Timed Automata with Observers under Energy Constraints*

Patricia Bouyer**
Lab. Spécification et Vérification
CNRS & ENS Cachan, France
bouyer@lsv.ens-cachan.fr

Uli Fahrenberg§
Dept. of Computer Science
Aalborg University, Denmark
uli@cs.aau.dk

Kim G. Larsen§
Dept. of Computer Science
Aalborg University, Denmark
kgl@cs.aau.dk

Nicolas Markey**
Lab. Spécification et Vérification
CNRS & ENS Cachan, France
markey@lsv.ens-cachan.fr

ABSTRACT

In this paper we study one-clock priced timed automata in which prices can grow linearly ($\dot{p} = k$) or exponentially ($\dot{p} = kp$), with discontinuous updates on edges. We propose EXPTIME algorithms to decide the existence of controllers that ensure existence of infinite runs or reachability of some goal location with non-negative observer value all along the run. These algorithms consist in computing the optimal delays that should be elapsed in each location along a run, so that the final observer value is maximized (and never goes below zero).

1. INTRODUCTION

Priced timed automata [5, 3] are emerging as a useful formalism for formulating and solving a broad range of real-time resource allocation problems of importance in application areas such as, *e.g.*, embedded systems. In [7] we began the study of a new class of resource scheduling problems, namely that of constructing infinite schedules or strategies subject to boundary constraints on the accumulated use of resources.

More specifically, we proposed priced timed automata with *positive* as well as *negative* price-rates. This extension allows for the modelling of systems where resources are not only consumed but also occasionally produced or regained, *e.g.* for scheduling the behaviour of an autonomous robot which, during operation, occasionally may need to return to its base in order not to run out of energy. As an example consider the priced timed automaton in Fig. 1 with infinite behaviours repeatedly delaying in ℓ_0, ℓ_1 and ℓ_2 for a total duration of one time unit. The negative weights (-3 and -6) in ℓ_0 and ℓ_2 indicate the rates by which energy will be consumed, and the positive rate ($+6$) in ℓ_1 indicates the rate by which energy will be gained. Thus, for a given iteration the effect on the energy remaining will highly depend on the distribution of the one time unit over the three locations.

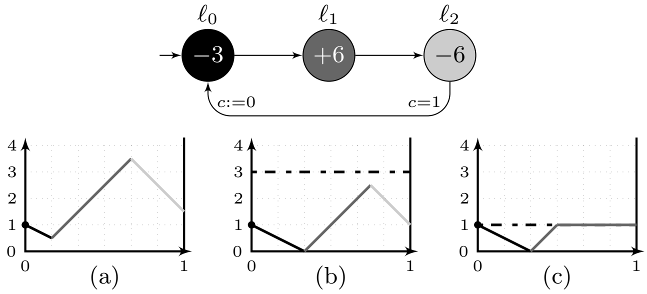

Figure 1: One-clock priced timed automaton and three types of infinite schedules: lower-bound (a), lower-upper-bound (b) and lower-weak-upper-bound (c).

In [7] three infinite scheduling problems for one-clock priced timed automata have been considered: the existence of an infinite schedule (run) during which the energy level never goes below zero (*lower-bound*), never goes below zero nor above an upper bound (*interval-bound*), and never goes below zero nor above a *weak* upper bound, which does not prevent energy-increasing behaviour from proceeding once the upper bound is reached but merely maintains the energy level at the upper bound (*lower-weak-upper-bound*). Fig. 1 illustrates the three types of schedules given an initial energy level of one.

For one-clock priced timed automata both the lower-bound and the lower-weak-upper-bound problems are shown decidable (in polynomial time) [7], whereas the interval-bound problem is proved to be undecidable in a game setting. Decidability of the interval-bound problem for one-clock priced timed automata as well as decidability of all of the considered scheduling problems for priced timed automata with two or more clocks are still unsettled.

*This work is partly supported by project FP7 Quasimodo.
**These authors are partially supported by project DOTS (ANR-06-SETI-003).
§These authors are partially supported by MT-LAB, a VKR Center of Excellence.

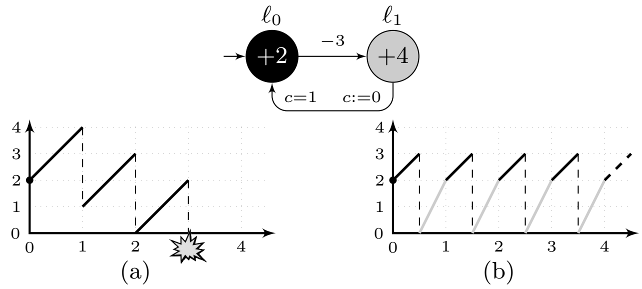

Figure 2: One-clock priced timed automaton with discrete updates. Infeasibility of region-stable lower-bound schedule (a) and optimal lower-bound schedule (b).

In this paper, we extend the decidability result of [7] for the lower-bound problem to "$1\frac{1}{2}$-clock" priced timed automata and with prices growing either linearly (*i.e.* $\dot{p} = k$) or exponentially (*i.e.* $\dot{p} = kp$). By "$1\frac{1}{2}$-clock" priced timed automata we refer to one-clock priced timed automata augmented with discontinuous (discrete) updates (*i.e.*, $p := p + c$) of the price on edges: discrete updates can be encoded using a second clock but do not provide the full expressive power of two clocks.

Surprisingly, the presence of discrete updates makes the lower-bound problem significantly more intricate. In particular, whereas region-based strategies suffice in the search for infinite lower-bound schedules for one-clock priced timed automata, this is no longer the case when discrete updates are permitted. For the priced timed automaton in Figure 1 the infinite lower-bound schedule in Figure 1(a) could be replaced by the region-based schedule in which the entire one time unit is spent in the location with the highest price-rate, *i.e.*, ℓ_1. In contrast, given initial energy-level of two, the (only possible) region-based schedule for the "$1\frac{1}{2}$-clock" priced timed automaton of Figure 2 requires the one time unit to be spent in location ℓ_0 and will eventually lead to energy-level below zero as indicated in Figure 2(a). However, choosing to leave for location ℓ_1 after having spent 0.5 time units in ℓ_0—and thus having achieved an energy-level of 3 matching the subtracting update of the edge—provides an infinite lower-bound schedule (Figure 2(b)).

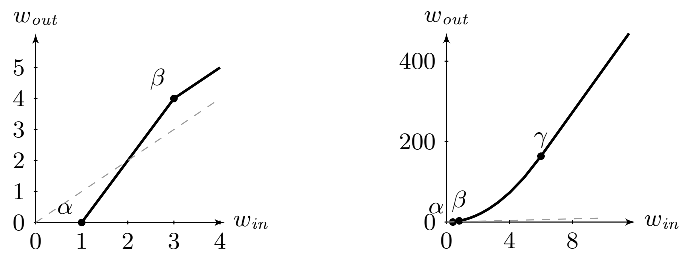

point	w_{in}	w_{out}
α	1	0
β	3	4

(a)

pt	w_{in}	w_{out}
α	≈ 0.406006	0
β	≈ 0.812012	3
γ	6	≈ 163.794

(b)

Figure 3: Energy functions for $\langle \ell_0, \ell_1 \rangle$ of the priced timed automaton of Figure 2 with linear rates (a) and exponential rates (b).

Not being able to rely on the classical region construction, the key to our decidability result is the notion of an *energy function* providing an abstraction of a path in the priced timed automaton. Given a path π, the energy function f_π maps an initial energy level w_{in} at the beginning of the path to the maximal energy w_{out} which may remain after the path. For the path $\pi = \langle \ell_0, \ell_1 \rangle$ of the priced timed automaton of Figure 2, we have already seen that $f_\pi(2) = 2$. We note that $f_\pi(1) = 0$ as the full one time-unit needs to be spent in location ℓ_0 in order to allow for a feasible run. Also $f_\pi(w_{in})$ is undefined for $w_{in} < 1$. On the other hand, $f_\pi(3) = 4$ as the full one time-unit may be spent in location ℓ_1; in fact $f_\pi(w_{in}) = w_{in} + 1$ whenever $w_{in} \geq 3$. Figure 3(a) shows the energy-function f_π. Also Figure 3(b) shows the energy function g_π for the same path π, but with exponential rates (*i.e.*, $\dot{p} = 2p$ in ℓ_0 and $\dot{p} = 4p$ in ℓ_1).

