Costs and Rewards in Priced Timed Automata

Martin Fränzle
Department of Computing Science, University of Oldenburg, Germany
martin.fraenzle@informatik.uni-oldenburg.de

Mahsa Shirmohammadi
CNRS & LIS, France
mahsa.shirmohammadi@lis-lab.fr

Mani Swaminathan
Department of Computing Science, University of Oldenburg, Germany
mani.swaminathan@informatik.uni-oldenburg.de

James Worrell
Department of Computer Science, University of Oxford, UK
james.worrell@cs.ox.ac.uk

Abstract

We consider Pareto analysis of reachable states of multi-priced timed automata (MPTA): timed automata equipped with multiple observers that keep track of costs (to be minimised) and rewards (to be maximised) along a computation. Each observer has a constant non-negative derivative which may depend on the location of the MPTA.

We study the Pareto Domination Problem, which asks whether it is possible to reach a target location via a run in which the accumulated costs and rewards Pareto dominate a given objective vector. We show that this problem is undecidable in general, but decidable for MPTA with at most three observers. For MPTA whose observers are all costs or all rewards, we show that the Pareto Domination Problem is PSPACE-complete. We also consider an ε-approximate Pareto Domination Problem that is decidable without restricting the number and types of observers.

We develop connections between MPTA and Diophantine equations. Undecidability of the Pareto Domination Problem is shown by reduction from Hilbert's 10^{th} Problem, while decidability for three observers is shown by a translation to a fragment of arithmetic involving quadratic forms.

2012 ACM Subject Classification Theory of computation → Timed and hybrid models

Keywords and phrases Priced Timed Automata, Pareto Domination, Diophantine Equations

Digital Object Identifier 10.4230/LIPIcs.ICALP.2018.125

Related Version A full version of this paper may be found at [10], `https://arxiv.org/abs/1803.01914`.

Acknowledgements to Dan Segal for discussions on [12] and the reviewers for their feedback.

1 Introduction

Multi Priced Timed Automata (MPTA) [5, 7, 8, 11, 17, 18, 19] extend priced timed automata [2, 3, 4, 6, 16] with *multiple observers* that capture the accumulation of costs and rewards along a computation. This extension allows to model multi-objective optimization problems beyond the scope of timed automata [1]. MPTA lie at the frontier between timed automata (for which reachability is decidable [1]) and linear hybrid automata (for which reachability is undecidable [14]). The observers exhibit richer dynamics than the clocks of timed automata

45th International Colloquium on Automata, Languages, and Programming (ICALP 2018).
Editors: Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella;
Article No. 125; pp. 125:1–125:14
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

by not being confined to unit slope in locations, but may neither be queried nor reset while taking edges. This *observability restriction* has been exploited in [17] (under a cost-divergence assumption) for carrying out a *Pareto analysis* of reachable values of the observers.

In this paper we distinguish between observers that represent *costs* (to be minimised) and those that represent *rewards* (to be maximised). Formally, we partition the set $\mathcal{Y}$ of observers into cost and reward variables and say that $\gamma \in \mathbb{R}^{\mathcal{Y}}_{\geq 0}$ *Pareto dominates* $\gamma' \in \mathbb{R}^{\mathcal{Y}}_{\geq 0}$ if $\gamma(y) \leq \gamma'(y)$ for each cost variable y and $\gamma(y) \geq \gamma'(y)$ for each reward variable y. Then the *Pareto curve* corresponding to an MPTA consists of all undominated vectors γ that are reachable in an accepting location. While cost and reward variables are syntactically identical in the underlying automaton model, distinguishing between them changes the notion of Pareto domination and the associated decision problems.

We introduce in Section 3 a decision version of the problem of computing Pareto curves for MPTA, called the *Pareto Domination Problem*. Here, given a target vector $\gamma \in \mathbb{R}^{\mathcal{Y}}_{\geq 0}$, one asks to reach an accepting location with a valuation $\gamma' \in \mathbb{R}^{\mathcal{Y}}_{\geq 0}$ that Pareto dominates γ. This has not been addressed in prior work on Pareto analysis of MPTA [17], which considers only costs or only rewards. Other works on MPTA either do not address Pareto analysis [5, 8, 11, 18, 19], or have only discrete costs updated on edges [22], or are confined to a single clock [7].

Our first main result is that the Pareto Domination Problem is undecidable in general. The undecidability proof in Section 4 is by reduction from Hilbert's 10^{th} problem. Owing to the existence of so-called "universal Diophantine equations" (of degree 4 with 58 variables [15]), our proof shows undecidabililty of the Pareto Domination Problem for some fixed but large number of observers. Undecidability of the Pareto Domination Problem entails that one cannot compute an exact Pareto curve for an arbitrary MPTA.

We consider three different approaches to recover decidability of the Pareto Domination Problem, which all have a common foundation, namely a *monotone* VASS described in Sections 2 and 5, which simulates integer runs of a given MPTA. By analysing the semi-linear reachability set of this VASS we can reduce the Pareto Domination Problem to satisfiability of a class of bilinear mixed integer-real constraints. We then consider restrictions on MPTA and variants of the Pareto Domination Problem that allow us to solve this class of constraints.

We first show in Section 6 that restricting to MPTA with only costs or only rewards yields PSPACE-completeness of the Pareto Domination Problem. Here we are able to eliminate integer variables from our bilinear constraints, resulting in a formula of linear real arithmetic. This strengthens [17, Theorem 1 and Corollary 1], whose decision procedures (that exploit well-quasi-orders for termination) do not yield complexity bounds.

Next we confine the MPTA in Section 7 to at most three observers, but allow a mix of costs and rewards. Decidability is now achieved by eliminating real variables from the bilinear constraint system, thus reducing the Pareto Domination Problem to deciding the existence of positive integer zeros of a quadratic form, which is known to be decidable from [12].

We consider in Section 8 another method to restore decidability for general MPTA with arbitrarily many costs and rewards, by studying an approximate version of the Pareto Domination Problem, called the *Gap Domination Problem*. Similar to the setting of [9], the Gap Domination Problem represents the decision version of the problem of computing ε-Pareto curves. This problem, whose input includes a tolerance $\varepsilon > 0$ and a vector $\gamma \in \mathbb{R}^{\mathcal{Y}}_{\geq 0}$, places no requirement on the answer if γ is dominated and all solutions dominating γ do so with slack at most ε. We solve the Gap Domination Problem by relaxation and rounding applied to our bilinear system of constraints.

In this paper we consider only MPTA with non-negative rates. Our approach can be generalised to obtain decidability results also in the case of negative rates by extending our foundation in Sections 2 and 5 from monotone VASS to $\mathbb{Z}$-VASS [13].

2 Background

Quadratic Diophantine Equations. For later use we recall a decidable class of non-linear Diophantine problems. Consider the quadratic equation

$$\sum_{i,j=1}^{n} a_{ij} X_i X_j + \sum_{j=1}^{n} b_j X_j + c = 0 \tag{1}$$

whose coefficients a_{ij}, b_j, and c are rational numbers. Consider also the family of constraints

$$f_1(X_1,\dots,X_n) \sim c_1 \wedge \dots \wedge f_k(X_1,\dots,X_n) \sim c_k \,, \tag{2}$$

where $f_1,\dots,f_k$ are linear forms with rational coefficients, $c_1,\dots,c_k \in \mathbb{Q}$, and $\sim \in \{<,\leq\}$.

▶ **Theorem 1** ([12]). *There is an algorithm that decides whether a given quadratic equation (1) and a family of linear inequalities (2) have a solution in $\mathbb{Z}^n$.*

Let us emphasize that in Theorem 1 at most one quadratic constraint is permitted. It is clear (e.g., by introducing a slack variable) that the theorem remains true if the equality symbol in (1) is replaced by any comparison operator in $\{<,\leq,>,\geq\}$.

Monotone VASS. A *monotone vector addition system with states* (monotone VASS) is a tuple $\mathcal{Z} = \langle n, Q, q_0, Q_f, \Sigma, \Delta\rangle$, where $n \in \mathbb{N}$ is the *dimension*, Q is a set of *states*, $q_0 \in Q$ is the *initial state*, $Q_f \subseteq Q$ is a set of *final states*, Σ is the set of *labels*, and $\Delta \subseteq Q \times \mathbb{N}^n \times \Sigma \times Q$ is the set of *transitions*.

