

Computing maximally-permissive strategies in acyclic timed automata^{*}

Emily Clement^{1,2}, Thierry Jéron¹, Nicolas Markey¹, and
David Mentré²

¹ IRISA, Inria & CNRS & Univ. Rennes, France
`firstname.lastname@inria.fr`

² Mitsubishi Electric R&D Centre Europe – Rennes, France
`initial-of-firstname.lastname@fr.mercede.mee.com`

Abstract. Timed automata are a convenient mathematical model for modelling and reasoning about real-time systems. While they provide a powerful way of representing timing aspects of such systems, timed automata assume arbitrary precision and zero-delay actions; in particular, a state might be declared reachable in a timed automaton, but impossible to reach in the physical system it models.

In this paper, we consider *permissive* strategies as a way to overcome this problem: such strategies propose *intervals of delays* instead of single delays, and aim at reaching a target state whichever delay actually takes place. We develop an algorithm for computing the optimal permissiveness (and an associated maximally-permissive strategy) in acyclic timed automata and games.

1 Introduction

Timed automata [AD94] are a powerful formalism for modelling and reasoning about real-time computer systems: they offer a convenient way of modelling timing conditions (not relying on discretization) while allowing for efficient verification algorithms; as a consequence, they have been widely studied by the formal-verification community, and have been applied to numerous industrial case studies thanks to advanced tools such as Uppaal [BDL⁺06], TChecker [HPT19] or Chronos [BDM⁺98].

One drawback of timed automata is that they are a mathematical model, assuming infinite precision in the measure of time; this does not correspond to physical devices such as computers. As a consequence, properties that are proven to hold on the model may fail to hold on any implementation. As a very simple example, consider two (or even infinitely-many) consecutive actions that have to be performed at the exact same time: while this would be possible in a mathematical model, this would not be possible on a physical device.

Several approaches have attempted to address such problems, depending on the property to be checked. When considering safety properties, timing imprecisions

^{*} This work was partially funded by ANR project Ticktac (ANR-18-CE40-0015).

may add new behaviours, which have to be taken into account in the safety check. In that setting, *guard enlargement* [Pur00, DDMR04] has been proposed as a way to model the fact that some timing conditions might be considered true even if they are (slightly) violated: the existence of an enlargement value for which the set of executions is safe is decidable. When dealing with reachability properties, timing imprecisions may prevent a run to be valid. A topological approach has been proposed, where a state is declared reachable only if there is a *tube of trajectories* reaching the target state [GHJ97]. Game-based approaches have also been proposed, where a state is said reachable if there is a strategy to reach this state when an opponent player is allowed to modify (up to a certain point) the values of the delays [BMS15, BFM15].

In this paper, we build on the approach of [BFM15], where the authors aim at computing *maximally-permissive* strategies for reaching a target state. While in classical timed automata, reachability is witnessed by a sequence of delays and transitions leading to a target state, here the aim is to propose *intervals* of delays, leaving it to an opponent player to decide which delay will indeed take place. Of course, the strategy has to be able to respond to any choice of the opponent, eventually reaching the target state.

We can then have several ways of measuring permissiveness of a strategy, the general idea being that larger intervals of delays are preferred. In [BFM15], each interval is associated with a penalty, which is the inverse of the length of the interval. Penalties are summed up along paths, and maximally-permissive strategies are those having minimal worst-case penalty. This favours both large intervals and short paths, but computing optimal strategies could only be achieved in the case of one-clock timed automata in [BFM15].

In the present paper, permissiveness of a strategy is defined as the size of the smallest interval proposed by that strategy. We develop an algorithm to compute the permissiveness of any (winning) configuration in acyclic timed automata and games, with any number of clocks. Consider for instance a scheduling problem, where a number of tasks have to be performed in a certain order within a given delay. Classical reachability algorithms would just say whether a given set of tasks are schedulable (in the mathematical model); this then requires launching some of the tasks at very precise dates, as the computed schedule need not be correct if delays are slightly modified. Instead, our algorithm could compute the permissiveness of the best schedule, thereby measuring the amount of imprecision that can be allowed, depending on the deadline by which all tasks have to be finished.

This paper is organized as follows: in Section 2, we introduce the necessary definitions, in particular of timed automata and permissiveness of strategies, and prove basic results. Section 3 is devoted to solving the case of *linear* timed automata, where all states have at most one outgoing transition, thereby focusing only on choices of delays. Section 4 extends this to acyclic timed automata and games.

By lack of space, most of the proofs could not be included in this version of the paper. They can be found in the long version [CJMM20] of this article.

2 Definitions

2.1 Piecewise-affine functions

A valuation for a set \mathcal{C} of variables is a mapping $v: \mathcal{C} \rightarrow \mathbb{R}_{\geq 0}$, assigning a nonnegative real value to each variable. We write $\mathbf{0}$ for the valuation defined as $\mathbf{0}(c) = 0$ for any $c \in \mathcal{C}$. We write $(\mathbb{R}_{\geq 0})^{\mathcal{C}}$ for the set of valuations for \mathcal{C} , which we identify with $(\mathbb{R}_{\geq 0})^n$ when \mathcal{C} has exactly n variables. We write $\overline{\mathbb{R}}$ for $\mathbb{R} \cup \{-\infty; +\infty\}$.

Definition 1. An n -dimensional affine function is a mapping $f: \mathbb{R}_{\geq 0}^n \rightarrow \overline{\mathbb{R}}$ s.t.

- either there exists a vector $(F_k)_{0 \leq k \leq n} \in \mathbb{R}^{n+1}$ such that $f(v) = F_0 + \sum_{1 \leq i \leq n} F_i \cdot v_i$;
- or $f(v) = -\infty$ for all $v \in \mathbb{R}_{\geq 0}^n$; in that case we can still write $f(v) = F_0 + \sum_{1 \leq i \leq n} F_i \cdot v_i$, by setting $F_0 = -\infty$ and $F_i = 0$ for all $1 \leq i \leq n$;
- or $f(v) = +\infty$ for all $v \in \mathbb{R}_{\geq 0}^n$; similarly, this corresponds to setting $F_0 = +\infty$ and $F_i = 0$ for all $1 \leq i \leq n$.

A linear function f is an affine function for which $f(\mathbf{0}) = 0$.

If $\Phi = (\varphi_k)_{1 \leq k \leq m}$ is a set of n -dimensional affine functions and $b = (b_k)_{1 \leq k \leq m}$ is a set of intervals, we write $\llbracket \Phi, b \rrbracket$ for the intersection $\bigcap_{1 \leq k \leq m} \varphi_k^{-1}(b_k)$. This defines a convex polyhedron of $\mathbb{R}_{\geq 0}^n$.

An n -dimensional piecewise-affine function is a mapping $f: \mathbb{R}_{\geq 0}^n \rightarrow \overline{\mathbb{R}}$ for which there exists a partition $S = (S_k)_{1 \leq k \leq m}$ of $\mathbb{R}_{\geq 0}^n$ into convex polyhedra, and a family $(f_k)_{1 \leq k \leq m}$ of affine functions such that for any $x \in \mathbb{R}_{\geq 0}^n$, writing k for the (unique) index in $[1; m]$ such that $x \in S_k$, it holds $f(x) = f^k(x)$.

2.2 Timed automata

Given a valuation v and a nonnegative real d , we denote with $v+d$ the valuation w such that $w(c) = v(c) + d$ for all $c \in \mathcal{C}$. For any subset $I \subseteq \mathbb{R}_{\geq 0}$, we write $v+I$ for the set of valuations $\{v+d \mid d \in I\}$. Given a valuation v and a subset $r \subseteq \mathcal{C}$, we write $v[r \rightarrow 0]$ for the valuation w such that $w(c) = 0$ if $c \in r$ and $w(c) = v(c)$ if $c \notin r$.

The set of linear constraints over \mathcal{C} is defined as $\mathcal{G}(\mathcal{C}) \ni g ::= c \sim n \mid g \wedge g$ where c ranges over \mathcal{C} , n ranges over \mathbb{N} , and $\sim \in \{<, \leq, =, \geq, >\}$. That a clock valuation v satisfies a clock constraint g , denoted $v \models g$ (and sometimes $v \in g$, seeing g as a convex polyhedron), is defined inductively as

$$v \models c \sim n \Leftrightarrow v(c) \sim n \quad v \models g_1 \wedge g_2 \Leftrightarrow v \models g_1 \text{ and } v \models g_2$$

For the rest of this paper, we fix a finite alphabet Σ .

Definition 2 ([AD94]). A timed automaton over Σ is a tuple $\mathcal{A} = (\mathcal{C}, L, T, I)$ where \mathcal{C} is a finite set of clocks, L is a finite set of states (or locations), and $T \subseteq L \times \mathcal{G}(\mathcal{C}) \times \Sigma \times 2^{\mathcal{C}} \times L$ is a finite set of transitions, and $I: S \rightarrow \mathcal{G}(\mathcal{C})$ define the invariant constraints in locations.

A *configuration* of a timed automaton is a pair (ℓ, v) where ℓ is a location of the automaton and v is a clock valuation such that $v \models I(\ell)$. The semantics of timed automata can be defined as an infinite-state labelled transition system whose states are the set of configurations, and whose transitions are of two kinds:

- *delay transitions* model time elapsing: no transitions of the timed automaton are taken, but the values of all clocks are augmented by the same value. For any configuration (ℓ, v) and any delay $d \in \mathbb{R}_{\geq 0}$, there is a transition $(\ell, v) \xrightarrow{d} (\ell, v + d)$, provided that $v + d \models I(\ell)$;
- *action transitions* represent the effect of taking a transition in the timed automaton. For any configuration (ℓ, v) and any transition $t = (\ell, g, a, r, \ell')$, if $v \models g$, then there is a transition $(\ell, v) \xrightarrow{a} (\ell', v[r \rightarrow 0])$, provided that $v[r \rightarrow 0] \models I(\ell')$.

We write $(\ell, v) \xrightarrow{d,a} (\ell', v')$ when there exists (ℓ'', v'') such that $(\ell, v) \xrightarrow{d} (\ell'', v'')$ and $(\ell'', v'') \xrightarrow{a} (\ell', v')$. A run of a timed automaton is a sequence of configurations $(\ell_i, v_i)_i$ such that there exists $d \in \mathbb{R}_{\geq 0}$ and $a \in \Sigma$ such that $(\ell_i, v_i) \xrightarrow{d,a} (\ell_{i+1}, v_{i+1})$ for all i . Even if it means adding a sink state and corresponding transitions, we assume that from any configuration, there always exists a transition $\xrightarrow{d,a}$ for some $d \in \mathbb{R}_{\geq 0}$ and some $a \in \Sigma$. This way, any finite run can be extended into an infinite run (in terms of its number of transitions). We also assume that, from any location ℓ and any action a , there is at most one transition from ℓ labelled with a .

One of the most basic problems concerning timed automata is that of reachability of a location: given a timed automaton \mathcal{A} , a source configuration (ℓ_0, v_0) (usually assuming $v_0 = \mathbf{0}$) and a target location ℓ_f , it amounts to deciding whether there exists a run from (ℓ_0, v_0) to some configuration (ℓ_f, v_f) in the infinite-state transition system defining the semantics of \mathcal{A} . This problem has been proven decidable (and PSPACE-complete) in the early 1990s [AD94], using *region equivalence*, which provides a finite-state automaton that is (time-abstracted) bisimilar to the original timed automaton.

2.3 Permissive strategies in timed automata

Solving reachability using the algorithm above, we can obtain a sequence of delays and transitions to be taken for reaching the target location. Playing this sequence of delays and transitions however requires infinite precision in order to meet all timing constraints, which might not be possible on physical devices.

In this paper, we address this problem by building on the setting studied in [BFM15]: in that setting, the delays that are played may be slightly perturbed, and it can be required to adapt the future delays (and possibly actions) so as to make sure that the target is indeed reached.

We encode the imprecisions using a game setting: the player proposes an interval of possible delays (together with the action to be played), and its opponent selects, in the proposed interval, the exact delay that will take place.

Formally, in our setting, a *move* from some configuration (ℓ, v) is a pair (I, a) , where $I \subseteq \mathbb{R}_{\geq 0}$ is a closed¹ interval, possibly right-unbounded, and $a \in \Sigma$, such that there is a transition (ℓ, g, a, r, ℓ') for which $v + I \subseteq g$ (i.e., for any valuation $w \in v + I$, it holds $w \models g$). We write $\text{moves}(\ell, v)$ for the set of moves from (ℓ, v) .

A *permissive strategy* is a function σ mapping finite runs $(\ell_i, v_i)_{0 \leq i \leq n}$ to moves in $\text{moves}(\ell_n, v_n)$. A run $\rho = (\ell_i, v_i)_i$ is *compatible* with a permissive strategy σ if, for any finite prefix $\pi = (\ell_i, v_i)_{0 \leq i \leq j}$ of ρ , $\sigma(\pi)$ is defined and, writing $\sigma(\pi) = (I, a)$, there exists $d \in I$ such that $(\ell_j, v_j) \xrightarrow{d, a} (\ell_{j+1}, v_{j+1})$. A permissive strategy σ is *winning* from a given configuration (ℓ_0, v_0) if any infinite run originating from (ℓ_0, v_0) that is compatible with σ is winning (which, in our setting, means that it visits the target location ℓ_f). Notice that classical strategies (which propose single delays instead of intervals of delays) are special cases of permissive strategies. It follows that, as soon as there is a path from some configuration (ℓ, v) to ℓ_f , there exists a winning permissive strategy from (ℓ, v) (possibly proposing punctual intervals). Such configurations are said *winning*, and the *winning zone* is the set of all winning configurations.