As we shall demonstrate in the remainder of the paper—for both linearly and exponentially priced timed automata—the energy function f_π for an arbitrary path π is a piecewise collection of rational power functions satisfying $f_\pi(x) - f_\pi(x') \geq x - x'$ whenever $x \geq x'$. The key for finding infinite lower-bound schedules now reduces to identifying (minimal) fixpoints $w = f_\pi(w)$, indicating that an initial energy level of w suffices for an infinite repetition of the path π. For the two energy functions f_π and g_π of Figure 2 the minimal fixpoints are $f_\pi(2) = 2$ and $g_\pi(\frac{3}{e^2-1}) = \frac{3}{e^2-1} \approx 0.47$, respectively, indicating the minimal initial energy level for infinite lower-bound schedules under the relevant linear or exponential interpretation.

Due to space restrictions, most of the proofs had to be omitted from this paper. They can be found in the long version [6].

2. TIMED AUTOMATA WITH OBSERVERS

The general formalism we introduce below, *timed automata with observers*, is intended to model control problems where resources may grow or decrease linearly or exponentially and with discontinuous updates. This includes oil tanks with pipes and drains which may be shut and opened using valves, electronic devices which may instantaneously lose energy when turned on or off, and bank accounts, or an investment portfolio, where the amount of money increases exponentially with time but where the transfer of money from one account to another typically has a fixed fee associated.

The formalism is quite general and unifies several concepts of timed automata with hybrid information found in the literature, *e.g.* in [9, 2]. In particular, it generalizes the notion of *priced*, or *weighted, timed automata* introduced in [4, 5]. In the definition, $\Phi(C)$ denotes the set of *clock constraints* on C given by the grammar $\varphi ::= c \bowtie k \mid \varphi_1 \wedge \varphi_2$ with $c \in C$, $k \in \mathbb{Z}$ and $\bowtie \in \{\leq, <, \geq, >, =\}$.

Definition 1. *A* timed automaton with observers *is a tuple* $(L, C, I, urg, E, X, fl, upd)$ *consisting of a finite set* L *of locations, a finite set* C *of clocks, location invariants* $I\colon L \to \Phi(C)$*, an urgency mapping* $urg\colon L \to \{\top, \bot\}$*, a finite set* $E \subseteq L \times \Phi(C) \times 2^C \times L$ *of edges, a finite set* X *of variables, flow conditions* $fl\colon L \to (\mathbb{R}^X \to \mathbb{R}^X)$*, and update conditions* $upd\colon E \to (\mathbb{R}^X \to \mathbb{R}^X)$*.*

Note that a timed automaton with observers has indeed an underlying timed automaton (L, C, I, E). In the following we shall write $\ell \xrightarrow{g,r} \ell'$ instead of (ℓ, g, r, ℓ') for edges.

A timed automaton (with observers) is said to be *closed* if only non-strict inequalities $\leq$ and $\geq$ are used in guards and invariants. We shall later restrict development to closed timed automata: this case contains the important aspects of our algorithm and makes exposition easier.

Using a standard construction for timed automata, urgency of locations (for which we use the *urg* mapping above) can be encoded using an extra clock, hence is not strictly necessary in the above definition. However we shall later consider the special case of *one-clock* timed automata with observers, and for these, urgency indeed adds expressivity.

In the definition below, we use the standard reset and delay operators $v[r]$, $v+d$ on valuations given by $v[r](x)=0$ if $x \in r$, $v[r](x)=v(x)$ if $x \notin r$, and $(v+d)(x)=v(x)+d$. Also, $\mathcal{D}([0,d],\mathbb{R}^X)$ denotes the set of continuous functions $[0,d] \to \mathbb{R}^X$ which are differentiable on the open interval $]0,d[$.

Definition 2. *The* semantics *of a timed automaton A with observers is given by the (infinite) transition system $[\![A]\!]=(S,T)$ with*

$$S = \{(\ell,v,w) \in L \times \mathbb{R}^C_{\geq 0} \times \mathbb{R}^X \mid v \models I(\ell)\}$$

$$\begin{aligned} T = \ &\{(\ell,v,w) \xrightarrow{e} (\ell',v',w') \mid \exists e = \ell \xrightarrow{g,r} \ell' \in E : v \models g, \\ &\qquad v' = v[r], w' = \mathit{upd}(e)(w)\} \\ \cup \ &\{(\ell,v,w) \xrightarrow{d} (\ell,v+d,w') \mid \mathit{urg}(\ell)=\bot, \\ &\qquad d \in \mathbb{R}_{\geq 0}, \exists f \in \mathcal{D}([0,d],\mathbb{R}^X): \\ &\qquad f(0)=w,\ f(d)=w',\ \text{and } \forall t \in]0,d[: \\ &\qquad v+t \models I(\ell) \text{ and } \dot f(t) = \mathit{fl}(\ell)(f(t))\} \end{aligned}$$

A run of a timed automaton A with observers is a path in its semantics $[\![A]\!]$. Hence A admits both *discrete* behaviour, indicated by transitions $\xrightarrow{e}$, and *continuous* behaviour indicated by *delay* transitions $\xrightarrow{d}$. Note that whether or not a discrete or continuous transition is available does not depend on the value w of the observer variables; the semantics defined above is indeed just the semantics of the underlying timed automaton, augmented with observer values.

We shall henceforth mostly write $\dot p$ instead of the more cumbersome $\mathit{fl}(\ell)(w)(p)$, provided that the location ℓ is clear from the context, and similarly p' instead of $\mathit{upd}(e)(w)(p)$. We also write p instead of $w(p)$ when no ambiguity arises from such an abuse of notation.

Note also that timed automata with observers form a special class of hybrid automata [1] in which the clock variables c have the restricted flow $\dot c = 1$ customary for timed automata.

In the sequel we shall consider two special classes of observers: *linear* and *exponential* ones. For a linear observer p, flow conditions are restricted to be of the form $\dot p = k$ for some constants k (possibly depending on the current location), hence linear observers admit a constant derivative (and linear growth) in locations. For an exponential observer p, flow conditions are restricted to be of the form $\dot p = kp$; that is, exponential observers have linear derivatives (hence exponential growth) in locations. We also restrict development to timed automata with *one* linear or exponential observer, with *additive* updates of the form $p' = p + c$, and with *one clock* only.

3. PROBLEMS AND RESULTS

The general problems with which we are concerned in this paper concern the existence of paths along which the observer value always remains positive:

Definition 3. *A run ρ in a timed automaton A with observers is* feasible *if the values of all the observers remain nonnegative all along ρ.*

Our problems can then be defined as follows:

Problem 1. (Reachability) Given a timed automaton A with observers X, an initial location ℓ_0, an initial valuation $w_0 \colon X \to \mathbb{R}$, and a set of goal locations $L_G \subseteq L$, either exhibit a feasible finite run in A with initial location ℓ_0, initial clock values $v(c)=0$ for all $c \in C$, and initial valuation w_0, and ending in one of the locations in L_G, or establish that no such run exists.

Problem 2. (Infinite runs) Given a timed automaton A with observers X, an initial location ℓ_0, and an initial valuation $w_0 \colon X \to \mathbb{R}$, either exhibit a feasible infinite run in A with initial location ℓ_0, initial clock values $v(c)=0$ for all $c \in C$, and initial valuation w_0, or establish that no such run exists.

In the case of linear observers, we also deal with a stronger notion of being feasible, in which the value of the observer must be larger than a given value m. The problem of *interval* bounds $m \leq w(p) \leq M$ appears to be much more difficult to handle however, see [7].