Given such a monotone VASS $\mathcal{Z}$ as above, the family of sets $\mathrm{Reach}_{\mathcal{Z},q} \subseteq \mathbb{N}^n$, for $q \in Q$, is the minimal family (w.r.t. to set inclusion) of integer vectors such that $\mathbf{0} \in \mathrm{Reach}_{\mathcal{Z},q_0}$ and for all $q \in Q$, if $\boldsymbol{u} \in \mathrm{Reach}_{\mathcal{Z},q}$ and $(q,\boldsymbol{v},\ell,p) \in \Delta$ for some $\ell \in L$, then $\boldsymbol{u}+\boldsymbol{v} \in \mathrm{Reach}_{\mathcal{Z},p}$. Finally we define the *reachability set* of $\mathcal{Z}$ to be $\mathrm{Reach}_{\mathcal{Z}} := \bigcup_{q \in Q_f} \mathrm{Reach}_{\mathcal{Z},q}$.

For every vector $\boldsymbol{v} \in \mathbb{N}^n$ and every finite set $P = \{\boldsymbol{u}_1,\dots,\boldsymbol{u}_m\}$ of vectors in $\mathbb{N}^n$, we define the $\mathbb{N}$-*linear set* $S(\boldsymbol{v},P) := \{\boldsymbol{v} + \sum_{i=1}^{m} a_i \boldsymbol{u}_i : a_1,\dots,a_m \in \mathbb{N}\}$. We call $\boldsymbol{v}$ the *base vector* and $\boldsymbol{u}_1,\dots,\boldsymbol{u}_m \in P$ the *period vectors* of the set.

The following proposition follows from [20, Proposition 4.3] (see [10, Appendix B.1]).

▶ **Proposition 2.** *Let $\mathcal{Z} = \langle n, Q, q_0, Q_f, \Sigma, \Delta\rangle$ be a monotone VASS. Then the set $\mathrm{Reach}_{\mathcal{Z}}$ can be written as a finite union of $\mathbb{N}$-linear sets $S(\boldsymbol{v}_1,P_1),\dots,S(\boldsymbol{v}_k,P_k)$, where for $i=1,\dots,k$ the components of $\boldsymbol{v}_i$ and of each vector in P_i are bounded by $poly(n,|Q|,M)^n$ in absolute value, where M is maximum absolute value of the entries of vectors in $\mathbb{N}^n$ occurring in Δ.*

3 Multi-Priced Timed Automata and Pareto Domination

Let $\mathbb{R}_{\geq 0}$ denote the set of non-negative real numbers. Given a set $\mathcal{X} = \{x_1,\dots,x_n\}$ of *clocks*, the set $\Phi(\mathcal{X})$ of *clock constraints* is generated by the grammar $\varphi ::= \mathtt{true} \mid x \leq k \mid x \geq k \mid \varphi \wedge \varphi$, where $k \in \mathbb{N}$ is a natural number and $x \in \mathcal{X}$. A *clock valuation* is a mapping $\nu : \mathcal{X} \to \mathbb{R}_{\geq 0}$ that assigns to each clock a non-negative real number. We denote by $\mathbf{0}$ the valuation such that $\mathbf{0}(x) = 0$ for all clocks $x \in \mathcal{X}$. We write $\nu \models \varphi$ to denote that ν satisfies the constraint φ. Given $t \in \mathbb{R}_{\geq 0}$, we let $\nu + t$ be the clock valuation such that $(\nu+t)(x) = \nu(x) + t$ for all clocks $x \in \mathcal{X}$. Given $\lambda \subseteq \mathcal{X}$, let $\nu[\lambda \leftarrow 0]$ be the clock valuation such that $\nu[\lambda \leftarrow 0](x) = 0$ if $x \in \lambda$, and $\nu[\lambda \leftarrow 0](x) = \nu(x)$ otherwise.

A *multi-priced timed automaton* (MPTA) is a tuple $\mathcal{A} = \langle L, \ell_0, L_f, \mathcal{X}, \mathcal{Y}, E, R\rangle$, where L is a finite set of *locations*, $\ell_0 \in L$ is an *initial location*, $L_f \subseteq L$ is a set of *accepting locations*,

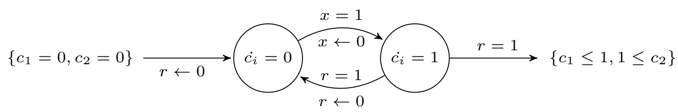

Figure 1 Predicates in curly brackets denote observer values enforced by initialisation, $c_i = 0$ with $i \in \{1, 2\}$, and by the Pareto constraint upon exit $\{c_1 \le 1, 1 \le c_2\}$. Denoting the initial value of clock x by x^*, the value of both c_1 and c_2 after n full traversals of the central cycle is nx^*. Meeting the final Pareto constraint from initial values thus requires that x^* be $\frac{1}{n}$ for some positive integer n.

$\mathcal{X}$ is a finite set of *clock variables*, $\mathcal{Y}$ is a finite set of *observers*, $E \subseteq L \times \Phi(\mathcal{X}) \times 2^{\mathcal{X}} \times L$ is the set of *edges*, $R : L \to \mathbb{N}^{\mathcal{Y}}$ is a *rate function*. Intuitively $R(\ell)$ is a vector that gives the rates of each observer in location ℓ.

A *state* of $\mathcal{A}$ is a triple (ℓ, ν, t) where ℓ is a location, ν a clock valuation, and $t \in \mathbb{R}_{\ge 0}$ is a *time stamp*. A *run* of $\mathcal{A}$ is an alternating sequence of states and edges $\rho = (\ell_0, \nu_0, t_0) \xrightarrow{e_1} (\ell_1, \nu_1, t_1) \xrightarrow{e_2} \ldots \xrightarrow{e_m} (\ell_m, \nu_m, t_m)$, where $t_0 = 0$, $\nu_0 = \mathbf{0}$, $t_{i-1} \le t_i$ for all $i \in \{1, \ldots, m\}$, and $e_i = \langle \ell_{i-1}, \varphi, \lambda, \ell_i \rangle \in E$ is such that $\nu_{i-1} + (t_i - t_{i-1}) \models \varphi$ and $\nu_i = (\nu_{i-1} + (t_i - t_{i-1}))[\lambda \leftarrow 0]$ for $i = 1, \ldots, m$. The run is *accepting* if $\ell_m \in L_f$ and said to have *granularity* $\frac{1}{g}$ for a fixed $g \in \mathbb{N}$ if all $t_i \in \mathbb{Q}$ are positive integer multiples of $\frac{1}{g}$. The *cost* of such a run is a vector $\mathrm{cost}(\rho) \in \mathbb{R}^{\mathcal{Y}}$, defined by $\mathrm{cost}(\rho) = \sum_{j=0}^{m-1} (t_{i+1} - t_i) R(\ell_i)$.

Henceforth we will assume that the set $\mathcal{Y}$ of observers of a given MPTA is partitioned into a set $\mathcal{Y}_c$ of *cost variables* and a set $\mathcal{Y}_r$ of *reward variables*. With respect to this partition we define a *domination ordering* $\preccurlyeq$ on the set of valuations $\mathbb{R}^{\mathcal{Y}}$, where $\gamma \preccurlyeq \gamma'$ if $\gamma(y) \le \gamma'(y)$ for all $y \in \mathcal{Y}_r$ and $\gamma'(y) \le \gamma(y)$ for all $y \in \mathcal{Y}_c$. Intuitively $\gamma \preccurlyeq \gamma'$ (read γ' dominates γ) if γ' is at least as good as γ in all respects.

Given $\varepsilon > 0$ we define an *ε-domination ordering* $\preccurlyeq_\varepsilon$, where $\gamma \preccurlyeq_\varepsilon \gamma'$ (read γ' ε-dominates γ) if $\gamma(y) + \varepsilon \le \gamma'(y)$ for all $y \in \mathcal{Y}_r$ and $\gamma'(y) + \varepsilon \le \gamma(y)$ for all $y \in \mathcal{Y}_c$. We can think of $\gamma \preccurlyeq_\varepsilon \gamma'$ as denoting that γ' is better than γ by an additive factor of ε in all dimensions. In particular we clearly have that $\gamma \preccurlyeq_\varepsilon \gamma'$ implies $\gamma \preccurlyeq \gamma'$.