Our aim is to compute *maximally-permissive* winning strategies. In this work, we measure the permissiveness of a strategy σ in a configuration (ℓ, v) , denoted $\text{Perm}_\sigma(\ell, v)$, as the length of the smallest interval it may return. Formally:

Definition 3. *Let σ be a permissive strategy, and (ℓ, v) be a configuration of \mathcal{A} . The permissiveness of σ in (ℓ, v) , denoted with $\text{Perm}_\sigma(\ell, v)$, is defined as follows:*

- if σ is not winning from (ℓ, v) , the permissiveness of σ in (ℓ, v) is $-\infty$;
- otherwise, $\text{Perm}_\sigma(\ell, v) = \inf\{|I| \mid \exists \pi. \sigma(\pi) = (I, a) \text{ for some } a\}$.

The permissiveness of configuration (ℓ, v) is then defined as

$$\text{Perm}(\ell, v) = \sup_{\sigma} \text{Perm}_\sigma(\ell, v).$$

In this paper, we prove that Perm is a piecewise affine function, and develop an algorithm for computing that function. Intuitively, this corresponds to computing how much precision is needed in order to reach the target configuration.

Remark 4. Notice that our definition of permissiveness is similar in spirit with that of [BFM15]. However, in [BFM15], each move (I, a) was associated a penalty (namely $1/|I|$), and penalties are summed up along the execution. This tends to make the player favour shorter paths with possibly small intervals (hence demanding more accuracy when playing) over long paths with larger intervals. Our setting only aims at maximizing the size of the smallest interval to be played.

Our work can also be seen as a kind of quantitative extension of *tubes of trajectories* of [GHJ97]: permissiveness could be seen as the minimal width of such a tube. However, we are in a game-based setting, and (except in Section 3) the strategy could suggest to take different transitions if they allow for more permissiveness.

¹ We only consider closed intervals here to simplify the presentation.

Finally, and perhaps more importantly, our setting is quite close to that of [BMS15], but with a more quantitative focus: we aim at computing the optimal permissiveness for all winning configurations (with reachability objective), while only a global lower bound (of the form $1/m$ where m is doubly-exponential in the size of the input) is obtained in [BMS15]. Similar results to those of [BMS15] are obtained in [SBMR13, BGMRS19] for Büchi objectives; such extensions are part of our future work.

Remark 5. The term *permissive* strategy is sometimes used to refer to *non-deterministic*, returning the set of all moves that lead to winning configurations. In particular, Uppaal-Tiga [BCD⁺07] can compute maximally-permissive strategies in that sense. But this is only a local view of permissiveness, while our aim is to allow for high permissiveness all along the execution.

2.4 Iterative computation of permissiveness

Towards computing Perm, we define:

$$\begin{aligned} \mathcal{P}_i(\ell_f, v) &= +\infty && \text{for all valuations } v \text{ and all } i \geq 0; \\ \mathcal{P}_0(\ell, v) &= -\infty && \text{for all valuations } v, \text{ and for all } \ell \neq \ell_f; \\ \mathcal{P}_{i+1}(\ell, v) &= \begin{cases} \sup_{(I, a) \in \text{moves}(\ell, v)} \min(|I|, \inf\{\mathcal{P}_i(\ell', v') \mid \exists d \in I. (\ell, v) \xrightarrow{d, a} (\ell', v')\}) & \text{if } \text{moves}(\ell, v) \neq \emptyset \\ -\infty & \text{otherwise} \end{cases} \end{aligned}$$

In the rest of section, we prove some basic properties of this sequence of functions, and in particular its link with permissiveness. The next sections will be devoted to its computation on acyclic timed automata.

Our first two results are concerned with the evolution of the sequence with i . They are proved by straightforward inductions.

Lemma 6. *For any (ℓ, v) , the sequence $(\mathcal{P}_i(\ell, v))_{i \in \mathbb{N}}$ is nondecreasing.*

Lemma 7. *If the longest path from ℓ to ℓ_f has at most i transitions, then for any v and any $j \geq 0$, it holds $\mathcal{P}_{i+j}(\ell, v) = \mathcal{P}_i(\ell, v)$.*

The following lemma ties the link between the sequence (\mathcal{P}_i) and permissiveness:

Proposition 8. *For any $i \in \mathbb{N}$ and for any configuration (ℓ, v) , it holds:*

1. $\mathcal{P}_i(\ell, v) = -\infty$ if, and only if, there are no runs of length at most i from (ℓ, v) to ℓ_f ;
2. for any $p \in \mathbb{R}_{\geq 0}$, and any $i \in \mathbb{N}$, it holds $\mathcal{P}_i(\ell, v) > p$ if, and only if, there is a permissive strategy with permissiveness larger than p that is winning from (ℓ, v) within i steps.

Proof. We begin with the first equivalence, which we prove by induction on i . The result is trivial for $i = 0$. Now, assume that the result holds up to index i . There may be two reasons for having $\mathcal{P}_{i+1}(\ell, v) = -\infty$ for some (ℓ, v) : either $\text{moves}(\ell, v)$ is empty, or it is not empty and for any $(I, a) \in \text{moves}(\ell, v)$, it holds

$$\inf\{\mathcal{P}_i(\ell', v') \mid \exists d \in I. (\ell, v) \xrightarrow{d, a} (\ell', v')\} = -\infty.$$

This is true in particular when $I = \{d\}$ is punctual: for any (d, a) , the successor (ℓ', v') such that $(\ell, v) \xrightarrow{d, a} (\ell', v')$ is such that $\mathcal{P}_i(\ell', v') = -\infty$. From the induction hypothesis, there can be no path from those (ℓ', v') to ℓ_f with i steps or less. Hence there are no paths from (ℓ, v) to ℓ_f with at most $i + 1$ steps.

Conversely, if there are no paths having at most $i + 1$ steps from (ℓ, v) to ℓ_f , then either this is because $\text{moves}(\ell, v) = \emptyset$, or this is because all moves lead to a configuration from which there are no paths of length at most i to ℓ_f . By induction hypothesis, all successor configurations have infinite \mathcal{P}_i , hence also $\mathcal{P}_{i+1}(\ell, v) = -\infty$.

We now prove the second claim, still by induction. The base case is again trivial. Now, assume that the result holds up to some index i . We fix some $p \in \mathbb{R}_{\geq 0}$, and first consider a configuration (ℓ, v) with $\mathcal{P}_{i+1}(\ell, v) > p$. This entails that $\text{moves}(\ell, v)$ is non-empty, and that there is a move (I, a) with $|I| > p$ such that $\mathcal{P}_i(\ell', v') > p$ for all (ℓ', v') such that $(\ell, v) \xrightarrow{d, a} (\ell', v')$ with $d \in I$. Applying the induction hypothesis, there is an i -step winning strategy with permissiveness larger than p from each successor configuration (ℓ', v') , from which we can build an $i + 1$ -step winning strategy with permissiveness larger than p from (ℓ, v) .

Conversely, pick an $i + 1$ -step winning strategy σ_p from (ℓ, v) with permissiveness larger than p . Write $\sigma_p(\ell, v) = (I_0, a_0)$. Then for any $d \in I_0$, in the location (ℓ', v') such that $(\ell, v) \xrightarrow{d, a_0} (\ell', v')$, strategy σ_p is an i -step winning strategy with permissiveness larger than p , so that, following the induction hypothesis, $\mathcal{P}_i(\ell', v') > p$. It immediately follows that $\mathcal{P}_{i+1}(\ell, v) > p$. \square

Our next three results focus on properties of the functions \mathcal{P}_i . First, we identify zones on which \mathcal{P}_i is constant. This will be useful for proving correctness of our algorithm computing \mathcal{P}_i in the next section:

Lemma 9. *Let \mathcal{A} be a timed automaton, with maximal constant M . Let ℓ be a location, and $i \in \mathbb{N}$. Take two valuations v and v' such that, for any clock c , we have either $v(c) = v'(c)$, or $v(c) > M$ and $v'(c) > M$. Then $\mathcal{P}_i(\ell, v) = \mathcal{P}_i(\ell, v')$.*

Next we prove that the functions \mathcal{P}_i are 2-Lipschitz continuous (on the zone where they take finite values):

Proposition 10. *For any integer $i \in \mathbb{N}$ and any location ℓ , the function $\tau_\ell: v \mapsto \mathcal{P}_i(\ell, v)$ is 2-Lipschitz on the set of valuations where it takes finite values.*

Finally, the following lemma shows the (rather obvious) fact that $\mathcal{P}_i(\ell, v+t) \leq \mathcal{P}_i(\ell, v)$. A consequence of this property is that, for any non-resetting transition, the optimal choice for the opponent is the largest delay in the interval proposed by

the player². This corresponds to the intuition that by playing later, the opponent will force the player to react faster at the next step. As Example 1 below shows, this is not the case in general: in that example, from $(\ell_0, \langle x = 0; y = 0 \rangle)$, if the player proposes interval $[1/4; 1]$, the optimal choice for the opponent is $d = 1/4$.

Lemma 11. *Let (ℓ, v) be a configuration, $t \in \mathbb{R}_{\geq 0}$ such that $(\ell, v + t)$ is a configuration of the automaton, and $i \in \mathbb{N}$. Then $\mathcal{P}_i(\ell, v) - t \leq \mathcal{P}_i(\ell, v + t) \leq \mathcal{P}_i(\ell, v)$.*

Example 1. Consider the automaton of Fig. 1. We compute the optimal permissiveness (and corresponding strategies) for this small example. First, $\mathcal{P}_i(\ell_f, v) = +\infty$ for all i , and $\mathcal{P}_0(\ell_0, v) = \mathcal{P}_0(\ell_1, v) = -\infty$.




Fig. 1. A linear timed automaton and its permissiveness at ℓ_0 and ℓ_1

We first focus on ℓ_1 , with some valuation v : obviously, if $v(x) > 2$ or $v(y) > 1$, the set $\text{moves}(\ell_1, v)$ is empty, and $\mathcal{P}_1(\ell_1, v) = -\infty$ in that case; similarly if $v(y) > v(x)$. Since $\mathcal{P}_0(\ell_f, v)$ does not depend on v , the optimal move for the player is the largest possible interval satisfying the guard:

- if $v(x) \leq 1$ and $v(y) \leq 1$ (and $v(x) \leq v(y)$), the optimal interval of delays is $[1 - v(x); 1 - v(y)]$, whose length is $v(x) - v(y)$;
- if $v(y) \leq 1 \leq v(x) \leq 2$, the transition is immediately available, so that the lower bound of the interval will be 0. For the upper bound, there are two cases:
 - if $v(y) \geq v(x) - 1$, the optimal interval is $[0; 1 - v(y)]$;
 - if $v(y) \leq v(x) - 1$, the optimal interval is $[0; 2 - v(x)]$.

² This also holds for any transition in a one-clock timed automaton (because in case the clock is reset, the new valuation does not depend on the delay chosen by the opponent).

This defines the permissiveness for ℓ_1 .

We now look at ℓ_0 : first, $\mathcal{P}_1(\ell_0, v) = -\infty$ for all v , and only configurations (ℓ_0, v) where $v(x) \leq 1$ and $v(y) \leq 1$ are winning, so that $\mathcal{P}_2(\ell_0, v) = -\infty$ as soon as $v(x) > 1$ or $v(y) > 1$. Fix a valuation v for which $v(x) \leq 1$ and $v(y) \leq 1$. We have to find the interval $I = [\alpha, \beta]$ such that $v(x) + \beta \leq 1$ and $v(y) + \beta \leq 1$, and for which $\min\{\beta - \alpha, \inf_{\gamma \in [\alpha, \beta]} \mathcal{P}_1(\ell_1, (v + \gamma)[y \rightarrow 0])\}$ is maximized. Noticing that $(v + \gamma)[y \rightarrow 0]$ is the valuation $(x \mapsto v(x) + \gamma; y \mapsto 0)$, and that $\mathcal{P}_1(\ell_1, w) = w(x)$ for any w satisfying $w(x) \in [0; 1]$ and $w(y) = 0$, we have to maximize $\min\{\beta - \alpha, \inf_{\gamma \in [\alpha, \beta]} v(x) + \gamma\}$ over the domain defined by $0 \leq \alpha \leq \beta \leq \min(1 - v(x); 1 - v(y))$. Obviously, $\inf_{\gamma \in [\alpha, \beta]} v(x) + \gamma = v(x) + \alpha$, so we have to maximize $\min\{\beta - \alpha, v(x) + \alpha\}$ on the set $\{(\alpha, \beta) \mid 0 \leq \alpha \leq \beta \leq \min(1 - v(x); 1 - v(y))\}$.