Below we give a precise definition of the classes of timed automata with observers which we shall consider in this paper and state the decidability results whose proof the rest of the paper is devoted to:

Definition 4. *A* one-clock timed automaton with one linear observer and additive updates *is a timed automaton with $C=\{c\}$, $X=\{p\}$, and for which there exist rate and weight functions $\mathit{rate}\colon L \to \mathbb{Z}$, $\mathit{weight}\colon E \to \mathbb{Z}$ such that $\mathit{fl}(\ell)(w)(p) = \mathit{rate}(\ell)$ and $\mathit{upd}(e)(w)(p) = w(p) + \mathit{weight}(e)$ for all $\ell \in L$ and $e \in E$.*

Definition 5. *A* one-clock timed automaton with one exponential observer and additive nonpositive updates *is a timed automaton with $C=\{c\}$, $X=\{p\}$, and for which there exist rate and weight functions $\mathit{rate}\colon L \to \mathbb{Z}$ and $\mathit{weight}\colon E \to \mathbb{Z}_{\leq 0}$ such that $\mathit{fl}(\ell)(w)(p) = \mathit{rate}(\ell)\, w(p)$ and $\mathit{upd}(e)(w)(p) = w(p) + \mathit{weight}(e)$ for all $\ell \in L$ and $e \in E$.*

Hence a linear observer indeed has $\dot p = \mathit{rate}(\ell)$ in all locations, and an exponential one has $\dot p = \mathit{rate}(\ell) \cdot p$. Notice that we require additive updates to be *nonpositive* for exponential observers; the general case poses additional difficulties.

Theorem 6. *Problems 1 and 2 are decidable in* EXPTIME *for closed one-clock timed automata with one linear observer and additive updates.*

Theorem 7. *Problems 1 and 2 are decidable in* EXPTIME *for closed one-clock timed automata with one exponential observer and additive* non-positive *updates.*

The rest of this paper is devoted to the proofs of these theorems, as follows: In Sections 4 to 6, we prepare the proofs by showing how to abstract observer values along *paths* in a timed automaton. In Section 7, we show how to use this abstraction to translate a timed automaton with one linear or exponential observer into a finite automaton *with energy functions* as defined in Section 8, for which the problems then can be decided.

For the sake of readability, we only present our proofs for the case of *closed* timed automata, *i.e.*, timed automata in which guards and invariants to not involve strict inequalities. This case already comprises the important ideas of our constructions.

4. OPTIMIZATIONS ALONG PATHS

Before attempting the general problem, we solve an optimization problem for paths without clock resets, both for a linear and for an exponential observer: We compute optimal delays in order to maximize the exit value along a path, as a function of the initial observer value. This is a special case of our problem, and will be the keystone of our general algorithm.

More precisely, we assume we are given an *annotated unit path*, *i.e.*, a sequence (which should be seen as a timed automaton with one observer, as explained below)

$$\pi\colon\quad \ell_0 \xrightarrow[\{c\}\quad \geq b_0]{\varphi \quad p_0} \ell_1 \xrightarrow[\geq b_1]{p_1} \ell_2 \cdots \xrightarrow[\geq b_{n-1}]{p_{n-1}} \ell_n \xrightarrow[\{c\}\quad \geq b_n]{c=1 \quad p_n} \ell_{n+1}$$

along which all unspecified guards are $0 \leq c \leq 1$, clock c is only reset along the first and last edges, and there is a global invariant $c \leq 1$. We write r_i for the rate in location ℓ_i, and assume (w.l.o.g.) that $r_0 = r_{n+1} = 0$. Along each edge, p_i indicates the discrete update, and $\geq b_i$ is an *annotation*, which is a special guard on the observer value just before firing the transition: the transition is only fireable if observer value is larger than or equal to b_i. Notice that this kind of constraint can be encoded in our one-clock models thanks to urgency, by adding two transitions with weights $-b_i$ and $+b_i$ with an urgent location in-between. An annotated unit path can thus be seen as a special kind of one-clock timed automaton with one observer and additive updates.

Let π be an annotated unit path. A *run* along π with initial observer value w is a run

$$\begin{aligned}\rho\colon\quad &(\ell_0, 0, w'_0) \xrightarrow{e_0} (\ell_1, v_1, w_1) \xrightarrow{t_1} (\ell_1, v'_1, w'_1) \xrightarrow{e_1} \cdots \\ &\cdots (\ell_n, v_n, w_n) \xrightarrow{t_n} (\ell_n, v'_n, w'_n) \xrightarrow{e_n} (\ell_{n+1}, 0, w_{n+1})\end{aligned}$$

in the corresponding timed automaton with observer with $w'_0 = w$. (Note that the precise values of w_i and w'_i depend on the type of observer we are considering.)

We write $\rho = (w, t_1, \ldots, t_n)_\pi$ to denote the run along π with initial observer value w and elapsing t_i time units in ℓ_i. Such a run ρ is a *feasible run* if it satisfies the additional constraint that $w'_i \geq b_i$ for every $0 \leq i \leq n$. Notice that these constraints are more general than our original aim of keeping observer value above 0: it suffices to let $b_i = \max(0, -p_i)$ to ensure that the value will remain nonnegative all along the run.

The *energy function* along an annotated unit path π is defined as

$$f_\pi(w) = \sup\big\{w_{n+1} \bigm| (w, t_1, \ldots, t_n)_\pi \text{ feasible run along } \pi\big\}$$

with $f_\pi(w)$ being undefined in case no feasible run along π with $w'_0 = w$ exists.

In the sequel, we explain how to compute f_π for an annotated unit path π, first in the linear and then in the exponential setting.

The first step, common to linear and exponential observers, is to remove urgent locations from our paths. Clearly enough, as no time elapses in urgent locations, the following two sequence of transitions are equivalent (w.r.t. time and observer values):

$$\ell_i \xrightarrow[\geq b_i]{p_i} \ell_{i+1}^{\mathrm{urg}} \xrightarrow[\geq b_{i+1}]{p_{i+1}} \ell_{i+2} \quad\rightsquigarrow\quad \ell_i \xrightarrow[\geq \max(b_i, b_{i+1} - p_i)]{p_i + p_{i+1}} \ell_{i+2}$$

Hence:

Lemma 8. *For any annotated unit path π, an annotated unit path $\bar\pi$ containing no urgent locations can be computed in polynomial time with $f_\pi = f_{\bar\pi}$.*

5. PATHS WITH LINEAR OBSERVER

In the following two sections, we show how to turn an annotated path into a *normal form* and how to compute f_π for normal-form paths. Both the notion of normal form, and how to compute energy functions for normal-form paths, depend on whether the observer is linear or exponential.

From now on, we can assume that π has the form

$$\pi\colon\quad \ell_0 \xrightarrow[\{c\}\quad \geq b_0]{\varphi \quad p_0} \ell_1 \xrightarrow[\geq b_1]{p_1} \ell_2 \cdots \xrightarrow[\geq b_{n-1}]{p_{n-1}} \ell_n \xrightarrow[\{c\}\quad \geq b_n]{c=1 \quad p_n} \ell_{n+1}$$

with $n \geq 1$, and that it contains no urgent locations.

Normal form. An annotated unit path as above is said to be in *normal form* (for linear observers) if all locations are non-urgent, $n \geq 1$, and one of the following three conditions holds:

- $n = 1$ (*trivial normal form*);
- all rates are positive, and $r_i < r_{i+1}$ for $1 \leq i \leq n-1$, and for every $1 \leq i \leq n-1$, it holds that $b_{i-1} + p_{i-1} < b_i$ (*positive normal form*);
- all rates are negative, and $r_i > r_{i+1}$ for $1 \leq i \leq n-1$, and for every $2 \leq i \leq n$, it holds that $b_{i-1} + p_{i-1} > b_i$ (*negative normal form*).

The proof of the fact that any annotated path can be converted into normal form, and the kind of normal form one arrives at, depend on the path's maximal location rate $\max\{r_i \mid i = 1, \ldots, n\}$. There are three cases to consider:

Case $\max\{r_i \mid i = 1, \ldots, n\} = 0$. In this case, any run which maximizes observer value will delay in one of the locations with rate 0, hence all other locations can be removed from the path (and the corresponding edges contracted). As a matter of fact, one only needs to keep *one* of the locations with zero rate; all others can be removed as well. Hence one arrives at the *trivial normal form*:

Lemma 9. *For any annotated path π (without urgent locations) such that* $\max\{r_i \mid i = 1, \ldots, n\} = 0$, *an annotated path $\widetilde\pi$ in trivial normal form can be constructed in polynomial time with $f_\pi = f_{\widetilde\pi}$.*

Case $\max\{r_i \mid i = 1, \ldots, n\} > 0$. In this case, we can transform π into an equivalent path in positive (or trivial) normal form:

Lemma 10. *For any annotated path π (without urgent locations) such that* $\max\{r_i \mid i = 1, \ldots, n\} > 0$, *an annotated path $\widetilde{\widetilde\pi}$ in positive (or trivial) normal form can be constructed in polynomial time with $f_\pi = f_{\widetilde{\widetilde\pi}}$.*

Proof sketch. The intuition is as follows: As before, the aim is to spend time in the most profitable location. However, due to annotations, we may have to delay some time in earlier

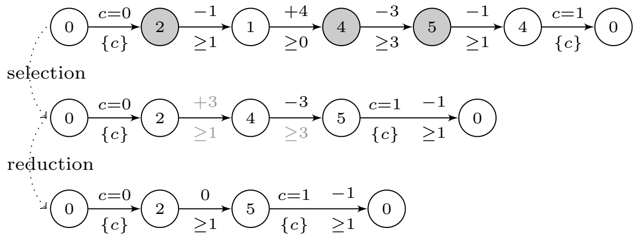

Figure 4: Conversion of annotated path into positive normal form

locations, in order to have high enough observer value to fire transitions up to this optimal location.