The *Pareto Domination Problem* is as follows. Given an MPTA $\mathcal{A}$ with a set $\mathcal{Y}$ of observers and a partition of $\mathcal{Y}$ into sets $\mathcal{Y}_c$ and $\mathcal{Y}_r$ of cost and reward variables, with a target $\gamma \in \mathbb{R}^{\mathcal{Y}}$, decide whether there is an accepting run ρ of $\mathcal{A}$ such that $\gamma \preccurlyeq \mathrm{cost}(\rho)$.

The *Gap Domination Problem* is a variant of the above problem in which the input additionally includes an accuracy parameter $\varepsilon > 0$. If there is some run ρ such that $\gamma \preccurlyeq_\varepsilon \mathrm{cost}(\rho)$ then the output should be "dominated" and if there is no run ρ such that $\gamma \preccurlyeq \mathrm{cost}(\rho)$ then the output should be "not dominated". In case neither of these alternatives hold (i.e., γ is dominated but not ε-dominated) then there is no requirement on the output.

In the (Pareto) Domination Problem the objective is to *reach* an accepting location while satisfying a family of upper-bound constraints on cost variables and lower-bound constraints on reward variables. We say that an instance of the problem is *pure* if all observers are cost variables or all are reward variables (and hence all constraints are upper bounds or all are lower bounds); otherwise we call the instance *mixed*. Our problem formulation involves only simple constraints on observers, i.e., those of the form $y \le c$ or $y \ge c$ for $y \in \mathcal{Y}$. However such constraints can be used to encode more general linear constraints of the form $a_1 y_1 + \cdots + a_k y_k \sim c$, where $y_1, \ldots, y_k \in \mathcal{Y}$, $a_1, \ldots, a_k, c \in \mathbb{N}$ and $\sim \in \{\le, \ge, =\}$. To do this one introduces a fresh observer to denote each linear term $a_1 y_1 + \cdots + a_k y_k$ (two fresh observers are needed for an equality constraint).

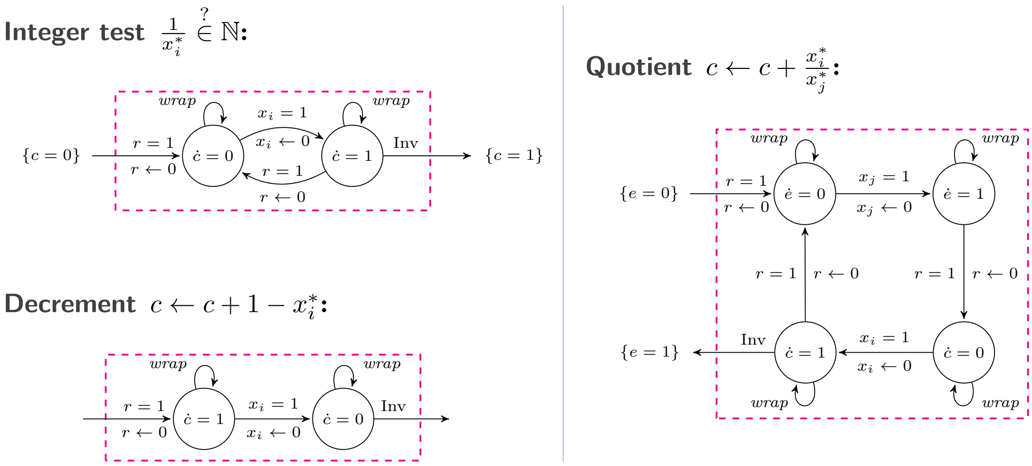

■ **Figure 2** The *wrap* self-loop denotes a family of m wrapping edges, as in [14, Fig. 14], where the j-th edge has guard $x_j = 1$ and resets x_j. In the quotient gadget, e is a fresh observer, as is c in the integer test. The integer test and quotient gadgets are annotated with predicates in curly brackets indicating the initial values of observers on entering and their target values on exiting the gadget. Enforcing these target values through a corresponding Pareto constraint guarantees the desired behaviour of the gadget.

Note that we consider timed automata without *difference constraints* on clocks, i.e., without clock guards of the form $x_i - x_j \sim k$, for $k \in \mathbb{N}$. As discussed in [10, Appendix A] all our decidability and complexity results hold also in case of such constraints.

4 Undecidability of the Pareto Domination Problem

In this section we prove undecidability of the Pareto Domination Problem. To give some insight we first give in Figure 1 an MPTA, in which the Pareto constraint $c_1 \leq 1, c_2 \geq 1$ is used to enforce that when control enters the MPTA the value of clock x is $\frac{1}{n}$ for some positive integer n.

We prove undecidability of the Pareto Domination Problem by reduction from the satisfiability problem for a fragment of arithmetic given by a language $\mathcal{L}$ that is defined as follows. There is an infinite family of variables $X_1, X_2, X_3, \ldots$ and formulas are given by the grammar $\varphi ::= X = Y + Z \mid X = YZ \mid \varphi \wedge \varphi$, where X, Y, Z range over the set of variables. The satisfiability problem for $\mathcal{L}$ asks, given a formula φ, whether there is an assignment of positive integers to the variables that satisfies φ. In [10, Appendix B.2] we show that the satisfiability problem for $\mathcal{L}$ is undecidable by reduction from Hilbert's Tenth Problem.

▶ **Theorem 3.** *The Pareto Domination Problem is undecidable.*

Proof. Consider the following problem of reaching a single valuation in $\mathbb{R}^{\mathcal{Y}}_{\geq 0}$: given an MPTA $\mathcal{A} = \langle L, \ell_0, L_f, \mathcal{X}, \mathcal{Y}, E, R \rangle$, and target valuation $\gamma \in \mathbb{R}^{\mathcal{Y}}_{\geq 0}$, decide whether there is an accepting run ρ of $\mathcal{A}$ such that $\mathrm{cost}(\rho) = \gamma$.

One can reduce the problem of reaching a given valuation to the Pareto Domination Problem as follows. Transform the MPTA $\mathcal{A}$ to an MPTA $\mathcal{A}'$ that has the same locations and edges as $\mathcal{A}$ but with two copies of each observer $y \in \mathcal{Y}$, with each copy having the same rate as y in each location. Formally $\mathcal{A}'$ has set of observers $\mathcal{Y}' = \{y_1, y_2 : y \in \mathcal{Y}\}$, where y_1 is a cost variable and y_2 is a reward variable. Then, defining $\gamma' \in \mathbb{R}^{\mathcal{Y}'}_{\geq 0}$ by $\gamma'(y_1) = \gamma'(y_2) = \gamma(y)$, we have that $\mathcal{A}'$ has an accepting run ρ' such that $\text{cost}(\rho')$ dominates γ' just in case A has an accepting run ρ such that $\text{cost}(\rho) = \gamma$.

Now we give a reduction from the satisfiability problem for $\mathcal{L}$ to the problem of reaching a single valuation. Consider an $\mathcal{L}$-formula φ over variables $X_1, \ldots, X_m$. We define an MPTA $\mathcal{A}$ over the set of clocks $\mathcal{X} = \{x_1, \cdots, x_m, r\}$. Clock x_i corresponds to the variable X_i, for $i = 1, \ldots, m$, while r is a *reference clock*. The reference clock is reset whenever it reaches 1 and is not otherwise reset—thus it keeps track of global time modulo one. After an initialisation phase the remaining clocks $x_1, \ldots, x_m$ are likewise reset in a cyclic fashion, whenever they reach 1 and not otherwise. We denote by x_i^* the value of clock x_i whenever r is 1. During the initialisation phase the values x_i^* are established non-deterministically such that $0 < x_i^* \leq 1$. The idea is that $\frac{1}{x_i^*}$ represents the value of variable X_i in φ; in particular, x_i^* is the reciprocal of a positive integer. For each atomic sub-formula in φ the automaton $\mathcal{A}$ contains a gadget that checks that the guessed valuation satisfies the sub-formula.

To present the reduction we first define three primitive gadgets. The first "integer test" gadget checks that the initial value x_i^* of clock x_i is a reciprocal of a positive integer, by adding wrapping edges on all clocks x_j other than x_i to the MPTA from Figure 1. The construction of each gadget is such that the precondition $r = 0$ holds when control enters the gadget and the postcondition $r = 1 \wedge \bigwedge_{j=1}^{m} x_j \leq 1$ holds on exiting the gadget. This last postcondition is abbreviated to Inv in the figures. For an observer c and $1 \leq i, j \leq m$, we define these three gadgets as in Figure 2.