We consider two cases:

- if $v(y) \leq v(x)$: clearly, it is optimal to maximize β , so we let $\beta = 1 - v(x)$. Hence we have to maximize $\min\{1 - (v(x) + \alpha), v(x) + \alpha\}$ over $0 \leq \alpha \leq 1 - v(x)$. Again, there are two cases, depending on whether $v(x)$ is larger or smaller than $1/2$; in the former case, $\min\{1 - (v(x) + \alpha), v(x) + \alpha\} = 1 - v(x) - \alpha$ when α ranges over $[0; 1 - v(x)]$; it is maximized for $\alpha = 0$, and we get $\mathcal{P}_2(\ell_0, v) = 1 - v(x)$. If $v(x) \leq 1/2$, the maximal value is reached when $\alpha = 1/2 - v(x)$, and $\mathcal{P}_2(\ell_0, v) = 1/2$.
- if $v(y) \geq v(x)$: then it is optimal to let $\beta = 1 - v(y)$. Again there are two cases for maximizing $\min\{1 - v(y) - \alpha, v(x) + \alpha\}$: if $1 - v(y) \leq v(x)$, then $\alpha = 0$ is optimal, and $\mathcal{P}_2(\ell_0, v) = 1 - v(y)$; otherwise, $\alpha = (1 - v(x) - v(y))/2$ is optimal, and $\mathcal{P}_2(\ell_0, v) = (1 - v(y) + v(x))/2$.

We end up with the diagram represented on the left of Fig. 1 (where for the sake of readability we write x and y in place of $v(x)$ and $v(y)$).

Our aim in the rest of this paper is to compute the sequence of functions \mathcal{P}_i , and to evaluate the complexity of this computation. Following Lemma 7, this will provide us with an algorithm for computing permissiveness in acyclic timed automata.

3 Computing optimal strategies in linear timed automata

In this section, we consider the simpler case of linear timed automata, where each location has at most one successor.

3.1 Optimal strategy for the opponent

We begin with focusing on the optimal choice of the opponent: given a configuration (ℓ, v) and an interval I of delays proposed by the player (there is a single outgoing transition, so the action to be played is fixed), what is the best delay that the opponent will choose so as to minimize the permissiveness of the resulting configuration?

As we already mentioned, Lemma 11 answers this question for non-resetting transitions: for such transitions, the best option for the opponent is to choose the maximal delay in the interval proposed by the player. On the other hand, Example 1 provides a situation where the opponent prefers to play as early as possible.

It turns out that, for linear timed automata, the optimal choice of the opponent is always one of these two extremal choices. This property will be a corollary of the following lemma, stating concavity of the permissiveness function in linear timed automata:

Proposition 12. *Let $i \in \mathbb{N}$. Let ℓ be a location of a linear timed automaton, let v_1 and v_2 be two clock valuations such that $\mathcal{P}_i(\ell, v_1)$ and $\mathcal{P}_i(\ell, v_2)$ are finite. Let $\lambda \in [0; 1]$, and $v_\lambda = \lambda \cdot v_1 + (1 - \lambda) \cdot v_2$. Then*

$$\mathcal{P}_i(\ell, v_\lambda) \geq \lambda \cdot \mathcal{P}_i(\ell, v_1) + (1 - \lambda) \cdot \mathcal{P}_i(\ell, v_2).$$

The aim of the opponent being to select the valuation in $V = \{v + \delta[r \rightarrow 0] \mid 0 \leq \delta \leq d\}$ that minimizes the permissiveness. Writing $v_1 = v[r \rightarrow 0]$ and $v_2 = v + d[r \rightarrow 0]$, we have $V = \{\lambda v_1 + (1 - \lambda)v_2 \mid 0 \leq \lambda \leq 1\}$. Proposition 12 entails that the permissiveness is minimized either in v_1 or in v_2 . This corresponds to our claim that the best choice for the opponent always is to select one of the bounds of the interval proposed by the player.

Corollary 13. *Let ℓ be a location of a linear timed automaton, v and v' be two clock valuations, $\lambda \in [0; 1]$, and $v_\lambda = \lambda \cdot v + (1 - \lambda) \cdot v'$. Then for all i :*

$$\mathcal{P}_i(\ell, v_\lambda) \geq \min\{\mathcal{P}_i(\ell, v), \mathcal{P}_i(\ell, v')\}.$$

In particular, for any valuation v , any bounded interval $[\alpha, \beta]$, and any transition $\ell \xrightarrow{g, a, r} \ell'$:

$$\inf\{\mathcal{P}_i(\ell', v') \mid \exists d \in [\alpha, \beta]. (\ell, v) \xrightarrow{d, a} (\ell', v')\} = \min\{\mathcal{P}_i(\ell', v'_\alpha), \mathcal{P}_i(\ell', v'_\beta)\}$$

where $(\ell, v) \xrightarrow{\alpha, a} (\ell', v'_\alpha)$ and $(\ell, v) \xrightarrow{\beta, a} (\ell', v'_\beta)$.

3.2 Computing the most-permissive strategy

Now that we have a better understanding of the optimal strategy of the opponent, we can compute the most-permissive strategy of the player for reaching the target location ℓ_f . We prove that for all i , \mathcal{P}_i is in fact a piecewise-affine function that can be computed in doubly-exponential time.

First notice that, following Lemma 7, for any location ℓ of a linear timed automaton with n locations, the sequence of functions $(\mathcal{P}_i)_i$ converges in at most n steps.

Theorem 14. *The permissiveness function for a linear timed automaton with d locations and n clocks is a piecewise-affine function. It can be computed in time $O((n + 1)^{8^d})$.*

The following technical lemma will be the central tool in the computation of \mathcal{P}_i :

Lemma 15. *Let $m_\alpha \leq M_\alpha$ and $m_\beta \leq M_\beta$, and $D = \{(\alpha, \beta) \in \mathbb{R}_{\geq 0}^2 \mid m_\alpha \leq \alpha \leq M_\alpha, m_\beta \leq \beta \leq M_\beta, \alpha \leq \beta\}$. Let $f: \alpha \mapsto a\alpha + b$ and $g: \beta \mapsto c\beta + d$ be two 1-dimensional affine functions, and $\mu: (\alpha, \beta) \mapsto \min\{\beta - \alpha, f(\alpha), g(\beta)\}$. Then the maximal value that μ may take over D is of one of the following five forms: $M_\beta - m_\alpha$, $\lambda \cdot f(\nu)$, $\lambda \cdot g(\mu)$, $\frac{ad-bc}{a-c}$ and $\frac{ad-bc}{(a+1)(1-c)-1}$, with $\lambda \in \{1, \frac{1}{1-c}, \frac{1}{a+1}\}$ and $\nu \in \{m_\alpha, M_\alpha, m_\beta, M_\beta\}$. This value can be computed by checking inequalities between expressions of the same forms.*

The following lemma corresponds to one step of our inductive computation of \mathcal{P}_i :

Lemma 16. *Let \mathcal{A} be a linear timed automaton with n clocks. Let (ℓ, g, a, z, ℓ') be a transition of \mathcal{A} , and assume that $v \mapsto \mathcal{P}_{i-1}(\ell', v)$ is piecewise affine, with m cells. Then $v \mapsto \mathcal{P}_i(\ell, v)$ is piecewise affine. It can be computed in time $O(m^4 \cdot (m+n)^4)$. It can be defined using a polyhedral partition of size $O(m^4 \cdot (m+n)^4)$, and with coefficients polynomial in those of \mathcal{P}_{i-1} .*

Proof. We assume that $\mathcal{P}_{i-1}(\ell', v)$ is not constantly $-\infty$ (if it were the case, then also $\mathcal{P}_i(\ell, v) = -\infty$ for all v). Similarly, we assume that $\text{moves}(\ell, v)$ is non-empty for some v . Since $v \mapsto \mathcal{P}_{i-1}(\ell', v)$ is piecewise-affine: we can then fix a polyhedral partition $[\Phi, P]$ and, for each cell h in this partition, an affine functions f_h , such that $\mathcal{P}_{i-1}(\ell', v) = f_h(v)$ for the only cell h containing v .

Our procedure for computing \mathcal{P}_i in ℓ consists in listing the possible pairs of cells defining \mathcal{P}_{i-1} in ℓ' where the left- and right-bounds of the interval to be proposed lie. For each pair (h_α, h_β) of such cells, we perform the following three steps (illustrated on Figure 2):

- characterize the set $S_{(h_\alpha, h_\beta)}$ of all valuations from which those cells can be reached by taking the transition from ℓ to ℓ' . We compute this polyhedron using quantifier elimination;




Fig. 2. Three steps of our procedure: $S_{(h_\alpha, h_\beta)}$; then compute expressions for I_α^v and I_β^v (notice that we had to refine $S_{(h_\alpha, h_\beta)}$, because the expression for I_β^v would be different for the lower part of $S_{(h_\alpha, h_\beta)}$ since it ends of a different facet of h_β); finally select best values for α and β .

- compute the ranges for α and β that can be played in order to indeed end up respectively in h_α and h_β . These are intervals I_α and I_β , whose bounds are expressed as functions of v . Computing these bounds may require refining the polyhedron obtained at the previous step into several subpolyhedra, in order to express them as affine functions of $v \in S_{(h_\alpha, h_\beta)}$;
- for each subpolyhedron, compute the optimal values for α and β : following Corollary 13, this amounts to find values for $\alpha \in I_\alpha$ and $\beta \in I_\beta$ that maximize the following function:

$$\mu: (\alpha, \beta) \mapsto \min\{\beta - \alpha; \mathcal{P}_{i-1}(\ell', (v + \alpha)[z \rightarrow 0]); \mathcal{P}_{i-1}(\ell', (v + \beta)[z \rightarrow 0])\}.$$

This is performed by applying our technical Lemma 15; it may again require another refinement of the subpolyhedra, and returns an affine function for each subpolyhedron.

For each pair (h_α, h_β) , we end up with a (partial) piecewise-affine function, defined on $S_{(h_\alpha, h_\beta)}$, returning the optimal permissiveness that can be obtained if playing interval $[\alpha, \beta]$ such that taking the transition to ℓ' after delay α (resp. β) leads to h_α (resp. h_β). Our final step to compute \mathcal{P}_i in ℓ consists in taking the maximum of all these partial functions on their (possibly overlapping) domains; this may introduce on more refinement of our polyhedron.

Notice that all these computations are performed symbolically w.r.t v : we manipulate affine functions of v , with conditions on v for our computations to be valid. \square

Assuming that $v \mapsto \mathcal{P}_{i-1}(\ell', v)$ has m cells, computing $v \mapsto \mathcal{P}_i(\ell, v)$ takes time $O(m^4 \cdot (m + n)^4)$, where n is the number of clocks, and this function has $O(m^4 \cdot (m + n)^4)$ many cells.

It follows that, for a linear timed automaton having d locations, we obtain the permissiveness function in the initial state as a piecewise-affine function in time $O((n + 1)^{8d})$, which proves Theorem 14.

4 Extension to acyclic timed automata and games

4.1 Adding branching

We extend the previous study to the case of acyclic timed automata (with branching). In that case, we can still apply our inductive approach, with a few changes: at each step, we would compute the optimal move of the player for each single action, and then select the optimal action by “superimposing” the resulting permissiveness functions and selecting the action that maximizes permissiveness. This however breaks the result of Prop. 12: the maximum of two concave functions need not be concave. Example 2, derived from Example 1, displays an example where the permissiveness function is not concave.

Example 2. Consider the automaton of Fig 5. The transition from ℓ_0 to ℓ_f has the same constraint as that from ℓ_1 to ℓ_f ; hence the permissiveness offered by




Fig. 3. Automaton of Fig. 1 where the guard on the first transition has been slightly extended




Fig. 4. Permissiveness function of the automaton of Fig. 3 in ℓ_0

that action is the same as the one from ℓ_1 , which we already computed. Hence the global permissiveness from ℓ_0 is the (pointwise) maximal of the two piecewise-affine functions displayed on Fig. 1, which is depicted on Fig. 5. On this diagram, the blue area corresponds to points from where it is better (or only possible) to go via ℓ_1 , while the red area corresponds to valuations from where it is better (or only possible) to take the bottom transition.





Fig. 5. A timed automaton and its (non-concave) permissiveness function in ℓ_0

We prove by induction that the permissiveness functions still are piecewise-affine in that setting. Hence all four steps of our proof of Lemma 16 still apply, with some adaptations. For each location ℓ , for each transition t from ℓ to some ℓ' , the procedure now is as follows:

- for the first step, we again consider two cells h_α and h_β in the partition defining $\mathcal{P}_{i-1}(\ell')$, together with a set H of cells that will be visited between h_α and h_β . Again applying Fourier-Motzkin, we get a polyhedron $S_{(h_\alpha, h_\beta, H)}$ of valuations from which those cells can indeed be visited;
- the computation of the intervals I_α^v and I_β^v is unchanged;
- for each cell $h \in H$, we can compute the values d_h^{in} and d_h^{out} for which $(v + d_h^{\text{in}})[z \rightarrow 0]$ enters h and $(v + d_h^{\text{out}})[z \rightarrow 0]$ leaves h (notice that this may require further refinement of the polyhedron being considered). Since \mathcal{P}_{i-1} is affine on cell h , it reaches its maximum on this cell either at $(v + d_h^{\text{in}})[z \rightarrow 0]$ or at $(v + d_h^{\text{out}})[z \rightarrow 0]$. The function we need to maximize now looks like

$$\mu' : (\alpha, \beta) \mapsto \min(\{\beta - \alpha, \mathcal{P}_{i-1}(\ell', (v + \alpha)[z \rightarrow 0]), \mathcal{P}_{i-1}(\ell', (v + \beta)[z \rightarrow 0])\} \cup \{\mathcal{P}_{i-1}(\ell', (v + d_h^{\text{in}})[z \rightarrow 0]), \mathcal{P}_{i-1}(\ell', (v + d_h^{\text{out}})[z \rightarrow 0]) \mid h \in H\}).$$

Now, we notice that all values in the second set are constant, not depending on α and β . We can thus still apply Lemma 15 in order to maximize $\mu(\alpha, \beta)$, and then take the above constants into account (which may again refine the polyhedra).