First we construct a sequence $(n_j)_{j\geq 0}$ of *location indices with increasing rates* as follows:

- $n_0 = 0$
- assuming n_j has been computed for some $j \geq 0$, then
  - if $r_{n_j} = \max\{r_i \mid i = 1, \dots, n\}$ is the maximal rate along π, then the sequence stops there;
  - otherwise, we let n_{j+1} be the least index $i > n_j$ for which $r_i > r_{n_j}$.

Let m be the index of the last item in $(n_j)_{j\geq 0}$. We add another last item $n_{m+1} = n + 1$, and define an intermediary annotated path $\bar{\pi}$, having $m + 2$ locations $\bar{\ell}_0$ to $\bar{\ell}_{m+1}$, with rates $\bar{r}_k = r_{n_k}$ when $0 \leq k \leq m+1$. Notice that this sequence of locations satisfies the first part of the condition for being in *positive* normal form (or in trivial normal form if $m = 1$).

We now define the transitions of $\bar{\pi}$. For $0 \leq j \leq m$, the annotated edge $\bar{\ell}_j \xrightarrow[\geq \bar{b}_j]{\bar{p}_j} \bar{\ell}_{j+1}$ is defined by

$$\bar{p}_j = \sum_{k=n_j}^{n_{j+1}-1} p_k \qquad \bar{b}_j = \max\left\{b_k - \sum_{l=n_j}^{k-1} p_l \mid n_j \leq k \leq n_{j+1}-1\right\}$$

Hence $\bar{b}_j$ is the minimum observer value needed in $\bar{\ell}_j = \ell_{n_j}$ to complete the sub-path from ℓ_{n_j} to $\ell_{n_{j+1}}$ without delaying and under observance of the lower bounds b_k.

It remains to enforce the second condition ($\bar{b}_{i-1}+\bar{p}_{i-1} < \bar{b}_i$ for $1 \leq i \leq m-2$) on $\bar{\pi}$. This is achieved by inductively replacing any offending pair of consecutive annotated edges $\bar{\ell}_{i-1} \xrightarrow[\geq \bar{b}_{i-1}]{\bar{p}_{i-1}} \bar{\ell}_i \xrightarrow[\geq \bar{b}_i]{\bar{p}_i} \bar{\ell}_{i+1}$ by a single annotated edge $\bar{\ell}_{i-1} \xrightarrow[\geq \bar{b}_{i-1}]{\bar{p}_{i-1}+\bar{p}_i} \bar{\ell}_{i+1}$. The resulting annotated path $\widetilde{\pi}$ is in normal form, and satisfies $f_\pi = f_{\widetilde{\pi}}$. □

Figure 4 shows an example of a path being converted into positive normal form.

Case $\max\{r_i \mid i = 1, \dots, n\} < 0$**.** The case with only negative rates is dual to the above one and can be handled using similar techniques:

Lemma 11. *For any annotated path π such that $\max\{r_i \mid i = 1, \dots, n\} < 0$, an annotated path $\bar{\pi}$ in negative (or trivial) normal form can be constructed in polynomial time with $f_\pi = f_{\bar{\pi}}$.*

Energy function. We now turn to the computation of the function mapping initial to final observer value along a unit path in normal form for linear observers. For the trivial normal form this is easy, as there is only one possible run along π. For the positive normal form we detail the computations below, and the negative normal form can be handled in an analogous manner.

Let

$$\pi: \quad \ell_0 \xrightarrow[\{c\}\quad \geq b_0]{\varphi \quad p_0} \ell_1 \xrightarrow[\geq b_1]{p_1} \ell_2 \cdots \xrightarrow[\geq b_{n-1}]{p_{n-1}} \ell_n \xrightarrow[\{c\}\quad \geq b_n]{c=1 \quad p_n} \ell_{n+1}$$

be an annotated unit path in positive normal form, and define the n-tuple $t^{\text{opt}} = (t_i^{\text{opt}})_{1\leq i\leq n}$ by

$$t_i^{\text{opt}} = \begin{cases} 0 & \text{if } i = m \text{ and } b_m \leq b_{m-1} + p_{m-1} \\ \frac{b_i-(b_{i-1}+p_{i-1})}{r_i} & \text{otherwise} \end{cases}$$

Since the rates are all positive and $b_i > b_{i-1} + p_{i-1}$ for all $1 \leq i \leq m-1$, these values are well-defined and positive. An important equality to notice is the following:

$$b_{n-1} + p_{n-1} + r_n \cdot t_n^{\text{opt}} = \max(b_{n-1} + p_{n-1}, b_n)$$

We prove in the sequel that those delays represent the "optimal" delays one should wait in each location, and correspond to the policy where each transition is fired as soon as the observer value satisfies the lower-bound constraint ($\geq b_i$ for the transition leaving ℓ_i).

As it may be the case that the optimal delays collected in t^{opt} do not sum up to 1 (which is the total time to be spent along π), we define another tuple $t^\star$ containing the delays which (as we shall show) have to be spent on an optimal run.

- In case $\sum_{i=1}^n t_i^{\text{opt}} > 1$, we have to cut down on the time we delay in the locations. The more profitable locations are the ones with higher rates at the end of the path, hence this is where we shall spend the delays: Letting ι_π be the largest index for which $\sum_{i=\iota_\pi}^n t_i^{\text{opt}} > 1$ (so that $\sum_{i=\iota_\pi+1}^n t_i^{\text{opt}} \leq 1$), we set

$$t_i^\star = \begin{cases} 0 & \text{for } i < \iota_\pi \\ 1 - \sum_{i=\iota_\pi+1}^n t_i^{\text{opt}} & \text{for } i = \iota_\pi \\ t_i^{\text{opt}} & \text{for } i > \iota_\pi \end{cases}$$

- In case $\sum_{i=1}^n t_i^{\text{opt}} \leq 1$, we may have to spend some extra time in one of the locations. The most profitable location for this delay is the last, hence we define $t_i^\star = t_i^{\text{opt}}$ for $1 \leq i \leq n-1$, and $t_n^\star = 1 - \sum_{i=1}^{n-1} t_i^{\text{opt}}$. We also let $\iota_\pi = 0$ in this case.

Before we prove that those delays are indeed optimal, we first compute the initial observer value needed to traverse the whole path under this policy.

- in the first case ($\iota_\pi \geq 1$), the minimal initial observer value is

$$w_{\iota_\pi}^* = b_{\iota_\pi-1} - \sum_{k=0}^{\iota_\pi-2} p_k + (t_{\iota_\pi}^{\text{opt}} - t_{\iota_\pi}^\star) \cdot r_{\iota_\pi},$$

  and the final accumulated cost is $\omega_{\iota_\pi}^* = \max(b_n, b_{n-1} + p_{n-1}) + p_n$;
- if $\iota_\pi = 0$, the minimal initial observer value is $w_0^* = b_0$, and the final accumulated cost is $\omega_0^* = \max(b_n, b_{n-1} + p_{n-1}) + p_n + (t_n^\star - t_n^{\text{opt}}) \cdot r_n$. These values actually equal w_1^* and ω_1^* defined below.

