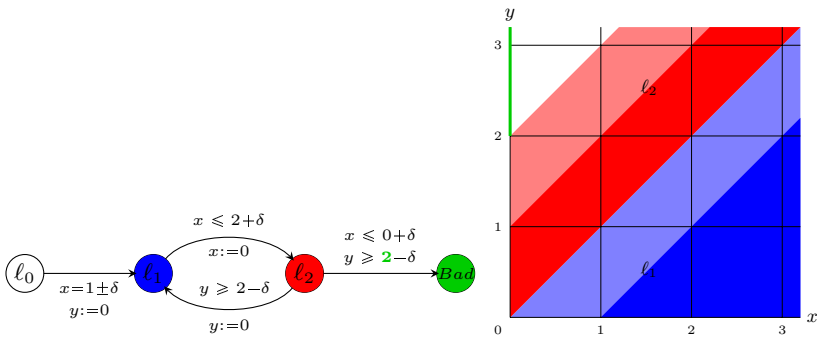

Fig. 4: $\mathcal{A}(\delta)$ is safe iff $\delta = 0$ ($\delta_{\min} = \delta^{\varrho}_{\neg\varnothing} = 0$)

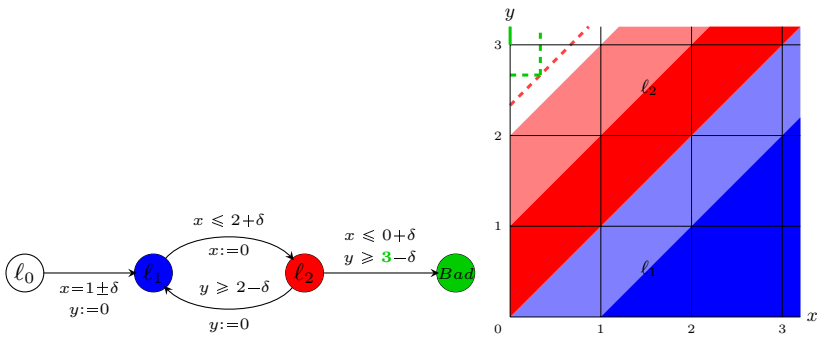

Fig. 5: $\mathcal{A}(\delta)$ is safe iff $\delta < \frac{1}{3}$ ($\delta_{\min} = \delta^{\varrho}_{\neg\varnothing} = 0$)

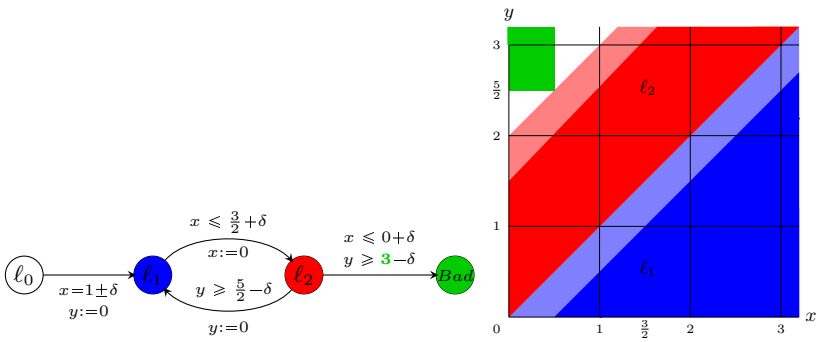

Fig. 6: $\mathcal{A}(\delta)$ is safe iff $\delta \leqslant \frac{1}{2}$ ($\delta_{\min} = \delta^{\varrho}_{\neg\varnothing} = \frac{1}{2}$)

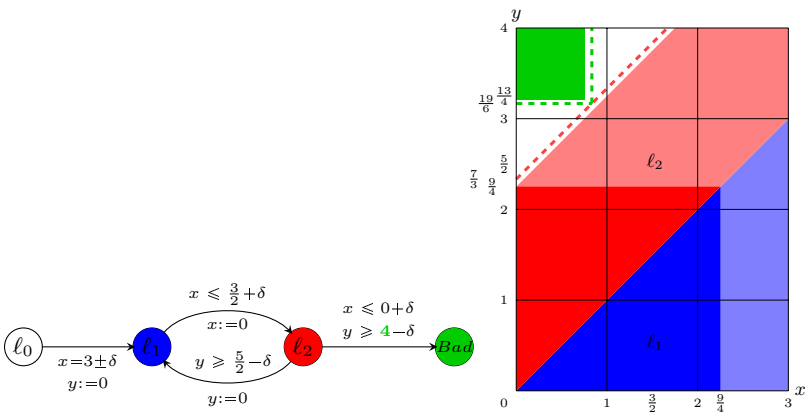

Fig. 7: $\mathcal{A}(\delta)$ is safe iff $\delta < \frac{5}{6}$ ($\delta_{\min} = \frac{3}{4}$ and $\delta^{\varrho}_{\neg\varnothing} = \frac{1}{2}$)