- the above three steps have to be performed for all outgoing transitions from the location ℓ being considered. The last step still consists in selecting the maximum of all the resulting functions.

The complexity of our procedure is much higher than that of linear automata: because we consider sets of cells already at the first step, we have $O(2^m)$ sets to consider. Assuming that \mathcal{P}_{i-1} is made of m cells, we may end up with \mathcal{P}_i having more than 2^m cells. Since we have to repeat this procedure up to $|T|$ times, so that the time complexity is in $O(|T|2)$ (where $^n a$ is tetration). Hence our procedure is non-elementary in the worst case. In the end:

Theorem 17. *The permissiveness function for acyclic timed automata is piecewise affine. It can be computed in non-elementary time.*

4.2 Adding uncontrollable states

We finally extend our approach to (acyclic) two-player turn-based timed games.

This setting is easily seen to preserve piecewise-affineness of the permissiveness function. Indeed, in order to compute \mathcal{P}_i in a location ℓ belonging to the opponent, it suffices to first compute the functions $\mathcal{P}_i^{\ell \rightarrow \ell'}$ for all outgoing transitions from ℓ to some ℓ' ; this follows the same procedure as above, and results in a piecewise-affine function, assuming (inductively) that \mathcal{P}_{i-1} is piecewise affine. We then

compute the (still piecewise-affine) minimum $\mathcal{M}_i(\ell, v)$ of all those functions, and finally

$$\mathcal{P}_i(\ell, v) = \min_{\substack{d \\ \text{s.t.} \\ v+d \models \text{Inv}(\ell)}} \mathcal{M}_i(\ell, v + d)$$

which is easily computed and remains piecewise-affine. The computation for locations that belong to the player is similar as in the case of plain timed automata. It follows:

Theorem 18. *The permissiveness function for acyclic turn-based timed games is piecewise affine, and can be computed in non-elementary time.*

5 Conclusions and perspectives

In this paper, we addressed the problem of measuring the amount of precision needed in a timed automaton to reach a given target location. We built on the formalism of permissive strategies defined in [BFM15], and developed an algorithm for computing the optimal permissiveness in acyclic timed automata and games.

There are several directions in which we will extend this work: as a first task, we will have a closer look at the complexity of our procedure, trying to either find examples where the number of cells indeed grows exponentially (for linear timed automata) or exponentially at each step (for acyclic timed automata). A natural continuation of our work consists in tackling cycles. We were unable to prove our intuition that there is no reason for the player to iterate a cycle. Following [BGMRS19], we might first consider fixing a timed automaton made of a single cycle, study how permissiveness evolves along one run in this cycle, and compute the optimal permissiveness for being able to take a cycle forever. Exploiting 2-Lipschitz continuity of the permissiveness function, we could also develop approximating techniques, both for making our computations more efficient in the acyclic case and to handle cycles. Finally, other interesting directions include extending our approach to linear hybrid automata, or considering a stochastic opponent, thereby modelling the fact that perturbations need not always be antagonist.

References

- [AD94] Rajeev Alur and David L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, April 1994.
- [BCD⁺07] Gerd Behrmann, Agnès Cougnard, Alexandre David, Emmanuel Fleury, Kim Guldstrand Larsen, and Didier Lime. UPPAAL-Tiga: Time for playing games! In Werner Damm and Holger Hermanns, editors, *Proceedings of the 19th International Conference on Computer Aided Verification (CAV'07)*, volume 4590 of *Lecture Notes in Computer Science*, pages 121–125. Springer-Verlag, July 2007.
- [BDL⁺06] Gerd Behrmann, Alexandre David, Kim Guldstrand Larsen, John Håkansson, Paul Pettersson, Wang Yi, and Martijn Hendriks. Uppaal 4.0. In *Proceedings of the 3rd International Conference on Quantitative Evaluation of Systems (QEST'06)*, pages 125–126. IEEE Comp. Soc. Press, September 2006.

- [BDM⁺98] Marius Bozga, Conrado Daws, Oded Maler, Alfredo Olivero, Stavros Tripakis, and Sergio Yovine. Kronos: A model-checking tool for real-time systems. In Alan J. Hu and Moshe Y. Vardi, editors, *Proceedings of the 10th International Conference on Computer Aided Verification (CAV'98)*, volume 1427 of *Lecture Notes in Computer Science*, pages 546–550. Springer-Verlag, June-July 1998.
- [BFM15] Patricia Bouyer, Erwin Fang, and Nicolas Markey. Permissive strategies in timed automata and games. In Gudmund Grov and Andrew Ireland, editors, *Proceedings of the 15th International Workshop on Automated Verification of Critical Systems (AVOCS'15)*, volume 72 of *Electronic Communications of the EASST*. European Association of Software Science and Technology, September 2015.
- [BGMRS19] Damien Busatto-Gaston, Benjamin Monmege, Pierre-Alain Reynier, and Ocan Sankur. Robust controller synthesis in timed büchi automata: A symbolic approach. In Işıl Dillig and Serdar Taşiran, editors, *Proceedings of the 31st International Conference on Computer Aided Verification (CAV'19)*, volume 11561 of *Lecture Notes in Computer Science*, pages 572–590. Springer-Verlag, July 2019.
- [BMS15] Patricia Bouyer, Nicolas Markey, and Ocan Sankur. Robust reachability in timed automata and games: A game-based approach. *Theoretical Computer Science*, 563:43–74, January 2015.
- [CJMM20] Emily Clement, Thierry Jéron, Nicolas Markey, and David Mentré. Computing maximally-permissive strategies in acyclic timed automata. Technical Report 2007.01815, arXiv, 2020.
- [DDMR04] Martin De Wulf, Laurent Doyen, Nicolas Markey, and Jean-François Raskin. Robustness and implementability of timed automata. In Yassine Lakhnech and Sergio Yovine, editors, *Proceedings of the Joint International Conferences on Formal Modelling and Analysis of Timed Systems (FORMATS'04) and Formal Techniques in Real-Time and Fault-Tolerant Systems (FTRTFT'04)*, volume 3253 of *Lecture Notes in Computer Science*, pages 118–133. Springer-Verlag, September 2004.
- [GHJ97] Vineet Gupta, Thomas A. Henzinger, and Radha Jagadeesan. Robust timed automata. In Oded Maler, editor, *Proceedings of the 1997 International Workshop on Hybrid and Real-Time Systems (HART'97)*, volume 1201 of *Lecture Notes in Computer Science*, pages 331–345. Springer-Verlag, March 1997.
- [HPT19] Frédéric Herbreteau, Gérald Point, and Thanh-Tung Tran. TChecker, an open-source model-checker for timed systems, 2019. Available at <https://github.com/fredher/tchecker>.
- [Pur00] Anuj Puri. Dynamical properties of timed systems. *Discrete Event Dynamic Systems*, 10(1-2):87–113, January 2000.
- [SBMR13] Ocan Sankur, Patricia Bouyer, Nicolas Markey, and Pierre-Alain Reynier. Robust controller synthesis in timed automata. In Pedro R. D'Argenio and Hernán C. Melgratt, editors, *Proceedings of the 24th International Conference on Concurrency Theory (CONCUR'13)*, volume 8052 of *Lecture Notes in Computer Science*, pages 546–560. Springer-Verlag, August 2013.

A Proofs of properties and lemmas about permissiveness

B Proofs of Section 2

This section is devoted to the proofs of the lemma of Section 2.

Lemma 6. *For any (ℓ, v) , the sequence $(\mathcal{P}_i(\ell, v))_{i \in \mathbb{N}}$ is nondecreasing.*

Proof. The proof is by induction. For any configuration (ℓ, v) , either $\mathcal{P}_0(\ell, v) = \mathcal{P}_1(\ell, v) = +\infty$, or $\mathcal{P}_0(\ell, v) = -\infty$. In both cases, $\mathcal{P}_0(\ell, v) \leq \mathcal{P}_1(\ell, v)$.

Then, assuming $\mathcal{P}_i(\ell, v) \leq \mathcal{P}_{i+1}(\ell, v)$ for all (ℓ, v) , we directly get the same property at step $i + 1$. The result follows. \square

Lemma 7. *If the longest path from ℓ to ℓ_f has at most i transitions, then for any v and any $j \geq 0$, it holds $\mathcal{P}_{i+j}(\ell, v) = \mathcal{P}_i(\ell, v)$.*

Proof. By induction on i : for $i = 0$, only ℓ_f satisfies the condition, and the result holds by definition of \mathcal{P}_i for ℓ_f .

Now, assume that the result holds for some index i , and consider a location ℓ such that the longest path to ℓ_f has at most $i + 1$ transitions. Then any successor location ℓ' of ℓ has longest path of length at most i , hence $\mathcal{P}_i(\ell', v) = \mathcal{P}_{i+1}(\ell', v)$. It immediately follows that $\mathcal{P}_{i+1}(\ell, v) = \mathcal{P}_i(\ell, v)$ for any v . \square

Lemma 9. *Let \mathcal{A} be a timed automaton, with maximal constant M . Let ℓ be a location, and $i \in \mathbb{N}$. Take two valuations v and v' such that, for any clock c , we have either $v(c) = v'(c)$, or $v(c) > M$ and $v'(c) > M$. Then $\mathcal{P}_i(\ell, v) = \mathcal{P}_i(\ell, v')$.*

Proof. The hypotheses ensure that any action- and delay transition performed from (ℓ, v) can be performed from (ℓ, v') , and the resulting configurations still satisfy the conditions of the lemma. The result follows by induction. \square

Lemma 11. *Let (ℓ, v) be a configuration, $t \in \mathbb{R}_{\geq 0}$ such that $(\ell, v + t)$ is a configuration of the automaton, and $i \in \mathbb{N}$. Then $\overline{\mathcal{P}}_i(\ell, v) - t \leq \mathcal{P}_i(\ell, v + t) \leq \mathcal{P}_i(\ell, v)$.*

Proof. For any move (I, a) that is available from $(\ell, v + t)$, the move $(I + t, a)$ is available from (ℓ, v) . Moreover, the set of valuations on which \mathcal{P}_{i-1} is minimized is the same in both cases, namely $\{(v + t[z \rightarrow 0]) \mid d \in I\}$. It follows that $\mathcal{P}_i(\ell, v + t) \leq \mathcal{P}_i(\ell, v)$.

Conversely, for any move (I, a) available from (ℓ, v) with $|I| \geq t$ (if any), the move $((I - t) \cap \mathbb{R}_{\geq 0}, a)$ is a valid (non-empty) move from $(\ell, v + t)$. The second inequality follows. \square

Finally, to prove proposition 10, we use the following lemmas:

Lemma 19. *Let v and v' be two clock valuations. Write $\eta = \|v' - v\|_\infty$. If $([\alpha, \beta], a) \in \text{moves}(\ell, v)$ with $\beta - \alpha \geq 2\eta$, then $([\alpha + \eta, \beta - \eta], a) \in \text{moves}(\ell, v')$.*

Proof. For any clock c , $0 \leq v(c) + \alpha = v'(c) + \alpha + (v(c) - v'(c)) \leq v'(c) + \alpha + \eta$. Similarly, $v(c) + \beta = v'(c) + \beta + (v(c) - v'(c)) \geq v'(c) + \beta - \eta$. Then for any interval J , if $v(c) + [\alpha, \beta] \subseteq J$, and also $v'(c) + [\alpha + \eta, \beta - \eta] \subseteq J$. It follows that for any guard g , if $v + [\alpha, \beta] \subseteq g$ with $\beta - \alpha \geq 2\|v' - v\|_\infty$, then $v' + [\alpha + \|v' - v\|_\infty, \beta - \|v' - v\|_\infty] \subseteq g$. \square

Corollary 20. *For any integer $i \in \mathbb{N}$ and any location ℓ , the function $\nu_\ell: v \mapsto \sup_{(I,a) \in \text{moves}(\ell,v)} |I|$ is 2-Lipschitz continuous on the set $\{v \mid \text{moves}(\ell,v) \neq \emptyset\}$.*

Proof. We first prove the result for the case where location ℓ has a single transition (ℓ, g, a, z, ℓ') . Take two valuations v and v' for which $\text{moves}(\ell, v)$ is non-empty. We prove that $\nu_\ell(v') - \nu_\ell(v) \geq -2\|v' - v\|_\infty$. By symmetry of the roles of v and v' , our result (for a single outgoing transition) follows.