We generalize the previous construction by letting, for

$\iota_\pi + 1 \leq i \leq n$:

$$w_i^* = b_{i-1} - \sum_{k=0}^{i-2} p_k$$

$$\omega_i^* = \max(b_n, b_{n-1} + p_{n-1}) + p_n + ((t_n^\star - t_n^{\mathrm{opt}}) + \sum_{j<i} t_j^\star) \cdot r_n$$

We claim that w_i^* is the minimal initial observer value for which it is possible to spend no delay in locations ℓ_0 to ℓ_{i-1} along a feasible run, and ω_i^* is the corresponding optimal observer value at the end of the run. This can be expressed as follows:

PROPOSITION 12. *The function f_π is a piecewise affine function defined on the interval $[w_{\iota_\pi}^*, \infty[$, visiting points (w_i^*, ω_i^*), for all $\iota_\pi \leq i \leq n$, with constant slope $\dot{f}_\pi(x) \geq 1$ between two consecutive such points, and with slope $\dot{f}_\pi(x) = 1$ after (w_n^*, ω_n^*).*

EXAMPLE 1. *We consider the following example, which is already in normal form. The corresponding function f_π then looks as depicted on Figure 5:*

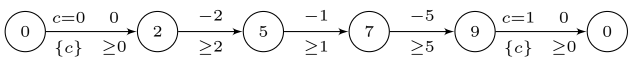

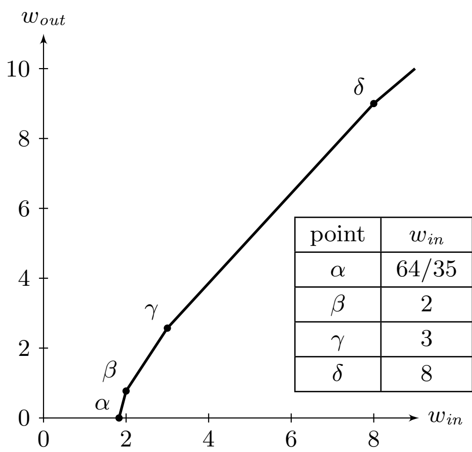

point	w_{in}	w_{out}
α	64/35	0
β	2	27/35
γ	3	18/7
δ	8	9

Figure 5: Function f_π for example with linear observer

For instance, if we enter the path with initial observer value 2, the optimal policy is to spend no time in the location with rate 2 (as we can leave it directly), then spend 1/5 time units in the next location (so that we have value 1 and can fire the outgoing transition), then spend 5/7 time units with rate 7, and the remaining 3/35 time units in the location with rate 9, ending with final observer value 27/35 (point β).

REMARK 1. *We note that the above considerations easily can be adapted to paths (without resets) with a general guard $c = k$ on the last transition (instead of $c = 1$), hence it is straight-forward to handle these. Also the restriction to closed timed automata can be lifted: we showed above how to handle paths with non-strict guards only, and the general case is similar. In this case, the energy function f_π gives, for each input value w, the* supremum *$f_\pi(w)$ of the observer values obtainable as output, an whether or not this value is actually attained for an input can be decided by looking at the delays spent in each location.*

6. PATHS WITH EXPONENTIAL OBSERVER

Normal form. As for linear observers, we are interested in computing the energy function along a unit path, by first transforming it into a normal form and then computing the energy function for normal-form paths. In this case however, we have to restrict the kinds of paths we can handle:

- We assume that the edge weights p_i are nonpositive, and that at least one of the rates r_i is nonnegative;
- paths are not annotated, *i.e.* we do not impose "local" constraints of the form "$\geq b_i$" in this case, and only require that observer value always be nonnegative along the run.

These restrictions amount to only considering the *positive* normal form, without local observer constraints. As in the previous case, we could handle the case where all rates are negative in a similar way (with a suitable notion of *negative* normal form). The other restrictions are purely technical: Currently we do not know how to handle paths with mixed positive and negative updates, or with local constraints, but we expect our techniques to also extend to these settings.

For the sequel, we again fix a unit path

$$\pi\colon \quad \ell_0 \xrightarrow[\{c\}]{\varphi \quad p_0} \ell_1 \xrightarrow{p_1} \ell_2 \cdots \xrightarrow{p_{n-1}} \ell_n \xrightarrow[\{c\}]{c=1 \quad p_n} \ell_{n+1}$$

satisfying the above constraints, and with $r_0 = r_{n+1} = 0$. As in the previous section, our aim is to compute f_π for such a path (but now with exponential observer), mapping initial to maximum final observer value.

A path as above is said to be in *normal form* (for exponential observers) if all locations are non-urgent, $m \geq 1$, and one of the following two conditions holds:

- $m = 1$ (*trivial normal form*);
- all rates are positive, and $r_i < r_{i+1}$ for $1 \leq i \leq m-1$, and for every $2 \leq i \leq m-1$, it holds that $\frac{p_{i-1} r_{i-1} r_i}{r_{i-1} - r_i} < \frac{p_i r_i r_{i+1}}{r_i - r_{i+1}}$ (*positive normal form*);

The last condition for being in positive normal form is the counterpart, for exponential observers, to the condition "$b_i > b_{i-1} + p_{i-1}$" which we had in the case of linear observers.

Such a normal form can be computed:

PROPOSITION 13. *Assume π is a unit path with nonpositive edge weights and such that* $\max\{r_i \mid i = 1, \ldots, n\} \geq 0$. *Then we can construct in polynomial time a path $\widetilde{\pi}$ in normal form for exponential observers so that $f_\pi = f_{\widetilde{\pi}}$.*

The proof relies on arguments similar to the ones we used for the linear case.

Energy function. Along a path in positive normal form, we can decide whether a given initial observer value is sufficient to reach the last location:

PROPOSITION 14. *Let π be a path in positive normal form (for exponential observers) and w an initial observer value. Then we can decide whether there is a feasible run along π with initial observer value w, and we can compute the value $f_\pi(w)$.*

Notice that contrary to the linear case, it is not sufficient to fire a transition as soon as the observer value can afford paying the nonpositive update: Consider the two-state automaton of Figure 2. If the initial observer value is 3, it is allowed to immediately fire the transition to ℓ_1, but this would set the energy level to 0, and the exponential growth would be annihilated.

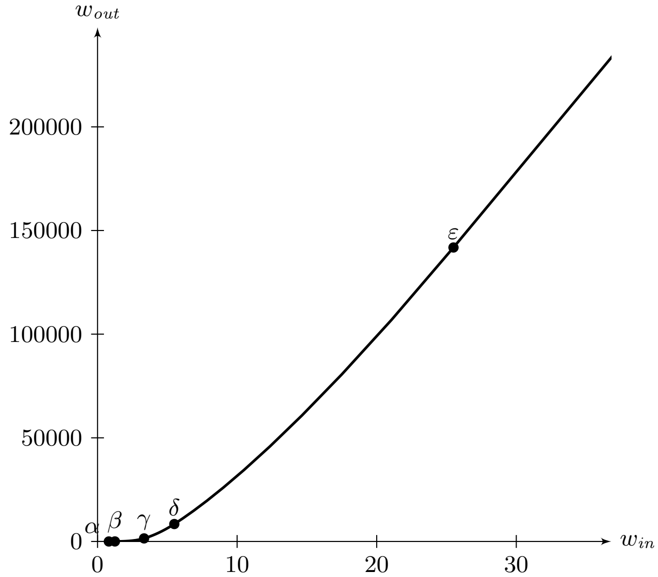

pt	w_{in}	w_{out}
α	$e^{-2}*10/3*(21/8)^{2/5}*(2)^{2/7}$ ≈ 0.809003	0
β	$e^{-2}*10/3*(21/8)^{2/5}*(9)^{2/7}$ ≈ 1.24332	$35/2$
γ	$10/3$	$e^{9}*35/2*(8/21)^{9/5}*(1/9)^{9/7}$ ≈ 1480.38
δ	$11/2$	$e^{9}*35/2*(1/9)^{9/7}$ ≈ 8410.18
ε	$51/2$	$e^{9}*35/2$ ≈ 141804

interval	equation of the curve
$\alpha - \beta$	$w_{\text{out}} = \frac{45}{2} \cdot \left(\frac{w_{\text{in}}}{e^{-2}*10/3*(21/8)^{2/5}*(9)^{2/7}}\right)^{7/2} - 5$
$\beta - \gamma$	$w_{\text{out}} = \frac{35}{2} \cdot \left(\frac{w_{\text{in}}}{e^{-2}*10/3*(21/8)^{2/5}*(9)^{2/7}}\right)^{9/2}$
$\gamma - \delta$	$w_{\text{out}} = \frac{35}{2} \cdot \left(\frac{w_{\text{in}}-2}{e^{-5}*7/2*(9)^{5/7}}\right)^{9/5}$
$\delta - \varepsilon$	$w_{\text{out}} = \frac{35}{2} \cdot \left(\frac{w_{\text{in}}-3}{e^{-7}*45/2}\right)^{9/7}$
$\varepsilon - +\infty$	$w_{\text{out}} = (w_{\text{in}} - 8) \cdot e^{9}$

