Model-Based Verification, Optimization, Synthesis and Performance Evaluation of Real-Time Systems

Uli Fahrenberg, Kim Guldstrand Larsen, Axel Legay, Claus Thrane

► To cite this version:

Uli Fahrenberg, Kim Guldstrand Larsen, Axel Legay, Claus Thrane. Model-Based Verification, Optimization, Synthesis and Performance Evaluation of Real-Time Systems. Manfred Broy; Doron Peled; Georg Kalus. Engineering Dependable Software Systems, 34, IOS Press, pp.67 - 108, 2013, NATO Science for Peace and Security Series - D: Information and Communication Security, 978-1-61499-206-6. 10.1007/978-3-642-39721-9_2 . hal-01087921

HAL Id: hal-01087921

<https://hal.inria.fr/hal-01087921>

Submitted on 27 Nov 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Model-Based Verification, Optimization, Synthesis and Performance Evaluation of Real-Time Systems

Uli Fahrenberg^a, Kim G. Larsen^{b,1}, Axel Legay^a Claus Thrane^b

^a *Irisa / INRIA Rennes, France*

^b *Department of Computer Science, Aalborg University, Denmark*

Abstract. This article aims at providing a concise and precise *Travellers Guide, Phrase Book or Reference Manual* to the timed automata modeling formalism introduced by Alur and Dill [8,9]. The paper gives comprehensive definitions of timed automata, priced (or weighted) timed automata, and timed games and highlights a number of results on associated decision problems related to model checking, equivalence checking, optimal scheduling, the existence of winning strategies, and then statistical model checking.

Keywords. Timed automaton, region, zone, reachability, bisimilarity; priced timed automaton, weighted timed automaton, optimal reachability, optimal infinite run, conditional optimality; timed game, winning strategy.

1. Introduction

The model of timed automata, introduced by Alur and Dill [8, 9], has by now established itself as a classical formalism for describing the behaviour of real-time systems. A number of important algorithmic problems has been shown decidable for it, including reachability, model checking and several behavioural equivalences and preorders.

By now, real-time model checking tools such as UPPAAL [20, 75] and KRONOS [37] are based on the timed automata formalism and on the substantial body of research on this model that has been targeted towards transforming the early results into practically efficient algorithms — *e.g.* [16, 17, 22, 24] — and data structures — *e.g.* [23, 72, 74].

The maturity of a tool like UPPAAL is witnessed by the numerous applications — *e.g.* [45, 52, 61, 65, 70, 73, 78, 79] — to the verification of industrial case-studies spanning real-time controllers and real-time communication protocols. More recently, model-checking tools in general and UPPAAL in particular have been ap-

¹Corresponding Author: Kim G. Larsen, Department of Computer Science, Aalborg University, Selma Lagerlöfs Vej 300, 9220 Aalborg Øst, Denmark. E-mail: e-mail: kgl@cs.aau.dk

plied to solve realistic scheduling problems by a reformulation as reachability problems — *e.g.* [1, 58, 64, 80].

Aiming at providing methods for performance analysis, a recent extension of timed automata is that of *priced* or *weighted* timed automata [10, 21], which makes it possible to formulate and solve *optimal* scheduling problems. Surprisingly, a number of properties have been shown to be decidable for this formalism [10, 21, 34, 53, 76]. The recently developed UPPAAL-CORA tool provides an efficient tool for solving cost-optimal reachability problems [71] and has been applied successfully to a number of optimal scheduling problems, *e.g.* [18, 25, 60].

Most recently, substantial efforts have been made on the automatic synthesis of (correct-by-construction) controllers from timed games for given control objectives. From early decidability results [13, 82] the effort has lead to efficient on-the-fly algorithms [41, 92] with the newest of the UPPAAL toolset, UPPAAL-TIGA [19], UPPAAL-SMC [48, 49], providing an efficient tool implementation with industrial applications emerging, *e.g.* [67].

This survey paper aims at providing a concise and precise *Travellers Guide*, *Phrase Book* or *Reference Manual* to the land and language of timed automata. The article gives comprehensive definitions of timed automata, weighted timed automata, and timed games and highlights a number of results on associated decision problems related to model checking, equivalence checking, optimal scheduling, the existence of winning strategies, and statistical model checking. The intention is that the paper should provide an easy-to-access collection of important results and overview of the field to anyone interested.