First, if $\text{moves}(\ell, v)$ contains no intervals of size at least $2\|v' - v\|_\infty$, then obviously $\nu_\ell(v') - \nu_\ell(v) \geq -2\|v' - v\|_\infty$.

Now, assume that there exists $[\alpha, \beta] \in \text{moves}(\ell, v)$ such that $\beta - \alpha \geq 2\|v' - v\|_\infty$. By Lemma 19, for any such interval, the interval $[\alpha + \|v' - v\|_\infty, \beta - \|v' - v\|_\infty] \in \text{moves}(\ell, v')$.

Fix $\varepsilon > 0$, and take $I = [\alpha, \beta] \in \text{moves}(\ell, v)$ such that $|I| \geq \nu_\ell(v) - \varepsilon$. Since $[\alpha + \|v' - v\|_\infty, \beta - \|v' - v\|_\infty] \in \text{moves}(\ell, v')$, it follows $\nu_\ell(v') \geq \nu_\ell(v) - 2\|v' - v\|_\infty - \varepsilon$. Since this holds for any $\varepsilon > 0$, we get the announced inequality.

Now, in case there are several outgoing transitions, we have

$$\nu_\ell(v) = \sup_{(I,a) \in \text{moves}(\ell,v)} |I| = \max_{a \in \Sigma} \sup_{(I,a) \in \text{moves}(\ell,v)} |I|.$$

Hence ν_ℓ is the pointwise maximum of 2-Lipschitz continuous functions, hence it is 2-Lipschitz continuous. \square

Proposition 10. *For any integer $i \in \mathbb{N}$ and any location ℓ , the function $\tau_\ell: v \mapsto \mathcal{P}_i(\ell, v)$ is 2-Lipschitz on the set of valuations where it takes finite values.*

Proof. The proof is again by induction on i . The case of $i = 0$ is trivial. Corollary 20 proves the result for $i = 1$.

Now, assume that the result holds for some index $i - 1$. Take a location ℓ , and again first assume that ℓ has a single outgoing transition (ℓ, g, a, z, ℓ') . As in the previous proof, the result for the general case directly follows.

Pick two valuations v and v' such that $\mathcal{P}_i(\ell, v)$ and $\mathcal{P}_i(\ell, v')$ are finite. In particular, $\text{moves}(\ell, v)$ and $\text{moves}(\ell, v')$ are non-empty. We follow the same approach as in the proof of Corollary 20, proving that $\tau_\ell(v') - \tau_\ell(v) \geq -2\|v' - v\|_\infty$. By symmetry, our result follows.

Again, in case $\text{moves}(\ell, v)$ contains no intervals of size larger than or equal to $\|v' - v\|_\infty$, the result is immediate. Otherwise, fix $\varepsilon > 0$, and take an interval $I = [\alpha, \beta]$ such that

$$\min(|I|, \inf_{d \in I} (\mathcal{P}_{i-1}(\ell', (v+d)[z \rightarrow 0]))) \geq \tau_\ell(v) - \varepsilon.$$

Then $|I| \geq \tau(v) - \varepsilon$ and for any $d \in I$, $\mathcal{P}_{i-1}(\ell', (v+d)[z \rightarrow 0]) \geq \tau_\ell(v) - \varepsilon$.

Let $I' = [\alpha + \|v' - v\|_\infty, \beta - \|v' - v\|_\infty]$. Then

$$|I'| \geq |I| - 2\|v' - v\|_\infty \geq \tau_\ell(v) - \varepsilon - 2\|v' - v\|_\infty.$$

Moreover, since $I' \subseteq I$, we have $\mathcal{P}_{i-1}(\ell', (v+d)[z \rightarrow 0]) \geq \tau_\ell(v) - \varepsilon$ also when $d \in I'$. Additionally, for any $d \in I'$,

$$\|(v' + d)[z \rightarrow 0] - (v + d)[z \rightarrow 0]\|_\infty \leq \|v' - v\|_\infty,$$

so that

$$\begin{aligned} \mathcal{P}_{i-1}(\ell', (v+d)[z \rightarrow 0]) - \mathcal{P}_{i-1}(\ell', (v'+d)[z \rightarrow 0]) \\ \leq 2\|(v' + d)[z \rightarrow 0] - (v + d)[z \rightarrow 0]\|_\infty \\ \leq 2\|v' - v\|_\infty. \end{aligned}$$

Thus for any $d \in I'$,

$$\begin{aligned} \mathcal{P}_{i-1}(\ell', (v' + d)[z \rightarrow 0]) &\geq \mathcal{P}_{i-1}(\ell', (v + d)[z \rightarrow 0]) - 2\|v' - v\|_\infty \\ &\geq \tau_\ell(v) - \varepsilon - 2\|v' - v\|_\infty. \end{aligned}$$

Since also $|I'| \geq \tau_\ell(v) - \varepsilon - 2\|v' - v\|_\infty$, we get

$$\tau_\ell(v') \geq \min(|I'|, \inf_{d \in I'} (\mathcal{P}_{i-1}(\ell', (v' + d)[z \rightarrow 0]))) \geq \tau_\ell(v) - \varepsilon - 2\|v' - v\|_\infty.$$

□

C Proof of lemma 16

D Proofs of Section 3

D.1 Proof of Proposition 12 and Corollary 13

Proposition 12. *Let $i \in \mathbb{N}$. Let ℓ be a location of a linear timed automaton, let v_1 and v_2 be two clock valuations such that $\mathcal{P}_i(\ell, v_1)$ and $\mathcal{P}_i(\ell, v_2)$ are finite. Let $\lambda \in [0; 1]$, and $v_\lambda = \lambda \cdot v_1 + (1 - \lambda) \cdot v_2$. Then*

$$\mathcal{P}_i(\ell, v_\lambda) \geq \lambda \cdot \mathcal{P}_i(\ell, v_1) + (1 - \lambda) \cdot \mathcal{P}_i(\ell, v_2).$$

Proof. The proof is by induction on i : it is trivial for $i = 0$, since $\mathcal{P}_0(\ell, v)$ does not depend on v . Assume that the result holds true for \mathcal{P}_i , and consider \mathcal{P}_{i+1} . Let ℓ be a state of the automaton, and (ℓ, g, a, r, ℓ') be its unique outgoing transition. Let $(I_j, a) \in \text{moves}(\ell, v_j)$ for $j \in \{1, 2\}$. By definition of moves , for $j \in \{1, 2\}$ we then have $v_j + d_j \models g$ for any $d_j \in I_j$. We can then define the set $I_\lambda = \{\lambda d_1 + (1 - \lambda)d_2 \mid d_1 \in I_1, d_2 \in I_2\}$. Moreover, pick any $d_\lambda \in I_\lambda$: then $d_\lambda = \lambda \cdot d_1 + (1 - \lambda) \cdot d_2$ for some $d_1 \in I_1$ and $d_2 \in I_2$. Then $v_\lambda + d_\lambda$ can be

written as $\lambda \cdot (v_1 + d_1) + (1 - \lambda) \cdot (v_2 + d_2)$. Since both $v_1 + d_1$ and $v_2 + d_2$ satisfy guard g , by convexity of g , we have that $v_\lambda + d_\lambda \models g$. This proves that $(I_\lambda, a) \in \text{moves}(\ell, v_\lambda)$. Moreover $|I_\lambda| = \lambda \cdot |I_1| + (1 - \lambda) \cdot |I_2|$.

Fix $\varepsilon > 0$, and take two intervals I_1 and I_2 such that $\min(|I_j|, \inf\{\mathcal{P}_i(\ell', (v_j + d_j[r \rightarrow 0]) \mid d_j \in I_j)\}) \geq \mathcal{P}_i(\ell, v_j) - \varepsilon$ for $j \in \{1, 2\}$. Define I_λ as above. Then:

$$\begin{aligned}
\mathcal{P}_{i+1}(\ell, v_\lambda) &= \sup_{(I, a) \in \text{moves}(\ell, v_\lambda)} \min(|I|, \inf\{\mathcal{P}_i(\ell', v') \mid \exists d \in I. (\ell, v) \xrightarrow{d, a} (\ell', v')\}) \\
&\geq \min(|I_\lambda|, \inf\{\mathcal{P}_i(\ell', v'_\lambda) \mid \exists d_\lambda \in I_\lambda. (\ell, v_\lambda) \xrightarrow{d_\lambda, a} (\ell', v'_\lambda)\}) \\
&\quad \text{(because the supremum over all moves is larger than or equal to the value for the particular move } (I_\lambda, a)) \\
&= \min(|I_\lambda|, \inf\{\mathcal{P}_i(\ell', (v_\lambda + d_\lambda)[r \rightarrow 0]) \mid d_\lambda \in I_\lambda\}) \\
&\quad \text{(by expanding the effect of transition } (\ell, g, a, r, \ell')) \\
&= \min(|I_\lambda|, \inf\{\mathcal{P}_i(\ell', (\lambda \cdot (v_1 + d_1) + (1 - \lambda) \cdot (v_2 + d_2))[r \rightarrow 0]) \\
&\quad \mid d_1 \in I_1, d_2 \in I_2\}) \\
&\quad \text{(by definition of } I_\lambda) \\
&= \min(|I_\lambda|, \inf\{\mathcal{P}_i(\ell', (\lambda \cdot ((v_1 + d_1)[r \rightarrow 0]) + (1 - \lambda) \cdot ((v_2 + d_2)[r \rightarrow 0])) \mid d_1 \in I_1, d_2 \in I_2\}) \\
&\quad \text{(by linearity of projection)} \\
&\geq \min(|I_\lambda|, \inf\{\lambda \cdot (\mathcal{P}_i(\ell', (v_1 + d_1)[r \rightarrow 0]) + (1 - \lambda) \cdot \mathcal{P}_i(\ell', (v_2 + d_2)[r \rightarrow 0])) \mid d_1 \in I_1, d_2 \in I_2\}) \\
&\quad \text{(by induction hypothesis)} \\
&= \min(\lambda \cdot |I_1| + (1 - \lambda)|I_2|, \lambda \cdot \inf\{\mathcal{P}_i(\ell', (v_1 + d_1)[r \rightarrow 0]) \mid d_1 \in I_1\} \\
&\quad + (1 - \lambda) \cdot \inf\{\mathcal{P}_i(\ell', (v_2 + d_2)[r \rightarrow 0]) \mid d_2 \in I_2\}) \\
&\geq \lambda \cdot \min(|I_1|, \inf\{\mathcal{P}_i(\ell', (v_1 + d_1)[r \rightarrow 0]) \mid d_1 \in I_1\}) \\
&\quad + (1 - \lambda) \cdot \min(|I_2|, \inf\{\mathcal{P}_i(\ell', (v_2 + d_2)[r \rightarrow 0]) \mid d_2 \in I_2\}) \\
&\quad \text{(as } \min(a + b, a' + b') \geq \min(a, a') + \min(b, b')) \\
&\geq \lambda \cdot \mathcal{P}_{i+1}(\ell, v_1) + (1 - \lambda) \cdot \mathcal{P}_{i+1}(\ell, v_2) - \varepsilon.
\end{aligned}$$

Since this holds for any $\varepsilon > 0$, our result follows. \square

Corollary 13. *Let ℓ be a location of a linear timed automaton, v and v' be two clock valuations, $\lambda \in [0, 1]$, and $v_\lambda = \lambda \cdot v + (1 - \lambda) \cdot v'$. Then for all i :*

$$\mathcal{P}_i(\ell, v_\lambda) \geq \min\{\mathcal{P}_i(\ell, v), \mathcal{P}_i(\ell, v')\}.$$

In particular, for any valuation v , any bounded interval $[\alpha, \beta]$, and any transition $\ell \xrightarrow{g, a, r} \ell'$:

$$\inf\{\mathcal{P}_i(\ell', v') \mid \exists d \in [\alpha, \beta]. (\ell, v) \xrightarrow{d, a} (\ell', v')\} = \min\{\mathcal{P}_i(\ell', v'_\alpha), \mathcal{P}_i(\ell', v'_\beta)\}$$

where $(\ell, v) \xrightarrow{\alpha, a} (\ell', v'_\alpha)$ and $(\ell, v) \xrightarrow{\beta, a} (\ell', v'_\beta)$.

Proof. The fact that $\mathcal{P}_i(\ell, v_\lambda) \geq \min\{\mathcal{P}_i(\ell, v), \mathcal{P}_i(\ell, v')\}$ is a direct consequence of Proposition 12.