In the following we show how to compose the three primitive operations in an MPTA to enforce the atomic constraints in the language $\mathcal{L}$. The initialisation automaton below is such that for $i = 1, \ldots, m$ the value x_i^* of clock x_i is such that $\frac{1}{x_i^*} \in \mathbb{N}$. Herein the Guess self-loop denotes a family of m edges, where the j-th edge non-deterministically resets clock x_j. Note that the incoming edge of the integer test gadget enforces $r = 1$ such that the initial guesses for the clocks x_i satisfy $x_i^* \in [0, 1]$. Of these, only reciprocals $\frac{1}{x_i^*} \in \mathbb{N}$ pass the subsequent series of integer tests.

Initialisation $X_1, \ldots, X_n \in \mathbb{N}$ **:**

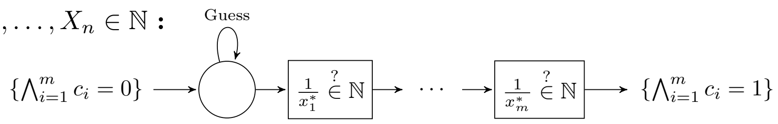

Sum $X_i = X_j + X_k$**:** According to the encoding of integer value X_n as clock value $x_n = \frac{1}{X_n}$, we have to enforce $\frac{1}{x_i^*} = \frac{1}{x_j^*} + \frac{1}{x_k^*}$, which is achieved by the following sequential combination of two quotient gadgets.

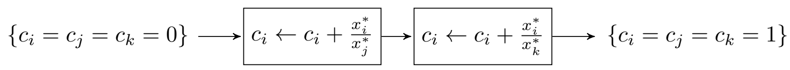

Product $X_i = X_j X_k$**:** The following gadget enforces $\frac{1}{x_i^*} = \frac{1}{x_j^*} \cdot \frac{1}{x_k^*}$:

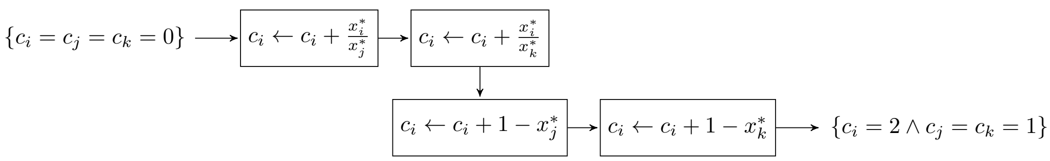

The satisfiability problem for a given $\mathcal{L}$ formula φ can now directly be reduced to the problem of reaching a single valuation $\gamma \in \mathbb{R}_{\geq 0}^{\mathcal{Y}}$ by translating each of the conjuncts of φ into the corresponding above MPTA gadget. The valuation γ encodes the target costs of the respective gadgets. ◀

Let us remark that the proof of Theorem 3 shows that undecidability of the Pareto Domination Problem already holds in case all observers have rates in $\{0,1\}$. Separately we observe that undecidability also holds in the special case that exactly one observer is a cost variable and the others are reward variables, and likewise when exactly one observer is a reward variable and the others are cost variables, when allowing rates beyond $\{0,1\}$. The idea is to reduce the problem of reaching a particular valuation $\gamma \in \mathbb{R}^{\mathcal{Y}}_{\geq 0}$ in an MPTA $\mathcal{A}$ to that of dominating a valuation $\gamma' \in \mathbb{R}^{\mathcal{Y}'}_{\geq 0}$ in a derived MPTA $\mathcal{A}'$ with set of observers $\mathcal{Y}' = \mathcal{Y} \cup \{y_{\text{sum}}\}$, where y_{sum} is a fresh variable. In $\mathcal{A}'$ we designate all $y \in \mathcal{Y}$ as cost variables and y_{sum} as a reward variable, or vice versa. Valuation γ' is specified by $\gamma'(y) = \gamma(y)$ for all $y \in \mathcal{Y}$ and $\gamma'(y_{\text{sum}}) = \sum_{y\in\mathcal{Y}} \gamma(y)$. Automaton $\mathcal{A}'$ has the same locations, edges, and rate function as those of $\mathcal{A}$ except that $R'(y_{\text{sum}}) = \sum_{y\in\mathcal{Y}} R(y)$.

5 The Simplex Automaton

This section introduces the basic construction from which we derive our positive decidability results and complexity upper bounds.

Let $\mathcal{A} = \langle L, \ell_0, L_f, \mathcal{X}, \mathcal{Y}, E, R\rangle$ be an MPTA. For a sequence of edges $e_1, \ldots, e_m \in E$, define $\mathit{Runs}(e_1, \ldots, e_m) \subseteq \mathbb{R}^m_{\geq 0}$ to be the collection of sequences of timestamps $(t_1, \ldots, t_m) \in \mathbb{R}^m_{\geq 0}$ such that $\mathcal{A}$ has a run $\rho = (\ell_0, \nu_0, t_0) \xrightarrow{e_1} (\ell_1, \nu_1, t_1) \xrightarrow{e_2} \ldots \xrightarrow{e_m} (\ell_m, \nu_m, t_m)$. Recalling that by convention $t_0 = 0$ and $\nu_0 = \mathbf{0}$, once the edges $e_1, \ldots, e_m$ have been fixed then the run ρ is determined solely by the timestamps $t_1, \ldots, t_m$. When the sequence of edges $e_1, \ldots, e_m$ is understood, we call such a sequence of timestamps a run.

▶ **Proposition 4.** *$\mathit{Runs}(e_1, \ldots, e_m) \subseteq \mathbb{R}^m_{\geq 0}$ is defined by a conjunction of difference constraints.*

The proof of Proposition 4 is in [10, Appendix B.3].

▶ **Proposition 5.** *$\mathit{Runs}(e_1, \ldots, e_m)$ is equal to the convex hull of the set of its integer points.*

Proof. Fix a positive integer M. From Proposition 4 it immediately follows that the set $\mathit{Runs}(e_1, \ldots, e_m) \cap [0, M]^m$ can be written as a conjunction of closed difference constraints $A\boldsymbol{t} \leq \boldsymbol{b}$, where A is an integer matrix, $\boldsymbol{t}$ the vector of time-stamps $t_1 \ldots t_m$, and $\boldsymbol{b}$ an integer vector. Given this, it follows that $\mathit{Runs}(e_1, \ldots, e_m) \cap [0, M]^m$, being a closed and bounded polygon, is the convex hull of its vertices. Moreover each vertex is an integer point since the matrix A here, being by Proposition 4 the incidence matrix of a balanced signed graph with half edges, is totally unimodular [21, Proposition 8A.5]. ◀

Proposition 6 shows that for Pareto reachability on an MPTA $\mathcal{A}$ with $|\mathcal{Y}| = d$ observers, it suffices to look at $d+1$-simplices of integer runs.

▶ **Proposition 6.** *For any run ρ of $\mathcal{A}$ there exists a set of at most $d+1$ integer-time runs S, all over the same sequence of edges as ρ, such that* $\text{cost}(\rho)$ *lies in the convex hull of* $\text{cost}(S)$.

Proof. Let ρ be a run of $\mathcal{A}$ over an edge-sequence $e_1, \ldots, e_m$ with time stamps $t_0, \ldots, t_m$, given by $\rho = (\ell_0, \nu_0, t_0) \xrightarrow{e_1} (\ell_1, \nu_1, t_1) \xrightarrow{e_2} \ldots \xrightarrow{e_m} (\ell_m, \nu_m, t_m)$. By Proposition 5, $(t_1, \ldots, t_m)$ lies in the convex hull of the set I of integer points in $\mathit{Runs}(e_1, \ldots, e_m)$.