The authors would like to thank the students of the Marktoberdorf and Quantitative Model Checking PhD schools for their useful comments and help in weeding out a number of errors in the first two editions of this survey [55, 57], as well as an anonymous reviewer who provided many useful remarks for the invited paper [56] at FSEN 2009.

2. Timed automata

In this section we review the notion of timed automata introduced by Alur and Dill [8, 9] as a formalism for describing the behaviour of real-time systems. We review the syntax and semantics and highlight the, by now classical, region construction underlying the decidability of several associated problems.

Here we illustrate how regions are applied in showing decidability of reachability and timed and untimed (bi)similarity. However, the notion of region does not provide the means for efficient tool implementations. The verification engine of UPPAAL instead applies so-called zones, which are *convex unions* of regions. We give a brief account of zones as well as their efficient representation and manipulation using difference-bound matrices.

2.1. Syntax and semantics

Definition 2.1. The set $\Phi(C)$ of *clock constraints* φ over a finite set (of *clocks*) C is defined by the grammar

Figure 1.: A light switch modelled as a timed automaton.

$$\varphi ::= x \bowtie k \mid \varphi_1 \wedge \varphi_2 \quad (x \in C, k \in \mathbb{Z}, \bowtie \in \{\leq, <, \geq, >\}).$$

The set $\Phi^+(C)$ of *extended clock constraints* φ is defined by the grammar

$$\varphi ::= x \bowtie k \mid x - y \bowtie k \mid \varphi_1 \wedge \varphi_2 \quad (x, y \in C, k \in \mathbb{Z}, \bowtie \in \{\leq, <, \geq, >\}).$$

Remark 2.2. The clock constraints in $\Phi(C)$ above are also called *diagonal-free* clock constraints, and the additional ones in $\Phi^+(C)$ are called *diagonal*. We restrict ourselves to diagonal-free clock constraints here; see Remark 2.44 for one reason. For additional modelling power, timed automata with diagonal constraints can be used, as it is shown in [9, 29] that any such automaton can be converted to a diagonal-free one; however the conversion may lead to an exponential blow-up.

Definition 2.3. A *timed automaton* is a tuple $(L, \ell_0, F, C, \Sigma, I, E)$ consisting of a finite set L of locations, an initial location $\ell_0 \in L$, a set $F \subseteq L$ of final locations, a finite set C of clocks, a finite set Σ of actions, a location invariants mapping $I : L \rightarrow \Phi(C)$, and a set $E \subseteq L \times \Phi(C) \times \Sigma \times 2^C \times L$ of edges.

Here 2^C denotes the set of subsets (*i.e.* the power set) of C . We shall write $\ell \xrightarrow{\varphi, a, r} \ell'$ for an edge $(\ell, \varphi, a, r, \ell') \in E$. In figures, resets are written as assignment to zero, *e.g.* $x := 0$.

Example 2.1. Figure 1 provides a timed automaton model of an intelligent light switch. Starting in the “Off” state, a press of the button turns the light on, and it remains in this state for 100 time units (*i.e.* until clock $x = 100$), at which time

the light turns off again. During this time, an additional press resets the clock x and prolongs the time in the state by 100 time units. Pressing the button twice, with at most three time units between the presses, triggers a special bright light.

Definition 2.4. A *clock valuation* on a finite set C of clocks is a mapping $v : C \rightarrow \mathbb{R}_{\geq 0}$. The *initial* valuation v_0 is given by $v_0(x) = 0$ for all $x \in C$. For a valuation v , $d \in \mathbb{R}_{\geq 0}$, and $r \subseteq C$, the valuations $v + d$ and $v[r]$ are defined by

$$(v + d)(x) = v(x) + d$$

$$v[r](x) = \begin{cases} 0 & \text{for } x \in r, \\ v(x) & \text{for } x \notin r. \end{cases}$$

Extending the notation for power set introduced above, we will in general write B^A for the set of mappings from a set A to a set B . The set of clock valuations on C is thus $\mathbb{R}_{\geq 0}^C$.