Additionally, we have

$$\inf\{\mathcal{P}_i(\ell', v') \mid \exists d \in [\alpha, \beta]. (\ell, v) \xrightarrow{d, a} (\ell', v')\} = \inf\{\mathcal{P}_i(\ell', v') \mid \exists \lambda \in [0, 1]. v' = \lambda \cdot v'_\alpha + (1 - \lambda) \cdot v'_\beta\}$$

because $(v + (\lambda\alpha + (1 - \lambda)\beta)[r \rightarrow 0]) = \lambda((v + \alpha)[r \rightarrow 0]) + (1 - \lambda)((v + \beta)[r \rightarrow 0])$. The second claim follows. \square

D.2 Proof of lemma 16

For this proof, we use a more precise definition of piecewise-affine functions:

Definition 21. Let $n \in \mathbb{N}$. An n -dimensional piecewise-affine function is a mapping $f: \mathbb{R}_{\geq 0}^n \rightarrow \overline{\mathbb{R}}$ for which there exist

- a finite family of n -dimensional linear functions $\Phi = (\varphi_k)_{1 \leq k \leq m}$, and a finite family of finite partitions $P = (P_k)_{1 \leq k \leq m}$ of \mathbb{R} ; these define the following partition of $\mathbb{R}_{\geq 0}^n$ into convex polyhedra (some of which may be empty):

$$[\Phi, P] = \{[\Phi, b] \mid b = (b_k)_{1 \leq k \leq m} \text{ s.t. for all } 1 \leq k \leq m, b_k \in P_k\}.$$

- for each convex polyhedron h of $[\Phi, P]$, an affine function f^h , which we write as $f^h(v) = F_0^h + \sum_{1 \leq k \leq n} F_k^h \cdot v_k$;

s.t. for all $v \in \mathbb{R}_{\geq 0}^n$, $f(v) = f^h(v)$ for the unique cell h of $[\Phi, P]$ containing v .

Example 3. We consider the 2-dimensional affine function f displayed on Fig. 6. Its underlying partition can be defined using two linear functions:

- $\varphi_1: (x, y) \mapsto y$, associated with the partition $P_1 = \{(-\infty; 1], (1; +\infty)\}$;
- $\varphi_2: (x, y) \mapsto x - 1$ associated with $P_2 = \{(-\infty; 0], (0; 1], (1; +\infty)\}$.




Fig. 6. An example of a 2-dimensional piecewise-affine function

This defines a partition of $\mathbb{R}_{\geq 0}^2$ into six cells: on three of them (namely C_4 , C_5 and C_6), for which $\varphi_1(x, y) \in (1; +\infty)$, our piecewise-affine function f constantly equals $-\infty$; for the other three cells:

- in $C_1 = \{(x, y) \mid \varphi_1(x, y) \in (-\infty; 1] \text{ and } \varphi_2(x, y) \in (-\infty; 0]\}$, the affine function f^{C_1} coincides with $(x, y) \mapsto \frac{x+y}{2}$;
- in $C_2 = \{(x, y) \mid \varphi_1(x, y) \in (-\infty; 1] \text{ and } \varphi_2(x, y) \in (0; 1]\}$, the affine function f^{C_2} coincides with $(x, y) \mapsto y$;
- in $C_3 = \{(x, y) \mid \varphi_1(x, y) \in (-\infty; 1] \text{ and } \varphi_2(x, y) \in (1; +\infty)\}$, the affine function f^{C_3} coincides with $(x, y) \mapsto x - 1$.

We now prove Lemma 16:

Lemma 16. *Let \mathcal{A} be a linear timed automaton with n clocks. Let (ℓ, g, a, z, ℓ') be a transition of \mathcal{A} , and assume that $v \mapsto \mathcal{P}_{i-1}(\ell', v)$ is piecewise affine, with m cells. Then $v \mapsto \mathcal{P}_i(\ell, v)$ is piecewise affine. It can be computed in time $O(m^4 \cdot (m+n)^4)$. It can be defined using a polyhedral partition of size $O(m^4 \cdot (m+n)^4)$, and with coefficients polynomial in those of \mathcal{P}_{i-1} .*

Proof. We assume that $v \mapsto \mathcal{P}_{i-1}(\ell', v)$ is not constantly $-\infty$ (if it were the case, then also $\mathcal{P}_i(\ell, v) = -\infty$ for all v). Similarly, we assume that $\text{moves}(\ell, v)$ is non-empty for some v . Since $v \mapsto \mathcal{P}_{i-1}(\ell', v)$ is piecewise-affine, we can then fix a polyhedral partition $\llbracket \Phi, P \rrbracket$ and, for each cell h in this partition, an affine functions f_h , such that $\mathcal{P}_{i-1}(\ell', v) = f_h(v)$ for the only cell h containing v .

Our procedure for computing \mathcal{P}_i in ℓ consists in listing the possible pairs of cells defining \mathcal{P}_{i-1} in ℓ' where the left- and right-bounds of the interval to be proposed lie. Our approach thus consists in listing each such pair of (possibly identical) cells (h_α, h_β) in the partition defining \mathcal{P}_{i-1} in ℓ' , and

- characterizing the set $S_{(h_\alpha, h_\beta)}$ of all valuations from which those cells can be reached by taking the transition from ℓ to ℓ' . We compute this polyhedron using quantifier elimination;
- computing the ranges for α and β that can be played in order to indeed end up respectively in h_α and h_β . These are intervals I_α and I_β , whose bounds are expressed as functions of v . Computing these bounds may require refining the polyhedron obtained at the previous step into several subpolyhedra, in order to express them as affine functions of $v \in S_{(h_\alpha, h_\beta)}$;
- for each subpolyhedra, compute the optimal values for α and β : following Corollary 13, this amounts to find values for $\alpha \in I_\alpha$ and $\beta \in I_\beta$ that maximize the following function

$$\mu: (\alpha, \beta) \mapsto \min\{\beta - \alpha; \mathcal{P}_{i-1}(\ell', (v + \alpha)[z \rightarrow 0]); \mathcal{P}_{i-1}(\ell', (v + \beta)[z \rightarrow 0])\}.$$

This is performed by applying our technical Lemma 15; it may again require another refinement of the subpolyhedra, and returns an affine function for each subpolyhedron.

For each pair (h_α, h_β) , we end up with a (partial) piecewise-affine function, defined on $S_{(h_\alpha, h_\beta)}$, returning the optimal permissiveness that can be obtained if playing interval $[\alpha, \beta]$ such that taking the transition to ℓ' after delay α (resp. β)

leads to h_α (resp. h_β). Our final step to compute \mathcal{P}_i in ℓ consists in taking the maximum of all these partial functions on their (possibly overlapping) domains.

Notice that all these computations are performed symbolically w.r.t v : we manipulate affine functions of v , with conditions on v for our computation to be valid. Figure 7 illustrates the main three steps of this procedure.




Fig. 7. Three steps of our procedure: $S_{(h_\alpha, h_\beta)}$; then compute expressions for I_α^v and I_β^v (notice that we had to refine $S_{(h_\alpha, h_\beta)}$, because the expression for I_β^v would be different for the lower part of $S_{(h_\alpha, h_\beta)}$ since it ends of a different facet of h_β); finally select best values for α and β .

We now detail the first three steps. For this, we fix to cells h_α and h_β of the partition defining \mathcal{P}_{i-1} in ℓ' . Following Definition 21, those cells can be characterized by two families, $(b_j^\alpha)_j$ and $(b_j^\beta)_j$, of cells in P , such that $h_\alpha = \bigcap_{1 \leq j \leq m} \varphi_j^{-1}(b_j^\alpha)$ and $h_\beta = \bigcap_{1 \leq j \leq m} \varphi_j^{-1}(b_j^\beta)$.

Computing $S_{(h_\alpha, h_\beta)}$. We assume that the conjunction of the guard g and the invariant $I(\ell)$ can be represented as the conjunction of one interval constraint $[L^c, U^c]$ per clock c . The set $S_{(h_\alpha, h_\beta)}$ of valuations that can reach h_α and h_β with delays $\alpha \leq \beta$ (and after taking the transition to ℓ') is defined as follows:

$$S_{h_\alpha, h_\beta} = \{v \in \mathbb{R}_{\geq 0}^n \mid \exists 0 \leq \alpha \leq \beta. \forall j. \varphi_j(v + \alpha[z \rightarrow 0]) \in b_j^\alpha \text{ and } \varphi_j(v + \beta[z \rightarrow 0]) \in b_j^\beta \text{ and } v + \alpha \in g \text{ and } v + \beta \in g\}.$$

Writing K_j for the sum of all coefficients in the linear function φ_j , and $b_j^\alpha = [l_j^\alpha, u_j^\alpha]$, condition $\varphi_j(v + \alpha[z \rightarrow 0]) \in b_j^\alpha$ can be rewritten either as $\varphi_j(v) \in b_j^\alpha$ if $K_j = 0$, or as

$$\frac{1}{K_j}(l_j^\alpha - \varphi_j(v)) \leq \alpha \leq \frac{1}{K_j}(u_j^\alpha - \varphi_j(v))$$

otherwise. The same applies for β . Using Fourier-Motzkin quantifier elimination, S_{h_α, h_β} can be written as the conjunction of the following four sets of constraints:

- existence of α is expressed as the conjunction of at most $(m+n) \cdot (m+n-1)$ conditions:
 - $m \cdot (m-1)$ conditions are of the form $\frac{1}{K_j}(l_j^\alpha - \varphi_j(v)) \leq \frac{1}{K_{j'}}(u_{j'}^\alpha - \varphi_{j'}(v))$. Notice that in those conditions, the coefficients of the linear functions sum up to zero (we name such linear functions *diagonal* in the sequel);
 - mn conditions of the form $\frac{1}{K_j}(l_j^\alpha - \varphi_j(v)) \leq U^c - v(c)$ and mn conditions of the form $L^c - v(c) \leq \frac{1}{K_j}(u_j^\alpha - \varphi_j(v))$. Again, this gives diagonal linear functions;
 - $n \cdot (n-1)$ conditions of the form $L^c - v(c) \leq U^{c'} - v(c')$; These also give rise to diagonal affine functions;
- existence of β is expressed similarly. In particular, it uses the same linear functions, which are all diagonal;
- that α is nonnegative is expressed as the conjunction of $\frac{1}{K_j}(u_j^\alpha - \varphi_j(v)) \geq 0$ for all j and $U_c - v(c) \geq 0$ for all c ;
- that $\alpha \leq \beta$ is expressed as $(m+n) \cdot (m+n-1)$ constraints similar to those in the first case. Again, no new linear functions are created in this step, compared to the previous ones, and only diagonal functions will be used.

In the end we have $3(m+n)(m+n-1) + (m+n)$ constraints (but defined with at most $(m+n)^2$ linear functions). Notice that at most $m+n$ of those linear functions may be non-diagonal, and those non-diagonal functions directly originate either from the guards or from the partition defining \mathcal{P}_{i-1} in ℓ' .

Computing the range for the bounds α and β . In case $S_{(h_\alpha, h_\beta)}$ is non-empty, we proceed with computing the values for α and β that indeed lead to h_α and h_β . For any $v \in S_{(h_\alpha, h_\beta)}$, the set I_α^v (resp. I_β^v) of values for α (resp. β) for which $v + \alpha \models g$ and $(v + \alpha)[z \rightarrow 0] \in h_\alpha$ (resp. $v + \beta \models g$ and $(v + \beta)[z \rightarrow 0] \in h_\beta$) then is an interval: from the conditions above, these intervals can be written³ $I_\alpha^v = [D_\alpha^v, E_\alpha^v]$ with

$$D_\alpha^v = \max \left(\left\{ \frac{1}{K_j}(l_j^\alpha - \varphi_j(v)) \mid 1 \leq j \leq m \right\} \cup \{L_c - v(C) \mid c \in \mathcal{C}\} \right)$$

$$E_\alpha^v = \min \left(\left\{ \frac{1}{K_j}(u_j^\alpha - \varphi_j(v)) \mid 1 \leq j \leq m \right\} \cup \{U_c - v(C) \mid c \in \mathcal{C}\} \right)$$

(and similarly for I_β). In order to have affine expressions for the bounds of I_α^v and I_β^v , we refine $S_{(h_\alpha, h_\beta)}$ into cells on which one of the affine functions in the expressions of D_α^v (resp. E_α^v) realizes the maximum (resp. minimum). This refinement is obtained by expressing the fact that the selected affine function is indeed larger than (resp. smaller than) all other functions. This may refine $S_{(h_\alpha, h_\beta)}$ into at most $(m+n)^4$ cells, defined with diagonal linear functions already introduced at the previous step.

³ Notice that the bounds of those intervals may be left- and/or right open; we only consider closed intervals to not blur the focus of our presentation, but we could handle open intervals easily.