Since the map $\mathit{cost} : \mathit{Runs}(e_1, \ldots, e_m) \to \mathbb{R}^d$ is linear we have that $\text{cost}(\rho)$ lies in the convex hull of $\text{cost}(I)$. Moreover by Carathéodory's Theorem there exists a subset $S \subseteq I$ of cardinality at most $d+1$ such that $\text{cost}(\rho)$ lies in the convex hull of $\text{cost}(S)$. ◀

We now exploit Proposition 6 by introducing the so-called *simplex automaton* $\mathcal{S}(\mathcal{A})$, which is a monotone VASS obtained from a given MPTA $\mathcal{A}$. The automaton $\mathcal{S}(\mathcal{A})$ generates $(d+1)$-tuples of integer-time runs of $\mathcal{A}$, such that each run in the tuple executes the same sequence of edges in $\mathcal{A}$ and the runs differ only in the times at which the edges are taken. The basic component underlying the definition of the simplex automaton is the *integer-time automaton* $\mathcal{Z}(\mathcal{A})$. This automaton is a monotone VASS that generates the integer-time runs of $\mathcal{A}$, using its counters to keep track of the running cost for each observer.

The definition of $\mathcal{Z}(\mathcal{A})$ is as follows. Let $\mathcal{A} = \langle L, \ell_0, L_f, \mathcal{X}, \mathcal{Y}, E, R\rangle$ be an MPTA. Let also $M_{\mathcal{X}} \in \mathbb{N}$ be a positive constant greater than the maximum clock constant in $\mathcal{A}$. We define a monotone VASS $\mathcal{Z}(\mathcal{A}) = \langle d, Q, q_0, Q_f, E, \Delta\rangle$, in which the dimension $d = |\mathcal{Y}|$, the set of states is $Q = L \times \{0, 1, \ldots, M_{\mathcal{X}}\}^{\mathcal{X}}$, the initial state is $q_0 = (\ell_0, \mathbf{0})$, the set of accepting states is $Q_f = L_f \times \{0, 1, \ldots, M_{\mathcal{X}}\}^{\mathcal{X}}$, the set of labels is E (i.e., the set of edges of the MPTA), and the transition relation $\Delta \subseteq Q \times \mathbb{N}^d \times E \times Q$ includes a transition $((\ell, \nu), t \cdot R(\ell), e, (\ell', \nu'))$ for every $t \in \{0, 1, \ldots, M_{\mathcal{X}}\}$ and edge $e = (\ell, \varphi, \lambda, \ell')$ in $\mathcal{A}$ s.t. $\nu \oplus t \models \varphi$ and $\nu' = (\nu \oplus t)[\lambda \leftarrow 0]$. Here $(\nu \oplus t)(x) = \min(\nu(x) + t, M_{\mathcal{X}})$ for all $x \in \mathcal{X}$. We then have:

▶ **Proposition 7.** *Given a valuation $\gamma \in \mathbb{R}^{\mathcal{Y}}_{\geq 0}$, there exists an integer-time accepting run ρ of $\mathcal{A}$ with* $\mathrm{cost}(\rho) = \gamma$ *if and only if* $\gamma \in \mathrm{Reach}_{\mathcal{Z}(\mathcal{A})}$.

The simplex automaton $\mathcal{S}(\mathcal{A})$ is built by taking $d+1$ copies of $\mathcal{Z}(\mathcal{A}) = \langle d, Q, q_0, Q_f, E, \Delta\rangle$ that synchronize on transition labels. Formally, $\mathcal{S}(\mathcal{A}) = \langle d(d+1), Q^{d+1}, \boldsymbol{q}_0, Q_f^{d(d+1)}, E, \boldsymbol{\Delta}\rangle$, where $\boldsymbol{q}_0 = (q_0, \ldots, q_0)$ and $\boldsymbol{\Delta} \subseteq Q^{d+1} \times \mathbb{Z}^{d(d+1)} \times E \times Q^{d+1}$ comprises those tuples $((q_1, \ldots, q_{d+1}), (\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{d+1}), e, (q'_1, \ldots, q'_{d+1}))$ s.t. $(q_i, \boldsymbol{v}_i, e, q'_i) \in \Delta$ for all $i \in \{1, \ldots, d+1\}$.

From Propositions 6 and 7 we have:

▶ **Proposition 8.** *Given $\gamma \in \mathbb{R}^{\mathcal{Y}}_{\geq 0}$, there exists an accepting run ρ of $\mathcal{A}$ with* $\mathrm{cost}(\rho) = \gamma$ *if and only if there exists $(\gamma_1, \ldots, \gamma_{d+1}) \in \mathrm{Reach}_{\mathcal{S}(\mathcal{A})}$ with γ in the convex hull of $\{\gamma_1, \ldots, \gamma_{d+1}\}$.*

We now introduce the following "master system" of bilinear inequalities that expresses whether $\gamma \preccurlyeq \mathrm{cost}(\rho)$ for some accepting run ρ of $\mathcal{A}$.

$$\begin{array}{lll} \gamma \preccurlyeq \lambda_1\gamma_1 + \cdots + \lambda_{d+1}\gamma_{d+1} & \quad 1 = \lambda_1 + \cdots + \lambda_{d+1} \\ (\gamma_1, \ldots, \gamma_{d+1}) \in \mathrm{Reach}_{\mathcal{S}(\mathcal{A})} & \quad 0 \leq \lambda_1, \ldots, \lambda_{d+1} \end{array} \tag{3}$$

The system has real variables $\lambda_1, \ldots, \lambda_{d+1} \in \mathbb{R}^{\mathcal{Y}}_{\geq 0}$ and integer variables $\gamma_1, \ldots, \gamma_{d+1} \in \mathbb{N}^{\mathcal{Y}}$. The key property of the master system is stated in the following Proposition 9, which follows immediately from Proposition 8.

▶ **Proposition 9.** *Given a valuation $\gamma \in \mathbb{R}^{\mathcal{Y}}_{\geq 0}$ there is an accepting run ρ of $\mathcal{A}$ such that $\gamma \preccurlyeq$* $\mathrm{cost}(\rho)$ *if and only if the system of inequalities (3) has a solution.*

Given Proposition 9, the results of Section 4 imply that satisfiability of the master system (3) is not decidable in general. In the rest of the paper we pursue different approaches to showing decidability of restrictions and variants of the Pareto Domination Problem by solving appropriately restricted versions of (3).

6 Pareto Domination Problem with Pure Constraints

In this section we show that the Pareto Domination Problem is decidable in polynomial space for the class of MPTA in which the observers are all costs. We prove this complexity upper bound by exhibiting for such an MPTA $\mathcal{A}$ and target $\gamma \in \mathbb{R}_{\geq 0}^{\mathcal{Y}}$ a positive integer M, whose bit-length is polynomial in the size of $\mathcal{A}$ and γ, such that there exists a run ρ of $\mathcal{A}$ reaching the target location with $\gamma \preccurlyeq \mathrm{cost}(\rho)$ iff there exists such a run of granularity $\frac{1}{M_1}$ for some $M_1 \leq M$. To show this we rewrite the bilinear system of inequalities (3) into an equisatisfiable disjunction of linear systems of inequalities. We thus obtain a bound on the bit-length of any satisfying assignment of (3) from which we obtain the above granularity bound. A similar bound in case of all reward variables is obtained in [10, Appendix C].

Consider an MPTA $\mathcal{A} = \langle L, \ell_0, L_f, \mathcal{X}, \mathcal{Y}, E, R\rangle$. Recall that the reachability set $\mathrm{Reach}_{\mathcal{S}(\mathcal{A})}$ can be written as a union of linear sets $S(\boldsymbol{v}_i, P_i)$, $i \in I$. More precisely, let $M_{\mathcal{Y}}$ be the maximum rate occurring in the rate function R of the given MPTA $\mathcal{A}$. We then have the following, see [10, Appendix B.4] for the proof.

▶ **Proposition 10.** *The set $\mathrm{Reach}_{\mathcal{S}(\mathcal{A})}$ can be written as a finite union of linear sets $\bigcup_{i\in I} S(\boldsymbol{v}_i, P_i)$ such that for each $i \in I$ the base vectors $\boldsymbol{v}_i$ and period vectors in P_i have entries of magnitude bounded by $poly(d, |L|, M_{\mathcal{Y}}, M_{\mathcal{X}})^{d(d+1)|\mathcal{X}|}$.*

Suppose that the set of observers $\mathcal{Y}$ with $|\mathcal{Y}| = d$ is comprised exclusively of cost variables. We will apply Proposition 10 to analyse the Pareto Domination Problem. The key observation is that in this case we can equivalently rewrite the bilinear system (3) as a disjunction of linear systems of inequalities.