Definition 2.5. The *zone* of an extended clock constraint in $\Phi^+(C)$ is a set of clock valuations $C \rightarrow \mathbb{R}_{\geq 0}$ given inductively by

$$\llbracket x \bowtie k \rrbracket = \{v : C \rightarrow \mathbb{R}_{\geq 0} \mid v(x) \bowtie k\},$$

$$\llbracket x - y \bowtie k \rrbracket = \{v : C \rightarrow \mathbb{R}_{\geq 0} \mid v(x) - v(y) \bowtie k\}, \text{ and}$$

$$\llbracket \varphi_1 \wedge \varphi_2 \rrbracket = \llbracket \varphi_1 \rrbracket \cap \llbracket \varphi_2 \rrbracket.$$

We shall write $v \models \varphi$ instead of $v \in \llbracket \varphi \rrbracket$.

Definition 2.6. The *semantics* of a timed automaton $A = (L, \ell_0, F, C, \Sigma, I, E)$ is the transition system $\llbracket A \rrbracket = (S, s_0, \Sigma \cup \mathbb{R}_{\geq 0}, T = T_s \cup T_d)$ given as follows:

$$S = \{(\ell, v) \in L \times \mathbb{R}_{\geq 0}^C \mid v \models I(\ell)\} \quad s_0 = (\ell_0, v_0)$$

$$T_s = \{(\ell, v) \xrightarrow{a} (\ell', v') \mid \exists \ell \xrightarrow{\varphi, a, r} \ell' \in E : v \models \varphi, v' = v[r]\}$$

$$T_d = \{(\ell, v) \xrightarrow{d} (\ell, v + d) \mid \forall d' \in [0, d] : v + d' \models I(\ell)\}$$

Remark 2.7. The transition system $\llbracket A \rrbracket$ from above is an example of what is known as a *timed transition system*, *i.e.* a transition system where the label set includes $\mathbb{R}_{\geq 0}$ as a subset and which satisfies certain additivity and time determinacy properties. We refer to [2] for a more in-depth treatment.

Also note that the semantics $\llbracket A \rrbracket$ contains no information about final states (derived from the final locations in F); this is mostly for notational convenience.

Definition 2.8. A (finite) *run* of a timed automaton $A = (L, \ell_0, F, C, \Sigma, I, E)$ is a finite path $\rho = (\ell_0, v_0) \rightarrow \dots \rightarrow (\ell_k, v_k)$ in $\llbracket A \rrbracket$. It is said to be *accepting* if $\ell_k \in F$.

Figure 2.: A timed automaton with two clocks.

Example 2.1 (continued). The light switch model from figure 1 has as state set

$$S = \{\text{Off}\} \times \mathbb{R}_{\geq 0} \cup \{\text{Light}, \text{Bright}\} \times [0, 100]$$

where we identify valuations with their values at x . A few example runs are given below; we abbreviate “press?” to “p”:

$$\begin{aligned} &(\text{Off}, 0) \xrightarrow{150} (\text{Off}, 150) \xrightarrow{\text{p}} (\text{Light}, 0) \xrightarrow{100} (\text{Light}, 100) \rightarrow (\text{Off}, 0) \\ &(\text{Off}, 0) \xrightarrow{\text{p}} (\text{Light}, 0) \xrightarrow{10} (\text{Light}, 10) \xrightarrow{\text{p}} (\text{Light}, 0) \xrightarrow{100} (\text{Light}, 100) \rightarrow (\text{Off}, 0) \\ &(\text{Off}, 0) \xrightarrow{\text{p}} (\text{Light}, 0) \xrightarrow{1} (\text{Light}, 1) \xrightarrow{\text{p}} (\text{Bright}, 0) \xrightarrow{100} (\text{Bright}, 100) \rightarrow (\text{Off}, 0) \end{aligned}$$

2.2. Reachability

We are concerned with the following problem: Given a timed automaton $A = (L, \ell_0, F, C, \Sigma, I, E)$, is any of the locations in F reachable? We shall later define the *timed language* generated by a timed automaton and see that this reachability problem is equivalent to *emptiness checking*: Is the timed language generated by A non-empty?

Example 2.2 (cf. [2, Ex. 11.7]). Figure 2 shows a timed automaton A with two clocks and a final location ℓ_1 . To ask whether ℓ_1 is reachable amounts for this automaton to the question whether there is a finite sequence of a - and b -transitions from ℓ_0 which brings clock values into accordance with the guard $x \geq 4 \wedge y \leq 2$ on the edge leading to ℓ_1 .