Computing the optimal values for α and β . We let $D = \{(\alpha, \beta) \mid \alpha \in I_\alpha^v, \beta \in I_\beta^v, \alpha \leq \beta\}$. It remains to find the optimal choices for α and β , i.e., the values that maximize

$$\min\{\beta - \alpha; \inf_{\gamma \in [\alpha; \beta]} \{\mathcal{P}_{i-1}(\ell', (v + \gamma)[r \rightarrow 0])\}\}$$

over D . Thanks to Corollary 13, this amounts to maximizing

$$\mu: (\alpha, \beta) \mapsto \min\{\beta - \alpha; \mathcal{P}_{i-1}(\ell', (v + \alpha)[z \rightarrow 0]); \mathcal{P}_{i-1}(\ell', (v + \beta)[z \rightarrow 0])\}$$

over that set. Since $(v + \alpha)[z \rightarrow 0] \in h_\alpha$, we have $\mathcal{P}_{i-1}(\ell', (v + \alpha)[z \rightarrow 0]) = f_{h_\alpha}((v + \alpha)[z \rightarrow 0])$. Function f_{h_α} is an n -dimensional affine function; writing $F_{h_\alpha}^z$ for the sum of the coefficients of the clocks that are not reset in z (i.e., $F_{h_\alpha}^z = f_{h_\alpha}(\mathbf{1}) - f_{h_\alpha}(\mathbf{1}_z)$), we have $f_{h_\alpha}((v + \alpha)[z \rightarrow 0]) = F_{h_\alpha}^z \cdot \alpha + f_{h_\alpha}(v[z \rightarrow 0])$. Similarly for β . We then have

$$\mu(\alpha, \beta) = \min\{\beta - \alpha; F_{h_\alpha}^z \cdot \alpha + f_{h_\alpha}(v[z \rightarrow 0]); F_{h_\beta}^z \cdot \beta + f_{h_\beta}(v[z \rightarrow 0])\},$$

which we want to minimize over D . In case I_β^v is unbounded (which may occur when all upper bounds u_j^β and U_c equal $+\infty$), by Lemma 9 we get that f_{h_β} is constant, and we can choose $\beta = +\infty$. It remains to maximize $\min\{F_{h_\alpha}^z \cdot \alpha + f_{h_\alpha}(v[z \rightarrow 0]), f_{h_\beta}(v[z \rightarrow 0])\}$ when α ranges over I_α^v . Again, if I_α^v is unbounded, f_{h_α} is constant, and α can be chosen arbitrarily in I_α^v ; otherwise, the maximum is obtained at one of the bounds of I_α^v , depending on the sign of $F_{h_\alpha}^z$.

Now, in case I_α^v and I_β^v are bounded, we apply Lemma 15 and directly get the optimal solution. This may again require refining the polyhedron being considered into at most 13 subpolyhedra, since there may be up to 13 different cases for minimizing $\mu(\alpha, \beta)$ (see Appendix F). We then get the optimal values for α and β , as well as the value of μ at that maximal point, depending on the signs of $F_{h_\alpha}^z$ and $F_{h_\beta}^z$. It can be checked that both the optimal choices for α and β , as well as the resulting permissiveness function, are linear functions of v : indeed, in our instance of the problem of Lemma 15, a and c are constant, while b and d , and m_x , M_x , m_y and M_y are affine functions of v ; the latter may only be multiplied by constants, and/or added with one another.

The coefficients of those affine functions can be computed from those of f_{h_α} and f_{h_β} , and from those of functions $(\varphi_j)_j$ defining the partition of the piecewise-linear permissiveness function $\mathcal{P}_i(\ell', v)$: in the worst case, the numerators are multiplied by the sum of all coefficients, and the denominators may be multiplied by the product of two sums of coefficients. In any case, the space needed to store one such function (assuming binary encoding) is at most linear in the space needed to store \mathcal{P}_{i-1} . This concludes the third step of our computation.

Finalizing the computation of \mathcal{P}_i in ℓ . We now have a collection of affine functions (at most $13m^2 \cdot (m + n)^4$), associated with a polyhedron on which they give a candidate expression for \mathcal{P}_i . The polyhedra may overlap, as for each valuation we considered m^2 possible cells in which the valuation may end up.

We thus have to refine one last time the partition we obtained, by considering subpolyhedra where one of the m^2 candidate functions is larger than the other ones. This may further refine each cell into m^2 subpolyhedra, defined with up to m^2 new linear functions.

In the end, this proves that the function \mathcal{P}_i in ℓ is piecewise affine, and that it can be computed from \mathcal{P}_{i-1} in time $O(m^4 \cdot (m+n)^4)$. The partition defining \mathcal{P}_i in ℓ may have up to $O(m^4 \cdot (m+n)^4)$ cells, defined with at most $O((m+n)^2)$ linear functions. The coefficients of the affine functions defining \mathcal{P}_i are polynomials in the coefficients of the affine functions defining \mathcal{P}_{i-1} . \square

E Example of computation of permissiveness

Example 4. We slightly modify the automaton of Fig. 1, by changing the guard on the first transition as displayed on Fig. 8. We develop the computation of the permissiveness function for this automaton.




Fig. 8. Automaton of Fig. 1 where the guard on the first transition has been slightly extended

Obviously, function $\mathcal{P}(\ell_1) = \mathcal{P}_1(\ell_1)$ is unchanged. We detail the computation of $\mathcal{P}_2(\ell_0)$. Following the proof of Lemma 16, we list the pairs of possible cells where the automaton may enter ℓ_1 after delays α and β : since the transition to ℓ_1 resets y , we have two possible cells, namely⁴ $C_0 = \{(x, 0) \mid 0 \leq x \leq 1\}$ and $C_1 = \{(x, 0) \mid 1 < x \leq 2\}$. Hence we have four possible situations to consider:

1. both $(v + \alpha)[y \rightarrow 0]$ and $(v + \beta)[y \rightarrow 0]$ in C_0 ;
2. both $(v + \alpha)[y \rightarrow 0]$ and $(v + \beta)[y \rightarrow 0]$ in C_1 ;
3. $(v + \alpha)[y \rightarrow 0] \in C_0$ and $(v + \beta)[y \rightarrow 0] \in C_1$;
4. $(v + \alpha)[y \rightarrow 0] \in C_1$ and $(v + \beta)[y \rightarrow 0] \in C_0$.

For each pair, we begin with computing the set of valuations v for which there are values $0 \leq \alpha \leq \beta$ satisfying the conditions:

1. having $(v + \alpha)[y \rightarrow 0]$ and $(v + \beta)[y \rightarrow 0]$ in C_0 can be written as

$$\begin{aligned} \exists \alpha \leq \beta. 0 \leq v(y) + \alpha \leq 1 \wedge 0 \leq v(y) + \beta \leq 1 \wedge \\ 0 \leq v(x) + \alpha \leq 1 \wedge 0 \leq v(x) + \beta \leq 1. \end{aligned}$$

⁴ For convenience in this 2-clock example, we may write valuations either as v or as pairs (x, y) , depending on the situation.

The constraints on y come from the guard of the transition, while those on x correspond to having the target valuations in C_0 . In this simple case, quantifier elimination returns $\{(x, y) \mid 0 \leq x \leq 1 \wedge 0 \leq y \leq 1\}$.

2. having $(v + \alpha)[y \rightarrow 0]$ and $(v + \beta)[y \rightarrow 0]$ in C_1 translates to

$$\begin{aligned} \exists \alpha \leq \beta. 0 \leq v(y) + \alpha \leq 1 \wedge 0 \leq v(y) + \beta \leq 1 \wedge \\ 1 < v(x) + \alpha \leq 2 \wedge 1 < v(x) + \beta \leq 2. \end{aligned}$$

This results in $\{(x, y) \mid 0 \leq x \leq 2 \wedge 0 \leq y \leq 1 \wedge y \leq x\}$.

3. the case where $(v + \alpha)[y \rightarrow 0] \in C_0$ and $(v + \beta)[y \rightarrow 0] \in C_1$ writes

$$\begin{aligned} \exists \alpha \leq \beta. 0 \leq v(y) + \alpha \leq 1 \wedge 0 \leq v(y) + \beta \leq 1 \wedge \\ 0 \leq v(x) + \alpha \leq 1 \wedge 1 < v(x) + \beta \leq 2. \end{aligned}$$

This corresponds to $\{(x, y) \mid 0 \leq x \leq 1 \wedge 0 \leq y \leq 1 \wedge y \leq x\}$.

4. Finally, the situation where $(v + \alpha)[y \rightarrow 0] \in C_1$ and $(v + \beta)[y \rightarrow 0] \in C_0$ translates as

$$\begin{aligned} \exists \alpha \leq \beta. 0 \leq v(y) + \alpha \leq 1 \wedge 0 \leq v(y) + \beta \leq 1 \wedge \\ 1 < v(x) + \alpha \leq 2 \wedge 0 \leq v(x) + \beta \leq 1. \end{aligned}$$

This in particular requires $1 - v(x) < \alpha$ and $\beta \leq 1 - v(x)$, which are incompatible with the condition $\alpha \leq \beta$. Hence this case is empty.

We now compute the intervals of possible values for α and β : this just amounts to writing the conditions for having $v + \alpha$ satisfy the guard and $(v + \alpha)[y \rightarrow 0]$ belong to the target cell (the computation is identical for β):

- having $(v + \alpha)[y \rightarrow 0]$ end up in C_0 requires $\alpha \in [0; 1 - v(x)] \cap [0; 1 - v(y)]$;
- having $(v + \alpha)[y \rightarrow 0]$ end up in C_1 requires $\alpha \in (1 - v(x); 2 - v(x)] \cap [0; 1 - v(y)]$.

We end up with the following situations:

1. having both $(v + \alpha)[y := 0]$ and $(v + \beta)[y := 0]$ in C_0 can be performed from $\{(x, y) \mid 0 \leq x \leq 1 \wedge 0 \leq y \leq 1\}$; from that zone:
 - (a) if $x \leq y$, we have $I_\alpha^v = I_\beta^v = [0; 1 - y]$;
 - (b) if $x > y$, we have $I_\alpha^v = I_\beta^v = [0; 1 - x]$.
2. having both $(v + \alpha)[y := 0]$ and $(v + \beta)[y := 0]$ in C_1 can be performed from $\{(x, y) \mid 0 \leq x \leq 2 \wedge 0 \leq y \leq 1 \wedge y \leq x\}$; from that zone:
 - (a) if $x \leq 1$ and $y < x$, we have $I_\alpha^v = I_\beta^v = (1 - x; 1 - y]$;
 - (b) if $1 \leq x \leq 1 + y$, we have $I_\alpha^v = I_\beta^v = [0; 1 - y]$;
 - (c) if $1 + y \leq x \leq 2$, we have $I_\alpha^v = I_\beta^v = [0; 2 - x]$.
3. having $(v + \alpha)[y := 0] \in C_0$ and $(v + \beta)[y := 0] \in C_1$ can be performed from $\{(x, y) \mid 0 \leq x \leq 1 \wedge 0 \leq y \leq 1 \wedge y \leq x\}$. We then have $I_\alpha^v = [0; 1 - x]$ and $I_\beta^v = (1 - x, 1 - y]$.

We now have to compute the optimal values of α and β in each of these six situations:

- (1a) for the first situation, we have to maximize $(\alpha, \beta) \mapsto \min\{\beta - \alpha, x + \alpha, x + \beta\}$ over $\{(\alpha, \beta) \mid \alpha \in [0; 1 - y], \beta \in [0; 1 - y], \alpha \leq \beta\}$. This corresponds to case “ $a \geq 0$ and $c \geq 0$ ” of Lemma 15. We get:
- if $\frac{1-y-x}{2} \leq 0$, the optimal interval for the player is $[0; 1 - y]$, yielding permissiveness $1 - y$;
 - if $0 \leq \frac{1-y-x}{2}$, the optimal interval is $[\frac{1-y-x}{2}; 1 - y]$, with permissiveness $\frac{1+x-y}{2}$.
- (1b) in the second situation, we maximize $(\alpha, \beta) \mapsto \min\{\beta - \alpha, x + \alpha, x + \beta\}$ over $\{(\alpha, \beta) \mid \alpha \in [0; 1 - x], \beta \in [0; 1 - x], \alpha \leq \beta\}$. The situation is the same as above, and we get:
- if $\frac{1}{2} - x \leq 0$, the optimal interval for the player is $[0; 1 - x]$, yielding permissiveness $1 - x$;
 - if $0 \leq \frac{1}{2} - x$, the optimal interval is $[\frac{1}{2} - x; 1 - x]$, with permissiveness $\frac{1}{2}$.
- (2a) in the third situation, we maximize $(\alpha, \beta) \mapsto \min\{\beta - \alpha, 2 - (x + \alpha), 2 - (x + \beta)\}$ over $\{(\alpha, \beta) \mid \alpha \in (1 - x, 1 - y], \beta \in (1 - x; 1 - y], \alpha \leq \beta\}$. We apply Lemma 15, with $a \leq 0$ and $c \leq 0$:
- the first condition corresponds to $x \leq \frac{1}{2} + y$; there the maximal point is $\min\{x - y, 1\}$, i.e. $x - y$, and is reached at $(1 - x, 1 - y)$. Since the bound at $1 - x$ is strict, we take $1 - x + \varepsilon$ instead of $1 - x$, for some arbitrarily small $\varepsilon > 0$.
 - the second condition is $x \geq \frac{1}{2} + y$, for which the maximal value is $\frac{1}{2}$ is reached at $(1 - x, \frac{3}{2} - x)$. Again, we have to take $1 - x + \varepsilon$ instead of $1 - x$.
- (2b) in this case, we maximize the same function over $\{(\alpha, \beta) \mid \alpha \in [0, 1 - y], \beta \in [0; 1 - y], \alpha \leq \beta\}$ over the zone $(1 \leq x \leq 1 + y \leq 2)$:
- the first condition is $y \geq \frac{x}{2}$: in that zone, the maximal point is $1 - y$, reached for $(0, 1 - y)$;
 - the second condition is $y \leq \frac{x}{2}$, and the maximal value $1 - \frac{x}{2}$ is reached at $(0, 1 - \frac{x}{2})$.
- (2c) we maximize the same function over $\{(\alpha, \beta) \mid \alpha \in [0, 2 - x], \beta \in [0; 2 - x], \alpha \leq \beta\}$. Again, the second condition holds, and the maximal value is $1 - \frac{x}{2}$, reached at $(0, 1 - \frac{x}{2})$.
- (3) finally, we have to maximize $(\alpha, \beta) \mapsto \min\{\beta - \alpha, x + \alpha, 2 - x - \beta\}$ over $\{(\alpha, \beta) \mid \alpha \in [0, 1 - x], \beta \in (1 - x; 1 - y], \alpha \leq \beta\}$. Hence we are in case $a \geq 0$ and $c \leq 0$ of Lemma 15. In no cases can the first and second condition hold. Then:
- when $y \geq 1 - x$ and $y \geq \frac{x}{2}$, then the third condition holds, and the maximal value $1 - y$ is reached at $(0, 1 - y)$.
- Now, let $T_x = \frac{2}{3} - x$ and $T_y = \frac{4}{3} - x$.
- the fourth condition rewrites as $y \geq x - \frac{1}{3}$ (and the complement of the previous condition). For those points, the maximal value is $\frac{1+x-y}{2}$, reached at $(\frac{1-x-y}{2}, 1 - y)$.
 - the fifth condition is $x \geq \frac{2}{3}$ (and the complement of the condition above). There the maximal point $1 - \frac{x}{2}$ is reached for $(0, 1 - \frac{x}{2})$.
 - for the remaining points: we have $ad = 2 - x \geq -x = bc$, and neither $T_y \leq m_y$ nor $T_x \geq M_x$ hold, so that the optimal point is $2/3$, corresponding to $(\frac{2}{3} - x, \frac{4}{3} - x)$.