As a first step we can rewrite the constraint $(\gamma_1, \ldots, \gamma_{d+1}) \in \mathrm{Reach}_{\mathcal{S}(\mathcal{A})}$ in (3) as a disjunction of constraints $(\gamma_1, \ldots, \gamma_{d+1}) \in S(\boldsymbol{v}_i, P_i)$, for $i \in I$. But since the period vectors in P_i are non-negative we can further observe that in order to satisfy the upper bound constraints on cost variables, the optimal choice of $(\gamma_1, \ldots, \gamma_{d+1}) \in S(\boldsymbol{v}_i, P_i)$ is the base vector $\boldsymbol{v}_i$. Thus we can treat $\gamma_1, \ldots, \gamma_{d+1}$ as a constant in (3).

Thus we rewrite (3) as a finite disjunction of systems of linear inequalities—one such system for each $i \in I$. For a given $i \in I$ let $\boldsymbol{v}_i = (\gamma_1^{(i)}, \ldots, \gamma_{d+1}^{(i)})$ be the base vector of the linear set $S(\boldsymbol{v}_i, P_i)$. The corresponding system of inequalities specialising (3) is

$$\gamma \preccurlyeq \lambda_1\gamma_1^{(i)} + \ldots + \lambda_{d+1}\gamma_{d+1}^{(i)},\ 1 = \lambda_1 + \cdots + \lambda_{d+1},\ 0 \leq \lambda_1, \ldots, \lambda_{d+1} \tag{4}$$

Recall that if a set of linear inequalities $A\boldsymbol{x} \geq \boldsymbol{a}$, $B\boldsymbol{x} > \boldsymbol{b}$ is feasible then it is satisfied by some $\boldsymbol{x} \in \mathbb{Q}^n$ of bit-length $poly(n, b)$, where b is the total bit-length of the entries of A, B, $\boldsymbol{a}$, and $\boldsymbol{b}$. Applying this bound and Proposition 10 we see that a solution of (4) can be written in the form $\lambda_1 = \frac{p_1}{g}, \ldots, \lambda_{d+1} = \frac{p_{d+1}}{g}$ for integers $p_1, \ldots, p_{d+1}, g$ of bit-length at most $poly(d, |\mathcal{X}|, |L|, \log(M_{\mathcal{Y}}), \log(M_{\mathcal{X}}))$. This entails that the cost vector $\lambda_1\gamma_1^{(i)} + \ldots + \lambda_{d+1}\gamma_{d+1}^{(i)}$ arises from a run of $\mathcal{A}$ with granularity $\frac{1}{g}$, thus indirectly addressing the open problem stated in [17, Section 8] on the granularity of optimal runs in MPTA.

Together with Proposition 10, this yields PSPACE-membership for the Pareto Domination Problem. As reachability in timed automata is already PSPACE-hard [1] we have:

▶ **Theorem 11.** *The Pareto Domination Problem with pure constraints is PSPACE-complete.*

7 Pareto Domination Problem with Three Mixed Observers

In this section we consider the Pareto Domination Problem for MPTA with three observers. In the case of three cost variables or three reward variables the results of Section 6 apply. Below we show decidability for two cost variables and one reward variable. The similar case of two reward variables and one cost variable is handled in [10, Appendix E].

Consider an instance of the Pareto Domination Problem given by an MPTA $\mathcal{A}$ with $|\mathcal{Y}| = 3$ observers, and a target vector $\gamma \in \mathbb{R}_{\geq 0}^{\mathcal{Y}}$. Our starting point is again Proposition 9. To apply

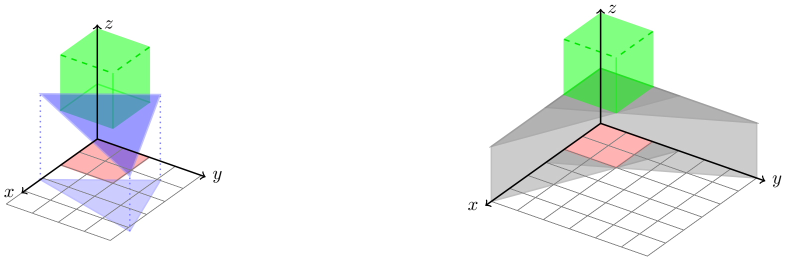

Figure 3 The target T is the green rectangular region and the blue region is S. The pink region is $\pi(T)$ and the light blue region $\pi(S)$. The grey region F is described in equation (5).

this proposition the idea is to eliminate the quantifiers over the real variables (the λ_i) in the system of equations (3) and thereby obtain a formula that lies in a decidable fragment of arithmetic (namely disjunctions of constraints of the form considered in Theorem 1).

To explain this quantifier-elimination step in more detail, let us identify $\mathbb{R}^{\mathcal{Y}}_{\geq 0}$ with $\mathbb{R}^3_{\geq 0}$. Denote by $T \subseteq \mathbb{R}^3_{\geq 0}$ the set of valuations that dominate a given fixed valuation $\gamma \in \mathbb{R}^3_{\geq 0}$. We can write $T = \{(x, y, z) \in \mathbb{R}^3_{\geq 0} : x \leq a \wedge y \leq b \wedge z \geq c\}$, where a, b, c are non-negative integer constants (see the left-hand side of Figure 3). We seek a quantifier-free formula of arithmetic that expresses that T meets a 4-simplex $S \subseteq \mathbb{R}^3_{\geq 0}$ given by the convex hull of $\{\gamma_1, \ldots, \gamma_4\}$, where $(\gamma_1, \ldots, \gamma_4) \in \mathrm{Reach}_{S(\mathcal{A})}$. However, since T is unbounded, it is clear that T meets a given 4-simplex S just in case it meets a face of S (which is a 3-simplex). Thus it will suffice to write a quantifier-free formula of arithmetic φ_T expressing that a 3-simplex in $\mathbb{R}^3_{\geq 0}$ meets T. Such a formula has nine free variables—one for each of the coordinates of the three vertices of S. We describe φ_T in the remainder of this section.

It is geometrically clear that S intersects T iff either S lies inside T, the boundary of S meets T, or the boundary of T meets S. More specifically we have the following proposition, whose proof is given in [10, Appendix B.5].

▶ **Proposition 12.** *Let $S \subseteq \mathbb{R}^3_{\geq 0}$ be a 3-simplex. Then $T \cap S$ is nonempty if and only if at least one of the following holds: (a) Some vertex of S lies in T; (b) Some bounding edge of S intersects either the face of T supported by the plane $x = a$ or the face of T supported by the plane $y = b$; (c) The bounding edge of T supported by the line $x = a \cap y = b$ intersects S.*

The following definition and proposition are key to expressing intersections of the form identified in Case (c) of Proposition 12 in terms of quadratic constraints. The idea is to identify a bounded region $F \subseteq \mathbb{R}^3_{\geq 0}$ such that in Case (c) one of the vertices of S lies in F. The proof of Proposition 13 can be found in [10, Appendix B.6].

Define a region $F \subseteq \mathbb{R}^3_{\geq 0}$ (depicted as the grey-shaded region on the right of Figure 3) by:

$$F = \{(x, y, z) \in \mathbb{R}^3_{\geq 0} \mid z < c \wedge (x + ay \leq a(b + 1) \vee y + bx \leq b(a + 1))\}. \tag{5}$$

Then we have:

▶ **Proposition 13.** *Let $S \subseteq \mathbb{R}^3_{\geq 0}$ be a 3-simplex such that $S \cap T$ is non-empty but none of the bounding edges of S meets T. Then some vertex of S lies in F.*

Denote by $\pi : \mathbb{R}^3 \to \mathbb{R}^2$ the projection of $\mathbb{R}^3$ onto the xy-plane, where $\pi(x, y, z) = (x, y)$ for all $x, y, z \in \mathbb{R}$. Write $\pi(T)$ and $\pi(S)$ for the respective images of T and S under π.

We write separate formulas $\varphi^{(1)}_T, \varphi^{(2)}_T, \varphi^{(3)}_T$, respectively expressing the three necessary and sufficient conditions for $T \cap S$ to be nonempty, as identified in Proposition 12. These are formulas of arithmetic whose free variables denote the coordinates of the three vertices of S.

Figure 4 Two cases for expressing that $\boldsymbol{c} \in \pi(S)$. The grey region is $\pi(F)$.