An immediate obstacle to reachability checking is the infinity of the state space of A . In general, the transition system $\llbracket A \rrbracket$ has uncountably many states, hence straight-forward reachability algorithms do not work for us.

Notation 2.9. The *derived transition relations* in a timed automaton $A = (L, \ell_0, F, C, \Sigma, I, E)$ are defined as follows: For $(\ell, v), (\ell', v')$ states in $\llbracket A \rrbracket$, we say that

- $(\ell, v) \xrightarrow{\delta} (\ell', v')$ if $(\ell, v) \xrightarrow{d} (\ell', v')$ in $\llbracket A \rrbracket$ for some $d > 0$,

- $(\ell, v) \xrightarrow{\alpha} (\ell', v')$ if $(\ell, v) \xrightarrow{a} (\ell', v')$ in $\llbracket A \rrbracket$ for some $a \in \Sigma$, and
- $(\ell, v) \rightsquigarrow (\ell', v')$ if $(\ell, v) (\xrightarrow{\delta} \cup \xrightarrow{\alpha})^* (\ell', v')$.

Definition 2.10. The set of *reachable locations* in a timed automaton $A = (L, \ell_0, F, C, \Sigma, I, E)$ is

$$\text{Reach}(A) = \{\ell \in L \mid \exists v : C \rightarrow \mathbb{R}_{\geq 0} : (\ell_0, v_0) \rightsquigarrow (\ell, v)\}.$$

Hence we can now state the reachability problem as follows:

Problem 2.1 (Reachability). Given a timed automaton $A = (L, \ell_0, F, C, \Sigma, I, E)$, is $\text{Reach}(A) \cap F \neq \emptyset$?

Definition 2.11. Let $A = (L, \ell_0, F, C, \Sigma, I, E)$ be a timed automaton. A relation $R \subseteq L \times \mathbb{R}_{\geq 0}^C \times L \times \mathbb{R}_{\geq 0}^C$ is a *time-abstracted simulation* provided that for all $(\ell_1, v_1) R (\ell_2, v_2)$,

- for all $(\ell_1, v_1) \xrightarrow{\delta} (\ell'_1, v'_1)$ there exists some (ℓ'_2, v'_2) such that $(\ell'_1, v'_1) R (\ell'_2, v'_2)$ and $(\ell_2, v_2) \xrightarrow{\delta} (\ell'_2, v'_2)$, and
- for all $a \in \Sigma$ and $(\ell_1, v_1) \xrightarrow{a} (\ell'_1, v'_1)$, there exists some (ℓ'_2, v'_2) such that $(\ell'_1, v'_1) R (\ell'_2, v'_2)$ and $(\ell_2, v_2) \xrightarrow{a} (\ell'_2, v'_2)$.

R is said to be *F-sensitive* if additionally, $(\ell_1, v_1) R (\ell_2, v_2)$ implies that $\ell_1 \in F$ if and only if $\ell_2 \in F$. A *time-abstracted bisimulation* is a time-abstracted simulation which is also symmetric; we write $(\ell_1, v_1) \approx (\ell_2, v_2)$ whenever $(\ell_1, v_1) R (\ell_2, v_2)$ for a time-abstracted bisimulation R .

Note that $\approx$ is itself a time-abstracted bisimulation, which is easily shown to be an equivalence relation and hence symmetric, reflexive and transitive. Observe also that a time-abstracted (bi)simulation on A is the same as a standard (bi)simulation on the transition system derived from $\llbracket A \rrbracket$ with transitions $\xrightarrow{\delta}$ and $\xrightarrow{a}$. Likewise, the quotient introduced below is just the standard bisimulation quotient of this derived transition system.