By superimposing those results and taking the maxima on cells where several solutions have been computed, we get the global permissiveness function depicted on Fig. 9.




Fig. 9. A linear timed automaton and its permissiveness at ℓ_0

F Proof of Lemma 15

Lemma 15. *Let $m_\alpha \leq M_\alpha$ and $m_\beta \leq M_\beta$, and $D = \{(\alpha, \beta) \in \mathbb{R}_{\geq 0}^2 \mid m_\alpha \leq \alpha \leq M_\alpha, m_\beta \leq \beta \leq M_\beta, \alpha \leq \beta\}$. Let $f: \alpha \mapsto a\alpha + b$ and $g: \beta \mapsto c\beta + d$ be two 1-dimensional affine functions, and $\mu: (\alpha, \beta) \mapsto \min\{\beta - \alpha, f(\alpha), g(\beta)\}$. Then the maximal value that μ may take over D is of one of the following five forms: $M_\beta - m_\alpha$, $\lambda \cdot f(\nu)$, $\lambda \cdot g(\mu)$, $\frac{ad-bc}{a-c}$ and $\frac{ad-bc}{(a+1)(1-c)-1}$, with $\lambda \in \{1, \frac{1}{1-c}, \frac{1}{a+1}\}$ and $\nu \in \{m_\alpha, M_\alpha, m_\beta, M_\beta\}$. This value can be computed by checking inequalities between expressions of the same forms.*

Proof. We write h for the function $(\alpha, \beta) \mapsto \beta - \alpha$. We assume that D is non-empty (i.e., $m_\alpha \leq M_\beta$). We split the proof into four cases, depending on the signs of a and c .

- ▶ *When $a \leq 0$ and $c \geq 0$, then all three functions defining μ are maximized when $\alpha = m_\alpha$ and $\beta = M_\beta$. It follows that the maximal value of μ over D is $\min\{M_\beta - m_\alpha, f(m_\alpha), g(M_\beta)\}$, and is reached at (m_α, M_β) .*
- ▶ *When $a \geq 0$ and $c \geq 0$, then for two points (α, β) and (α', β') in D with $\beta \leq \beta'$, it holds $(\alpha, \beta') \in D$ and we have $\mu(\alpha, \beta) \leq \mu(\alpha, \beta')$. Hence μ is maximized over D at a point where $\beta = M_\beta$. It remains to maximize $\alpha \mapsto \mu(\alpha, M_\beta)$ over $\{\alpha \in \mathbb{R} \mid (\alpha, M_\beta) \in D\}$. Over \mathbb{R} , this function is maximized for $\alpha_0 = \frac{M_\beta - b}{a+1}$, where $\mu(\alpha_0, M_\beta) = \min\{\frac{a \cdot M_\beta + b}{a+1}, c \cdot M_\beta + d\}$. If $(\alpha_0, M_\beta) \in D$, this is the maximum of μ over D , otherwise the maximum is reached on the border of $\{\alpha \in \mathbb{R} \mid (\alpha, M_\beta) \in D\}$, i.e. for $\alpha = m_\alpha$ or $\alpha = \min\{M_\beta, M_\alpha\}$.*




Fig. 10. Four cases for the proof of Lemma 15: when $a \leq 0$ and $c \geq 0$ (left); when $a \geq 0$ and $c \geq 0$, and symmetrically, $a \leq 0$ and $c \leq 0$ (center); when $a \geq 0$ and $c \leq 0$ (right). In each picture, the state space $\mathbb{R}_{\geq 0}^2$ is divided into three cells, depending on which function is minimal among f , g and h . In each cell, we also indicate the direction in which those functions increase.

- ▶ When $a \leq 0$ and $c \leq 0$, then by letting $\alpha' = -\beta$ and $\beta' = -\alpha$, the problem is transformed into that of maximizing $\mu'(\alpha', \beta') = \min\{\beta' - \alpha', -a\beta' + b, -c\alpha' + d\}$ over $D' = \{(\alpha', \beta') \mid -M_\alpha \leq \beta' \leq -m_\alpha, -M_\beta \leq \alpha' \leq -m_\beta, \alpha' \leq \beta'\}$. Now $-c \geq 0$ and $-a \geq 0$, and we have reduced this case to the previous one.
- ▶ When $a \geq 0$ and $c \leq 0$, we split $\mathbb{R}_{\geq 0}^2$ into three convex polyhedra, depending on which of $h(\alpha, \beta)$, $f(\alpha)$ and $g(\beta)$ is minimal, and look at the position of D w.r.t. these polyhedra (see right of Fig. 10):
 - if the upper-right point $(\min\{M_\alpha, M_\beta\}, M_\beta)$ is in the polyhedron where f is minimal (which can be checked by computing the values of $f(\min\{M_\alpha, M_\beta\})$, $g(M_\beta)$, and $h(\min\{M_\alpha, M_\beta\}, M_\beta)$), then it realizes the maximum over D ;
 - similarly, if the lower-left point $(m_\alpha, \max\{m_\alpha, m_\beta\})$ is in the polyhedron where g is minimal, then again it realizes the maximum over D ;
 - similarly, if the upper-left corner (m_α, M_β) is in the polyhedron where h is minimal, it realizes the maximum over D ;
 - if none of the above cases apply, we consider the tripoint of the diagram, whose coordinates are $(\frac{d-b(1-c)}{(a+1)(1-c)-1}, \frac{d(a+1)-b}{(a+1)(1-c)-1})$. Write T_α and T_β for those coordinates.
 - if $T_\beta > M_\beta$, then the maximal point is at $(\frac{M_\beta-b}{a+1}, M_\beta)$;
 - if $T_\alpha < m_\alpha$, then the maximal point is at $(m_\alpha, \frac{m_\alpha+d}{1-c})$;
 - if $T_\beta < T_\alpha$, i.e. $da < bc$, then (m_α, m_β) must be in the polyhedron where f is minimal among the three functions, and (M_α, M_β) must be in that where g is minimal. We then are in one of the following three situations:
 - * if $(\min(m_\beta, M_\alpha), m_\beta)$ is in the polyhedron where g is minimal, then the maximal point is at $(\frac{cm_\beta+d-b}{a}, m_\beta)$;
 - * otherwise, if $(M_\alpha, \max(m_\beta, M_\alpha))$ is in the polyhedron where g is minimal, then the maximal point is at $(\frac{b-d}{c-a}, \frac{b-d}{c-a})$;
 - * otherwise, the maximal point is at $(M_\alpha, \frac{aM_\alpha+b-d}{c})$;
 - finally, if $T_\beta \geq T_\alpha$, we again have three different cases:

- * if $T_\alpha > M_\alpha$, then the maximal point $f(M_\alpha)$ is at any point between $(M_\alpha, (a+1)M_\alpha + b)$ and $(M_\alpha, \frac{aM_\alpha + b - d}{c})$;
- * if $T_\beta < m_\beta$, then the maximal point $g(m_\beta)$ is reached at any point between $(\frac{cm_\beta + d - b}{a}, m_\beta)$ and $(m_\beta(1-c) - d, m_\beta)$;
- * otherwise, the maximal point $\frac{da - bc}{(a+1)(1-c) - 1}$ is reached at (T_α, T_β) .

□

Table 1 summarizes the various cases for computing the maximum value of function $(\alpha, \beta) \mapsto \min\{\beta - \alpha, a\alpha + b, c\beta + d\}$ over $D = \{(\alpha, \beta) \in \mathbb{R}_{\geq 0}^2 \mid m_\alpha \leq \alpha \leq M_\alpha, m_\beta \leq \beta \leq M_\beta, \alpha \leq \beta\}$.

when $a \leq 0$ and $c \geq 0$:

coordinates of maximal point	value of maximal point
(m_α, M_β)	$\min\{M_\beta - m_\alpha, am_\alpha + b, cM_\beta + d\}$

when $a \geq 0$ and $c \geq 0$:

Condition	coordinates of maximal point	value of maximal point
$\frac{M_\beta - b}{a+1} \leq m_\alpha$	(m_α, M_β)	$\min\{M_\beta - m_\alpha, cM_\beta + d\}$
$m_\alpha \leq \frac{M_\beta - b}{a+1} \leq \min\{M_\alpha, M_\beta\}$	$(\frac{M_\beta - b}{a+1}, M_\beta)$	$\min\{\frac{aM_\beta + b}{a+1}, cM_\beta + d\}$
$\min\{M_\alpha, M_\beta\} \leq \frac{M_\beta - b}{a+1}$	$(\min\{M_\alpha, M_\beta\}, M_\beta)$	$\min\{aM_\alpha + b, aM_\beta + b, cM_\beta + d\}$

when $a \leq 0$ and $c \leq 0$:

Condition	coordinates of maximal point	value of maximal point
$M_\beta \leq \frac{m_\alpha + d}{1-c}$	(m_α, M_β)	$\min\{M_\beta - m_\alpha, am_\alpha + b\}$
$\max\{m_\alpha, m_\beta\} \leq \frac{m_\alpha + d}{1-c} \leq M_\beta$	$(m_\alpha, \frac{m_\alpha + d}{1-c})$	$\min\{\frac{cm_\alpha + d}{1-c}, am_\alpha + b\}$
$\frac{m_\alpha + d}{1-c} \leq \max\{m_\alpha, m_\beta\}$	$(m_\alpha, \max\{m_\alpha, m_\beta\})$	$\min\{am_\alpha + b, cm_\beta + d, cm_\alpha + d\}$

when $a \geq 0$ and $c \leq 0$:

Condition	coordinates of maximal point	value of maximal point
$f \leq g, h$ at $(\min\{M_\alpha, M_\beta\}, M_\beta)$	$(\min\{M_\alpha, M_\beta\}, M_\beta)$	$\min\{aM_\alpha + b, aM_\beta + b\}$
$g \leq f, h$ at $(m_\alpha, \max\{m_\alpha, m_\beta\})$	$(m_\alpha, \max\{m_\alpha, m_\beta\})$	$\min\{cm_\alpha + d, cm_\beta + d\}$
$h \leq f, g$ at (m_α, M_β)	(m_α, M_β)	$M_\beta - m_\alpha$

If none of the above conditions hold: let $T_\alpha = \frac{d-b(1-c)}{(a+1)(1-c)-1}$ and $T_\beta = \frac{d(a+1)-b}{(a+1)(1-c)-1}$

$T_\beta \geq M_\beta$		$(\frac{M_\beta - b}{a+1}, M_\beta)$	$\frac{aM_\beta + b}{a+1}$
$T_\alpha \leq m_\alpha$		$(m_\alpha, \frac{m_\alpha + d}{1-c})$	$\frac{cm_\alpha + d}{1-c}$
$ad \leq bc$	$g \leq f, h$ at $(\min\{m_\beta, M_\alpha\}, m_\beta)$	$(\frac{cm_\beta + d - b}{a}, m_\beta)$	$cm_\beta + d$
	$g \leq f, h$ at $(M_\alpha, \max\{m_\beta, M_\alpha\})$	$(\frac{d-b}{a-c}, \frac{d-b}{a-c})$	$\frac{ad-bc}{a-c}$
	otherwise	$(M_\alpha, \frac{aM_\alpha + b - d}{c})$	$aM_\alpha + b$
$ad \geq bc$	$T_\beta \leq m_\beta$	$((1-c)m_\beta - d, m_\beta)$	$cm_\beta + d$
	$T_\alpha \geq M_\alpha$	$(M_\alpha, (a+1)M_\alpha + b)$	$aM_\alpha + b$
	otherwise	(T_α, T_β)	$\frac{ad-bc}{(a+1)(1-c)-1}$

Table 1. Solutions to the optimization problem of Lemma 15