Figure 6: Function f_π for example with exponential observer

The proof of Proposition 14 relies on the computation of optimal exit values for the observer: letting $w_i^{\text{opt}} = \frac{p_i \cdot r_{i+1}}{r_i - r_{i+1}}$, we prove that w_i^{opt} is the optimal value of the observer with which to exit location ℓ_i (as long as time permits). The optimal time to be spent in location ℓ_i is then

$$t_i^{\text{opt}} = \frac{1}{r_i} \ln\left(\frac{w_i^{\text{opt}}}{w_{i-1}^{\text{opt}} + p_{i-1}}\right),$$

with the convention that w_0^{opt} is the initial observer value. The technical condition for being in normal form implies that these values are positive (except possibly t_1^{opt}: having t_1^{opt} negative means that the initial observer value is sufficient to go with no delay to the next location).

This allows us to compute the exact value of f_π:

Proposition 15. *The function f_π is defined on an interval $[w_0^*, \infty[$, and there is a sequence $w_0^* < w_1^* < \cdots < w_n^*$ of algebraic numbers $w_0^*, w_1^*, \ldots, w_n^* \in \mathbb{R}_{\geq 0}$ such that on each interval $[w_i^*, w_{i+1}^*]$ and on $[w_n^*, \infty[$, f_π can be obtained in closed form as*

$$f_\pi(w) = \alpha_i \cdot (w - \beta_i)^{r_i/r'_i} + \gamma_i$$

where r_i and r'_i are rates of π with $r_i \geq r'_i$, and α_i, β_i and γ_i are algebraic numbers which can be computed from the constants appearing in π. Moreover, f_π is continuous and has continuous derivative $\dot{f}_\pi(x) \geq 1$ on its domain (also at the points w_i^*).

Example 2. *We consider the same path as depicted in Example 1, but assuming an exponential observer with additive updates. This automaton satisfies the restrictions and is already in normal form for exponential observer. The resulting function f_π is depicted in Figure 6. For instance, if we enter the path with initial observer value 1, the optimal policy is to exit the location with rate 2 when the value reaches $w_1^{opt} = 10/3$ (which occurs after $\ln(5/3)/2$ time units), then leave the next location when the value reaches $w_2^{opt} = 7/2$, and then spend the remaining $1 - \ln(5/3)/2 - \ln(21/8)/5$ time units in the location with rate 7 (there is not enough time remaining to reach $w_3^{opt} = 45/2$). When we leave this location, observer value then equals $(7/2 - 1) \cdot \exp(7 \cdot (1 - \ln(5/3)/2 - \ln(21/8)/5)) \approx 21/2$, which makes it possible to fire the last transition directly and end up with final observer value around $11/2$.*

7. THE DISCRETE ABSTRACTION

In this section we show how a timed automaton with one clock and one observer with additive updates can be converted into a form in which the analysis of paths without resets from the preceding sections can be applied. After applying this, we arrive at a *finite automaton with energy functions* which we subsequently analyse in the next section.

Clock bounded above by 1. As a first step, we show that A can be converted to a timed automaton with observers where the clock is bounded above by 1 and with only three different types of edges. A similar simplification technique for priced games was used in [8].

Lemma 16. *Let A be a closed one-clock timed automaton with observers, ℓ_0 a location of A, and $L_G \subseteq L$ a set of goal locations. One can construct in exponential time another one-clock timed automaton A' with observers, together with a new initial location ℓ'_0 and a new set of goal locations L'_G, such that*

- *in any state ℓ' of A', the invariant is $c' \leq 1$;*
- *for any edge $\ell \xrightarrow{g,r} \ell'$ in A', either $r = \emptyset$ and g is the constraint $0 \leq c' \leq 1$, or $r = \{c'\}$ and g is an equality constraint $c' = 0$ or $c' = 1$,*

and such that, for any w_0 and m, $\langle A, \ell_0, w_0, L_G, m\rangle$ is a positive instance of the reachability (resp. infinite-run) problem iff $\langle A', \ell'_0, w_0, L'_G, m\rangle$ is also a positive instance of that problem.

Note that the lemma applies to automata with an arbitrary number of observers, with arbitrary updates instead of only additive ones. For the following conversions however, we have to assume a single linear or exponential observer with additive updates.

Remark 2. *It is convenient for the sequel to have global invariant $c \leq 1$ for the clock, but not strictly necessary. As mentionned at Remark 1, it is possible to handle paths with general invariant $c \leq k$. A variant of the second property of the lemma can be ensured using the coarse clock regions of [10], and using this construction, one avoids exponential blowup.*

Eliminating cycles without resets. To be able to apply our analysis of paths without resets in Sections 5 and 6, we need to ensure that there are only finitely many such paths between any two locations. Using a partial unfolding of the timed automaton where we only unfold along edges without resets, and afterwards pruning infinite reset-free paths, we can construct a timed automaton without reset-free cycles. In order to be correct, this construction must be preceded by a detection of feasible Zeno runs in the original automaton, and this information has to be stored in the unfolding. For reachability, we also have to take into account positive Zeno cycles from which some final location is reachable.

Lemma 17. *Let A be a closed one-clock timed automaton with one linear or exponential observer and additive updates, and with clock bound $c \leq 1$. Let $\ell_0 \in L$ be a location of A, and $L_G \subseteq L$ be a set of goal locations. We can compute in exponential time*

- *two labelling functions $w_{Zeno}, w_{Zeno}^{L_G} : L \to \mathbb{R} \cup \{+\infty\}$,*
- *another such automaton A', with set of locations L', and a projection $\mathsf{lab}\colon L' \to L$,*
- *a location $\ell'_0 \in L'$*
- *a set of goal locations $L'_G \subseteq L'$,*

such that A' does not contain reset-free cycles, and for any initial observer value w_0, we have the following:

- *There is an infinite feasible run in A from $(\ell_0, c = 0)$ with initial observer value w_0 if and only if there is such a run in A' from $(\ell'_0, c = 0)$, or there is a feasible run in A' from $(\ell'_0, c = 0)$ with observer value w_0 to a configuration $(\ell, c = 0)$ with observer value w, and such that $w \geq w_{Zeno}(\mathsf{lab}(\ell))$.*
- *There is a feasible run in A from $(\ell_0, c = 0)$ with initial observer value w_0 to a location in L_G if and only if there is such a run in A' from $(\ell'_0, c = 0)$ to a location in L'_G, or there is a feasible run in A' from $(\ell'_0, c = 0)$ with observer value w_0 to a configuration $(\ell, c = 0)$ with observer value $w \geq w_{Zeno}^{L_G}(\mathsf{lab}(\ell))$.*

8. AUTOMATA WITH ENERGY FUNCTIONS

We are now left with a timed automaton A' in which all cycles have at least one resetting transition. We shall construct a discrete abstraction B of A' which will contain all the information we need for solving our problem. This abstraction will be a *finite automaton with energy functions*:

Definition 18. *A* finite automaton with energy functions *is a finite transition system (S, T) equipped with a function $f\colon T \to (\mathbb{R} \rightharpoonup \mathbb{R})$ decorating transitions with energy functions. The semantics $[\![B]\!]$ of such an automaton is given by an infinite transition system with states $(s, w) \in S \times \mathbb{R}$ and transitions $(s, w) \to (s', w')$ whenever there is $e = (s, s') \in T$ such that $f(e)$ is defined in w and $f(e)(w) = w'$.*

Discrete abstraction. For the discrete abstraction of A, we take as states of B the locations of A having at least one incoming resetting transition. It is intended that state ℓ of B represents configuration $(\ell, c = 0)$ of A. For each pair of states (ℓ, ℓ') in B, there is an edge from ℓ to ℓ' iff there is a reset-free path from ℓ to ℓ' in A. We label each edge (ℓ, ℓ') of B with a (partial) function $f_{\ell,\ell'} : \mathbb{R} \to \mathbb{R}$ computed as follows: $f_{\ell,\ell'}(w)$ is the maximal achievable observer value when entering ℓ', if starting from ℓ with observer value w and visiting only reset-free paths in A'. Since there are only exponentially many such paths, $f_{\ell,\ell'}$ can be computed in exponential time, using our procedures of Sections 5 or 6.