Some vertex of S lies in T. Denote the vertices of S by $\boldsymbol{p}, \boldsymbol{q}, \boldsymbol{r}$. Formula $\varphi_T^{(1)}$ expresses that $\boldsymbol{p} \in T$ or $\boldsymbol{q} \in T$ or $\boldsymbol{r} \in T$. This is clearly a formula of linear arithmetic.

Some bounding edge of S meets a face of T. It is straightforward to obtain $\varphi_T^{(2)}$ given a formula ψ expressing that an arbitrary line segment $\boldsymbol{xy}$ in $\mathbb{R}^3_{\geq 0}$ meets a given fixed face of T. We outline such a formula in the rest of this sub-section. For concreteness we consider the face of T supported by the plane $x = a$, which maps under π to the line segment $L = \{(a, y) : 0 \leq y \leq b\}$. Formula ψ has six free variables, respectively denoting the coordinates of $\boldsymbol{x} = (x_1, x_2, x_3)$ and $\boldsymbol{y} = (y_1, y_2, y_3)$.

Formula ψ is a conjunction of two parts. The first part expresses that $\pi(\boldsymbol{x})\pi(\boldsymbol{y})$ meets L. Since the complement of $\pi(F)$ is a convex region in $\mathbb{R}^2_{\geq 0}$ that excludes $\pi(T)$ we have that either $\pi(\boldsymbol{x}) \in \pi(F)$ or $\pi(\boldsymbol{y}) \in \pi(F)$. Moreover since $\pi(F)$ contains finitely many integer points, we can write separate sub-formulas expressing that $\pi(\boldsymbol{x})\pi(\boldsymbol{y})$ meets L for each fixed value of $\pi(\boldsymbol{x}) \in \pi(F)$ and each fixed value of $\pi(\boldsymbol{y}) \in \pi(F)$. Each of these sub-formulas can then be written in linear arithmetic [10, Appendix D].

Suppose now that $\pi(\boldsymbol{x})\pi(\boldsymbol{y})$ meets L. Then the line $\boldsymbol{xy}$ meets the face of T supported by the plane $x = a$ iff the line in xz-plane connecting (x_1, x_3) and (y_1, y_3) passes above (a, c). This requirement is expressed by the quadratic constraint (8) in [10, Appendix D].

A bounding edge of T meets S. We proceed to describe the formula $\varphi_T^{(3)}$ expressing that the bounding edge E of T, supported by the line $x = a \cap y = b$, meets S. Note that image of E under the projection π is the single point $\boldsymbol{c} = (a, b)$. Thus E meets S just in case $\boldsymbol{c} \in \pi(S)$ and the point (a, b, c) lies below the plane affinely spanned by S. We describe two formulas that respectively express these requirements.

Denote the vertices of S by $\boldsymbol{p}$, $\boldsymbol{q}$, and $\boldsymbol{r}$. We first give a formula of linear arithmetic expressing that $\boldsymbol{c} \in \pi(S)$. Notice that if $\boldsymbol{c} \in \pi(S)$ then at least one vertex of $\pi(S)$ must lie in $\pi(F)$. We now consider two cases. The first case is that exactly one vertex of $\pi(S)$ (say $\pi(\boldsymbol{p})$) lies in $\pi(F)$. The second case is that at least two vertices of of $\pi(S)$ (say $\pi(\boldsymbol{p})$ and $\pi(\boldsymbol{q})$) lie in $\pi(F)$. The two cases are respectively denoted in Figure 4, that we refer to in the following.

In the first case we can express that $\boldsymbol{c} \in \pi(S)$ by requiring that the line segment $\pi(\boldsymbol{p})\pi(\boldsymbol{q})$ crosses the edge $\boldsymbol{f}_2\boldsymbol{c}$ and $\pi(\boldsymbol{p})\pi(\boldsymbol{r})$ crosses the edge $\boldsymbol{f}_1\boldsymbol{c}$. By writing a separate constraint for each fixed value of $\pi(\boldsymbol{p}) \in \pi(F)$ the above requirements can be expressed in linear arithmetic.

In the second case we can express that $\boldsymbol{c} \in \pi(S)$ by requiring that $\boldsymbol{c}$ lies on the left of each of the directed line segments $\pi(\boldsymbol{p})\pi(\boldsymbol{q})$, $\pi(\boldsymbol{q})\pi(\boldsymbol{r})$, and $\pi(\boldsymbol{r})\pi(\boldsymbol{p})$. By writing such a constraint for each fixed value of $\pi(\boldsymbol{p})$ and $\pi(\boldsymbol{q})$ in $\pi(F)$ we obtain, again, a formula of linear arithmetic [10, Appendix D].

It remains to give a formula expressing that (a, b, c) lies below the plane affinely spanned by $\boldsymbol{p}$, $\boldsymbol{q}$, and $\boldsymbol{r}$ under the assumption that $\boldsymbol{c} \in \pi(S)$. Note here that the above-described formula expressing that $\pi(\boldsymbol{c}) \in \pi(S)$ specifies *inter alia* that $\pi(\boldsymbol{p})$, $\pi(\boldsymbol{q})$, and $\pi(\boldsymbol{r})$ are oriented counter-clockwise. Thus (a, b, c) lies below the plane affinely spanned by $\boldsymbol{p}$, $\boldsymbol{q}$, and $\boldsymbol{r}$ iff

$$\begin{vmatrix} q_1 - p_1 & r_1 - p_1 & a - p_1 \\ q_2 - p_2 & r_2 - p_2 & b - p_2 \\ q_3 - p_3 & r_3 - p_3 & c - p_3 \end{vmatrix} < 0$$

The above expression is cubic, but by Proposition 13 we may assume that $\boldsymbol{p}$ lies in the set F, which has finitely many integer points. Thus by a case analysis we may regard $\boldsymbol{p}$ as being fixed and so write the desired formula as a disjunction of atoms, each with a single quadratic term, whose satisfiability is known to be decidable from Theorem 1. This leads us to:

▶ **Theorem 14.** *The Pareto Domination Problem is decidable for at most three observers.*

Theorem 14 was proven by reduction to satisfiability of a system of arithmetic constraints with a *single* quadratic term. For the case of four observers this technique does not appear to yield arithmetic constraints in a known decidable class. Note that satisfiability of systems of constraints featuring two distinct quadratic terms is not known to be decidable in general.

In [10, Appendix F] we consider (a generalisation of) the Pareto Domination Problem for MPTA with at most two observers. In contrast to the case of three observers, we are able to show decidability for two observers by reduction to satisfiability in linear arithmetic.

8 Gap Domination Problem

In this section we give a decision procedure for the Gap Domination Problem. Given an MPTA $\mathcal{A}$, valuation $\gamma \in \mathbb{R}^{\mathcal{Y}}_{\geq 0}$, and a rational tolerance $\varepsilon > 0$, our procedure is such that

- if there is an accepting run ρ of $\mathcal{A}$ such that $\gamma \preccurlyeq_\varepsilon \mathrm{cost}(\rho)$ then we output "dominated";
- if there is no accepting run ρ of A such that $\gamma \preccurlyeq \mathrm{cost}(\rho)$ then we output "not dominated".

To do this, our approach is to find approximate solutions of the bilinear system (3) by relaxation and rounding.