Finally, we label states of B with their values of w_{Zeno} and $w_{\text{Zeno}}^{L_G}$ (as locations in A), obtained from Lemma 17: the values $w_{\text{Zeno}}(\ell)$ (respectively $w_{\text{Zeno}}^{L_G}(\ell)$) represent the minimal observer value needed in ℓ for which there exists a simple reset-free path from ℓ to a reset-free simple cycle with nonnegative accumulated update (respectively with positive accumulated update and from which L_G can be reached).

The following lemma directly follows from this construction and Lemma 17:

Lemma 19. *Let A be a closed one-clock timed automaton with one linear or exponential observer and additive updates, and with clock bound $c \leq 1$. Let $\ell_0 \in L$ be a location of A and $L_G \subseteq L$ a set of goal locations, and assume w.l.o.g. that locations in L_G only have resetting incoming transitions and have no outgoing transitions. Let B be the finite automaton with energy functions as constructed above and $w_0 \in \mathbb{R}+$ an initial observer value. Then*

- *there is a feasible infinite run in A from $(\ell_0, c = 0)$ with initial credit w_0 iff either there is an infinite path in $[\![B]\!]$ from (ℓ_0, w_0), or there is a finite path in $[\![B]\!]$ from (ℓ_0, w_0) to a configuration (ℓ, w) such that $w \geq w_{Zeno}(\ell)$;*
- *there is a feasible run in A reaching a location in L_G from $(\ell_0, c = 0)$ with initial credit w_0 iff either there is a finite path in $[\![B]\!]$ from (ℓ_0, w_0) to (ℓ, w) for some $\ell \in L_G$ and some w, or there is a finite path in $[\![B]\!]$ from (ℓ_0, w_0) to (ℓ, w) for some ℓ such that $w \geq w_{Zeno}^{L_G}(\ell)$.*

Energy functions. We now take advantage of the special shape of the energy functions that we get in the case of linear and exponential cases.

Definition 20. *A* rational power function *is a function of the form $f\colon x \mapsto \alpha \cdot x^r + \beta$ where r is rational and α and β are algebraic numbers.*

As a consequence of our results of Sect. 5 and 6, we have:

Lemma 21. *Let π be an annotated path (with linear or exponential observer, under the corresponding restrictions of Sections 5 and 6). Then the energy function f_π has the following property:*

($\star$) *there exists an increasing sequence $x_1 < x_2 < \ldots < x_n$ of algebraic numbers such that*
- *the domain of f_π, written $\mathrm{dom}(f_\pi)$, is $[x_1, +\infty[$;*
- *for all i, the restrictions $(f_\pi)_{\upharpoonright [x_i, x_{i+1}[}(x)$ and $(f_\pi)_{\upharpoonright [x_n, \infty[}$ are rational power functions;*
- *for all $x \geq x' \geq x_1$, it holds that $f_\pi(x) - f_\pi(x') \geq x - x'$.*

Notice that the last condition follows from the fact that $\dot{f_\pi}(x) \geq 1$. Also, functions satisfying this condition are injective.

We shall need operations of (binary) maximum and composition on functions with property ($\star$) above; these are defined in the standard way: Given partial functions f and f' with right-infinite domain,

- $\max(f, f')$ is the function with domain $\mathrm{dom}(f) \cup \mathrm{dom}(f')$ defined by $x \mapsto \max\{g(x) \mid g \in \{f, f'\}, x \in \mathrm{dom}(g)\}$;
- $f' \circ f$ is the function with domain $\mathrm{dom}(f) \cap f^{-1}(\mathrm{dom}(f'))$ defined by $x \mapsto f'(f(x))$.

LEMMA 22. *If f and f' are partial functions satisfying Property ($\star$) above, then $\max(f, f')$ and $f' \circ f$ also satisfy ($\star$).*

PROOF. Let (x_i) and (x'_i) be the corresponding sequences of algebraic numbers as of Lemma 21. Let (y_j) be the increasing sequence of algebraic numbers given by $\{y_j\} = \{x_i\} \cup \{x'_i\}$. Then $\max(f, f')$ is defined on $[y_1, +\infty)$ and is clearly a piecewise rational power function. The proof of the third property is straightforward.

For composition $f' \circ f$, we let (y_j) be the increasing sequence of algebraic numbers for which $\{y_j\} = \{f^{-1}(x'_i) \mid x'_i \geq f(x_1)\} \cup \{x_i \mid f(x_i) \geq x'_1\}$. (Note that, indeed, these numbers are roots of polynomials and hence algebraic.) Then $f' \circ f$ is a piecewise rational power function on $[y_1, +\infty)$. The third property is straightforward. □

We now study fixed points of those functions:

LEMMA 23. *Let f be a function satisfying Property ($\star$) of Lemma 21. Then*

1. *The set of fixed points of f is either empty or a left-closed interval.*
2. *Let $[x^*, x^\dagger\rangle$ be the set of fixed points of f (assuming it is not empty, and allowing $x^\dagger$ to equal $+\infty$). Then $f(x) < x$ for all $x < x^*$ and $f(x) > x$ for all $x > x^\dagger$.*
3. *If x^* exists, then for all $x \in \mathrm{dom}(f)$, there exists an infinite sequence $(f^n(x))_n$ of iterated values iff $x \geq x^*$.*

PROOF. For the first claim: From the fact that $f(x) - f(y) \geq x - y$ whenever $x \geq y$, we get that any point between two fixed points is a fixed point. Moreover, f is left-continuous, so that if $f(x) = x$ on a left-open interval, then also $f(x) = x$ at the left end-point.

For the second claim: Let $x < x^*$. Then $f(x^*) - f(x) \geq x^* - x$, which entails $f(x) \leq x$. Since x cannot be a fixed point, we must have $f(x) < x$. Similarly for the other claim.

Finally, for the third claim: For any $z \geq x^*$, we have $f(z) \geq z$. Hence if $x \geq x^*$, then for any k such that $f^k(x)$ is defined, we have $f^k(x) \geq x^*$, so that $f^{k+1}(x)$ is defined. On the other hand, assume that $x < x^*$, and that the infinite sequence $(f^n(x))_n$ is defined. Then $f(x) < x$ and, by induction, $f^{n+1}(x) < f^n(x)$ for all n. The sequence being decreasing and bounded by x_1, it converges. As f is left-continuous, the limit $\bar{x}$ satisfies $f(\bar{x}) = \bar{x}$, which contradicts the fact that x^* is the smallest fixed point. □

Algorithm. We now gather everything together in order to solve Problems 1 and 2. We first explain how we detect feasible non-Zeno runs of A: Assume that such a run exists in A, and let ρ be the corresponding infinite run in B. Then some simple cycle σ in ρ must be repeated infinitely often. For all $i \geq 0$, write w_{2i} for the observer value when entering the $(i+1)$-st occurrence of σ, and w_{2i+1} for the value when exiting the $(i+1)$-st occurrence.

Assume that $w_{2i+1} < w_{2i}$ for all i, and that either $w_{2i+2} < w_{2i+1}$, or $w_{2i+2} = w_{2i+1}$ and the $(i+2)$-nd occurrence of σ directly follows the $(i+1)$-st one. Notice that it cannot be the case that we are in the latter situation for all i, as this would give an infinite iteration $f^n(w)$ for an energy function with $f(w) < w$, contradicting Lemma 23. Hence there are two occurrences of σ having a non-empty sub-path in-between, and this sub-path has negative effect on observer value. Dropping the earliest such path yields a feasible infinite non-Zeno run ρ'. This procedure can be repeated recursively, as long as the sequence (w_i) satisfies the condition above.