Recall from Proposition 9 that (3) is satisfiable iff $\mathcal{A}$ has an accepting run ρ such that $\gamma \preccurlyeq \mathrm{cost}(\rho)$. Now we use the semi-linear decomposition of $\mathrm{Reach}_{\mathcal{S}(\mathcal{A})}$ to eliminate the constraints on integer variables from (3). In more detail, fix a decomposition of $\mathrm{Reach}_{\mathcal{S}(\mathcal{A})}$ as a union of linear sets and let $S := S(\boldsymbol{v}, P)$ be one such linear set, where $P = \{\boldsymbol{u}_1, \ldots, \boldsymbol{u}_k\}$. Then we replace the constraint $(\gamma_1, \ldots, \gamma_{d+1}) \in \mathrm{Reach}_{\mathcal{S}(\mathcal{A})}$ in (3) with

$$(\gamma_1, \ldots, \gamma_{d+1}) = \boldsymbol{v} + n_1\boldsymbol{u}_1 + \cdots + n_k\boldsymbol{u}_k\,,$$

where $n_1, \ldots, n_k$ are variables ranging over $\mathbb{N}$. We thus obtain for each choice of S a bilinear system of inequalities φ_S of the form (6), where I and J are finite sets and for each $i \in I$ and $j \in J$, it holds that f_i, g_j are linear forms (i.e., polynomials of degree one with no constant terms) with non-negative integer coefficients and c_i and d_j are rational constants.

$$\begin{array}{rclcrcl} f_i(n_1\lambda_1, n_1\lambda_2, \ldots, n_k\lambda_{d+1}) & \leq & c_i \quad (i \in I) & & \lambda_1, \ldots, \lambda_{d+1} & \geq & 0 \\ g_j(n_1\lambda_1, n_1\lambda_2, \ldots, n_k\lambda_{d+1}) & \geq & d_j \quad (j \in J) & & \lambda_1 + \cdots + \lambda_{d+1} & = & 1 \\ & n_1, \ldots, n_k \in \mathbb{N} & & & & & \end{array} \tag{6}$$

Fix a particular system φ_S, as depicted in (6). Let μ be the maximum coefficient of the f_i, $i \in I$. Given $T \subseteq \{1, \ldots, d+1\}$, we define the following constraint ψ_T on $\lambda_1, \ldots, \lambda_{d+1}$:

$$\psi_T := \bigwedge_{i \in T} \lambda_i \leq \tfrac{\varepsilon}{(d+1)k\mu} \wedge \bigwedge_{i \notin T} \lambda_i \geq \tfrac{\varepsilon}{(d+1)k\mu}\,.$$

Intuitively, ψ_T expresses that λ_i is "small" for $i \in T$ and "large" for $i \notin T$. Given any satisfying assignment of φ_S it is clear that $\lambda_1, \ldots, \lambda_{d+1}$ must satisfy φ_T for some $T \subseteq \{1, \ldots, d+1\}$.

Now fix a set $T \subseteq \{1, \ldots, d+1\}$ and consider the satisfiability of $\varphi_S \wedge \psi_T$. If $i \notin T$ then for any term $\lambda_i n_j$ that appears in an upper-bound constraint with right-hand side c in φ_S, we must have $n_j \leq \lceil \frac{c(d+1)\mu}{\varepsilon} \rceil$ in order for the constraint to be satisfied. Thus by enumerating all values of n_j we can eliminate this variable. By doing this we may assume that in $\varphi_S \wedge \psi_T$, for any term $\lambda_i n_j$ that appears on the left-hand side of an upper-bound constraint we have $i \in T$ and hence that λ_i must be "small" in any satisfying assignment.

The next step is relaxation—try to solve $\varphi_S \wedge \psi_T$ (after the above described elimination step), letting the variables $n_1, \ldots, n_k$ range over the non-negative reals. Recall here that the existential theory of real closed fields is decidable in polynomial space. If there is no real solution of $\varphi_S \wedge \psi_T$ for any S and T then there is certainly no solution over the naturals. and we can output "not dominated". On the other hand, if there is a run ρ with $\gamma \preccurlyeq_\varepsilon \mathrm{cost}(\rho)$ then for some S and T, the system $\varphi_S \wedge \psi_T$ will have a real solution in which moreover the inequalities $f_i(n_1\lambda_1, \ldots, n_k\lambda_{d+1}) \leq c_i$ for $i \in I$ all hold with slack at least ε. Given such a solution, replace n_j with $\lceil n_j \rceil$ for $j = 1, \ldots, k$. Consider the left-hand side $f_i(n_1\lambda_1, \ldots, n_k\lambda_{d+1})$ of an upper bound constraint in φ_S. Since the variables λ_i mentioned in such a linear form are small, the effect of rounding is to increase this term by at most ε. Hence the rounded valuation still satisfies φ_S thanks to the slack in the original solution. This then leads to Theorem 15 below:

▶ **Theorem 15.** *The Gap Domination Problem is decidable.*

References

1. R. Alur and D. Dill. A theory of timed automata. *TCS*, 126(2):183–235, 1994.
2. R. Alur, S. La Torre, and G. J. Pappas. Optimal paths in weighted timed automata. In M.-D. Di Benedetto and A. S-Vincentelli, editors, *HSCC*, volume 2034 of *LNCS*, pages 49–62. Springer, 2001.
3. G. Behrmann, A. Fehnker, T. Hune, K. G. Larsen, P. Pettersson, J. Romijn, and F. W. Vaandrager. Minimum-cost reachability for priced timed automata. In M.-D. Di Benedetto and A. S-Vincentelli, editors, *HSCC*, volume 2034 of *LNCS*, pages 147–161. Springer, 2001.
4. P. Bouyer, T. Brihaye, V. Bruyère, and J.-F. Raskin. On the optimal reachability problem of weighted timed automata. *Formal Methods in System Design*, 31(2):135–175, 2007.
5. P. Bouyer, E. Brinksma, and K. G. Larsen. Optimal infinite scheduling for multi-priced timed automata. *Formal Methods in System Design*, 32(1):3–23, 2008.
6. P. Bouyer, U. Fahrenberg, K. G. Larsen, N. Markey, and J. Srba. Infinite runs in weighted timed automata with energy constraints. In F. Cassez and C. Jard, editors, *FORMATS*, volume 5215 of *LNCS*, pages 33–47. Springer, 2008.
7. P. Bouyer, K. G. Larsen, and N. Markey. Model checking one-clock priced timed automata. *Logical Methods in Computer Science*, 4:1–28, 2008.
8. T. Brihaye, V. Bruyère, and J.-F. Raskin. On model-checking timed automata with stop-watch observers. *Inf. Comput.*, 204(3):408–433, 2006.
9. I. Diakonikolas and M. Yannakakis. Small approximate pareto sets for biobjective shortest paths and other problems. *SIAM J. Comput.*, 39(4):1340–1371, 2009.
10. M. Fränzle, M. Shirmohammadi, M. Swaminathan, and J. Worrell. Costs and rewards in priced timed automata. *CoRR*, 2018. URL: `https://arxiv.org/abs/1803.01914`.
11. M. Fränzle and M. Swaminathan. Revisiting decidability and optimum reachability for multi-priced timed automata. In J. Ouaknine and F. W. Vaandrager, editors, *FORMATS*, volume 5813 of *LNCS*, pages 149–163. Springer, 2009.

12 F. Grunewald and D. Segal. On the integer solutions of quadratic equations. *Journal für die reine und angewandte Mathematik*, 569:13–45, 2004.

13 C. Haase and S. Halfon. Integer vector addition systems with states. In J. Ouaknine, I. Potapov, and J. Worrell, editors, *RP*, volume 8762 of *LNCS*, pages 112–124. Springer, 2014.

14 T. A. Henzinger, P. W. Kopke, A. Puri, and P. Varaiya. What's decidable about hybrid automata? *J. Comput. Syst. Sci.*, 57(1):94–124, 1998.

15 J. P. Jones. Undecidable diophantine equations. *Bull. Amer. Math. Soc.*, 3:859–862, 1980.

16 K. G. Larsen, G. Behrmann, E. Brinksma, A. Fehnker, T. Hune, P. Pettersson, and J. Romijn. As cheap as possible: Efficient cost-optimal reachability for priced timed automata. In G. Berry, H. Comon, and A. Finkel, editors, *CAV*, volume 2102 of *LNCS*, pages 493–505. Springer, 2001.

17 K. G. Larsen and J. I. Rasmussen. Optimal reachability for multi-priced timed automata. *TCS*, 390(2-3):197–213, 2008.

18 V. Perevoshchikov. *Multi-weighted automata models and quantitative logics.* PhD thesis, University of Leipzig, 2015.

19 K. Quaas. *Kleene-Schützenberger and Büchi theorems for weighted timed automata.* PhD thesis, University of Leipzig, 2010.

20 A. W. To. Parikh images of regular languages: Complexity and applications. *CoRR*, 2010. URL: `http://arxiv.org/abs/1002.1464`.

21 T. Zaslavsky. Signed graphs. *Discrete Applied Mathematics*, 4(1):47–74, 1982.

22 Z. Zhang, B. Nielsen, K. G. Larsen, G. Nies, M. Stenger, and H. Hermanns. Pareto optimal reachability analysis for simple priced timed automata. In Z. Duan and L. Ong, editors, *ICFEM*, volume 10610 of *LNCS*, pages 481–495. Springer, 2017.