We first assume that the sequence $(w_i)_i$ is decreasing. This means that between the first and second occurrences of σ, observer value has decreased. Then this part between the first two occurrences of σ can be dropped, yielding a new feasible infinite non-Zeno run. Apply this procedure recursively, as long as the sequence $(w_i)_i$ in the resulting run is decreasing. This yields a sequence of feasible runs $(\rho_n)_n$ of the form $\rho_0 \cdot (\sigma)^n \cdot \pi_n$. There must exist an index at which the procedure cannot be repeated, since otherwise it would contradict Lemma 23. At that point, we end up with a run in which a simple cycle has a positive effect on observer value. Following Lemma 23, this cycle can be iterated from that point on, yielding a feasible lasso-shaped non-Zeno run.

Now, consider the non-periodic part of this run: if it contains a (simple) cycle with negative effect, then again we can drop this cycle while obtaining a feasible lasso-shaped infinite non-Zeno run. On the other hand, if the cycle has nonnegative effect, it can be iterated itself. In the end, we have proved the following lemma:

LEMMA 24. *If there is a feasible infinite non-Zeno run in A from $(\ell_0, c = 0)$ with observer value w_0, then there is a lasso-shaped one in which the initial part is acyclic and the periodic part is a simple cycle. In particular, both parts have linear size.*

Now consider the case of feasible Zeno runs in A. Such a run corresponds to a finite run ρ in B, ending in a location ℓ with observer value at least $w_{\mathrm{Zeno}}(\ell)$. Using similar arguments as above, if there is a simple cycle in ρ with nonnegative effect, then we can deduce a feasible infinite run in B, hence in A. If the effect of the cycle is negative, then dropping the cycle yields another feasible Zeno run. In the end, we have

LEMMA 25. *If there is a feasible infinite Zeno run in A from $(\ell_0, c = 0)$ with observer value w_0, then either there is an acyclic path in B reaching a configuration (ℓ, w) with $w \geq w_{Zeno}(\ell)$, or there is a lasso-shaped run in which the initial part is acyclic and the periodic part is a simple cycle.*

As B has size linear in A, we can enumerate in exponential time the possible witnesses given by the above two lemmas. Since constructing B is already in exponential time, our global procedure is in exponential time. For reachability, similar techniques give the following lemma:

LEMMA 26. *If there is a feasible run in A from $(\ell_0, c = 0)$ with observer value w_0 to a goal location in L_G, then there is an acyclic path in B reaching a configuration (ℓ, w) with $\ell \in L_G$ or $w \geq w^{L_G}_{Zeno}(\ell)$.*

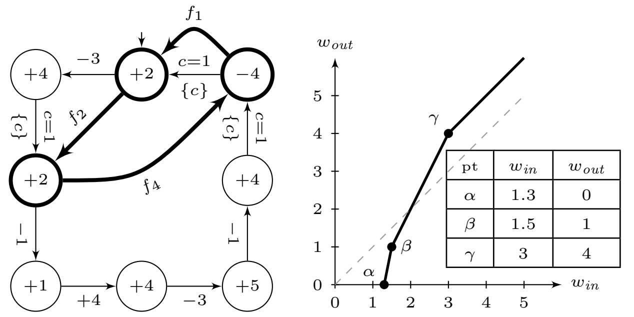

pt	w_{in}	w_{out}
α	1.3	0
β	1.5	1
γ	3	4

Figure 7: Solving the infinite-run problem for an example timed automaton

EXAMPLE 3. *The simple linearly priced timed automaton in Figure 7 is a collection of some of the paths we have seen earlier. Specifically, we have taken the paths from Figures 2 and 4 and connected them using a trivial path with one location with rate* -4. *The figure also displays the discrete abstraction of the automaton, consisting of three states and three transitions labeled with energy functions* f_2 *for the path from Figure 2,* f_4 *for the path from Figure 4, and* f_1 *for the trivial path. Note that no Zeno runs are possible in the automaton, hence the* w_{Zeno} *annotation has been omitted.*

To compute the minimum observer value necessary for having an infinite run in the example automaton, we need to find the least fixed point of the energy function $f_1 \circ f_4 \circ f_2$ *which is displayed in the right part of the figure. This point can be computed to be at* $w_{in} = 2$, *hence the automaton admits an infinite run if and only if initial observer value is at least* 2.

9. CONCLUSION AND FUTURE WORK

We have shown that the infinite-runs and reachability problems are decidable for closed one-clock timed automata with one linear or exponential observer and additive updates, using a novel technique of energy functions. We expect this technique, and the general notion of finite automata with energy functions, to have applications in other areas as well.

The restriction to a one-clock setting is essential in our approach, and currently we do not know how to extend it to timed automata with more than one clock. However one-clock models are expressive enough to model a large class of interesting examples, and in the context of priced timed automata, it is well-known that models with more than one clock are very difficult to handle.

The restriction to closed timed automata is not essential. It was mainly adapted to ease exposition, and can be lifted by during the analysis carefully taking note of which values of energy functions can actually be attained. On another note, our technique may also apply to the lower-soft-upper-bound problem mentioned in the introduction. We have concentrated on solving lower-bound problems here, but by adapting our analysis of paths, we may also solve the problem with soft upper bound. The interval-bound problem seems however much more difficult.

Considering more general observers than only linear or exponential ones is also of interest. The present work is part of a general project concerning "hybridization" of timed-automata technology, and linear and exponential flow conditions combined cover the whole class of first-order differentiable observers. We can easily handle observers which can have either linear or exponential behaviour, but the general first-order differential case is more difficult, partly because our technique relies on flows being monotonous, either increasing or decreasing, in locations.

10. REFERENCES

[1] R. Alur, C. Courcoubetis, N. Halbwachs, Th. A. Henzinger, P.-H. Ho, X. Nicollin, A. Olivero, J. Sifakis, and S. Yovine. The algorithmic analysis of hybrid systems. *Theoretical Computer Science*, 138(1):3–34, 1995.

[2] R. Alur, C. Courcoubetis, and Th. A. Henzinger. Computing accumulated delays in real-time systems. In *Proc. 5th International Conference on Computer Aided Verification (CAV'93)*, volume 697 of *Lecture Notes in Computer Science*, pages 181–193. Springer, 1993.

[3] R. Alur, S. La Torre, and G. J. Pappas. Optimal paths in weighted timed automata. In *Proc. 4th International Workshop on Hybrid Systems: Computation and Control (HSCC'01)*, volume 2034 of *Lecture Notes in Computer Science*, pages 49–62. Springer, 2001.

[4] R. Alur, S. La Torre, and G. J. Pappas. Optimal paths in weighted timed automata. *Theoretical Computer Science*, 318(3):297–322, 2004.

[5] G. Behrmann, A. Fehnker, Th. Hune, K. G. Larsen, P. Pettersson, J. Romijn, and F. Vaandrager. Minimum-cost reachability for priced timed automata. In *Proc. 4th International Workshop on Hybrid Systems: Computation and Control (HSCC'01)*, volume 2034 of *Lecture Notes in Computer Science*, pages 147–161. Springer, 2001.

[6] P. Bouyer, U. Fahrenberg, K. G. Larsen, and N. Markey. Timed automata with observers under energy constraints. Research Report LSV-10-01, Laboratoire Spécification et Vérification, ENS Cachan, France, 2010.

[7] P. Bouyer, U. Fahrenberg, K. G. Larsen, N. Markey, and J. Srba. Infinite runs in weighted timed automata with energy constraints. In *Proc. 6th International Conference on Formal Modeling and Analysis of Timed Systems (FORMATS'08)*, Lecture Notes in Computer Science, pages 33–47. Springer, 2008.

[8] P. Bouyer, K. G. Larsen, N. Markey, and J. I. Rasmussen. Almost optimal strategies in one-clock priced timed automata. In *Proc. 26th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'06)*, volume 4337 of *Lecture Notes in Computer Science*, pages 345–356. Springer, 2006.

[9] Y. Kesten, A. Pnueli, J. Sifakis, and S. Yovine. Decidable integration graphs. *Information and Computation*, 150(2):209–243, 1999.

[10] F. Laroussinie, N. Markey, and Ph. Schnoebelen. Model checking timed automata with one or two clocks. In *Proc. 15th International Conference on Concurrency Theory (CONCUR'04)*, volume 3170 of *Lecture Notes in Computer Science*, pages 387–401. Springer, 2004.