Definition 2.12. Let $A = (L, \ell_0, F, C, \Sigma, I, E)$ be a timed automaton and $R \subseteq L \times \mathbb{R}_{\geq 0}^C \times L \times \mathbb{R}_{\geq 0}^C$ a time-abstracted bisimulation which is also an equivalence. The *quotient* of $\llbracket A \rrbracket = (S, s_0, \Sigma \cup \mathbb{R}_{\geq 0}, T)$ with respect to R is the transition system $\llbracket A \rrbracket_R = (S_R, s_R^0, \Sigma \cup \{\delta\}, T_R)$ given by $S_R = S/R$, $s_R^0 = [s_0]_R$, and with transitions

- $\pi \xrightarrow{\delta} \pi'$ whenever $(\ell, v) \xrightarrow{\delta} (\ell', v')$ for some $(\ell, v) \in \pi$, $(\ell', v') \in \pi'$, and
- $\pi \xrightarrow{a} \pi'$ whenever $(\ell, v) \xrightarrow{a} (\ell', v')$ for some $(\ell, v) \in \pi$, $(\ell', v') \in \pi'$.

The following proposition expresses that F -sensitive quotients are sound and complete with respect to reachability.

Proposition 2.13 ([5]). Let $A = (L, \ell_0, F, C, \Sigma, I, E)$ be a timed automaton, $R \subseteq L \times \mathbb{R}_{\geq 0}^C \times L \times \mathbb{R}_{\geq 0}^C$ an F -sensitive time-abstracted bisimulation and $\ell \in F$. Then

Figure 3.: Time-abstracted bisimulation classes for the two-clock timed automaton from Example 2.2. Left: equivalence classes for switch transitions only; right: equivalence classes for switch and delay transitions.

$\ell \in \text{Reach}(A)$ if and only if there is a reachable state π in $\llbracket A \rrbracket_R$ and $v : C \rightarrow \mathbb{R}_{\geq 0}$ such that $(\ell, v) \in \pi$.

Example 2.2 (continued). We shall now try to construct, in a naïve way, a time-abstracted bisimulation R for the timed automaton A from Figure 2 which is as coarse as possible. Note first that we cannot have $(\ell_0, v) R (\ell_1, v')$ for any $v, v' : C \rightarrow \mathbb{R}_{\geq 0}$ because $\ell_1 \in F$ and $\ell_0 \notin F$. On the other hand it is easy to see that we can let $(\ell_1, v) R (\ell_1, v')$ for all $v, v' : C \rightarrow \mathbb{R}_{\geq 0}$, which leaves us with constructing R on the states involving ℓ_0 .

We handle switch transitions $\xrightarrow{\alpha}$ first: If $v, v' : C \rightarrow \mathbb{R}_{\geq 0}$ are such that $v(y) \leq 2$ and $v'(y) > 2$, the state (ℓ_0, v) has an a -transition available while the state (ℓ_0, v') has not, hence these cannot be related in R . Similarly we have to distinguish states (ℓ_0, v) from states (ℓ_0, v') where $v(x) \leq 2$ and $v'(x) > 2$ because of b -transitions, and states (ℓ_0, v) from states (ℓ_0, v') where $v(x) < 4$ and $v'(x) \geq 4$ because of c -transitions. Altogether this gives the five classes depicted to the left of Figure 3, where the shading indicates to which class the boundary belongs, and we have written the set of available actions in the classes.

When also taking delay transitions $\xrightarrow{\delta}$ into account, one has to partition the state space further: From a valuation v in the class marked $\{a, b\}$ in the left of the figure, a valuation in the class marked $\{a\}$ can only be reached by a delay transition if $v(y) < v(x)$; likewise, from the $\{a\}$ class, the $\{a, c\}$ class can only be reached if $v(y) \leq v(x) - 2$. Hence these two classes need to be partitioned as shown to the right of Figure 3.

It can easily be shown that no further partitioning is needed, thus we have defined the coarsest time-abstracted bisimulation relation for A , altogether with eight equivalence classes.

2.3. Regions

Motivated by the construction in the example above, we now introduce a time-abstracted bisimulation with a *finite quotient*. To ensure finiteness, we need the

maximal constants to which respective clocks are compared in the invariants and guards of a given timed automaton. These may be defined as follows.

Definition 2.14. For a finite set C of clocks, the *maximal constant* mapping $c_{\max} : C \rightarrow \mathbb{Z}^{\Phi(C)}$ is defined inductively as follows:

$$c_{\max}(x)(y \bowtie k) = \begin{cases} k & \text{if } y = x \\ 0 & \text{if } y \neq x \end{cases}$$

$$c_{\max}(x)(\varphi_1 \wedge \varphi_2) = \max(c(x)(\varphi_1), c(x)(\varphi_2))$$

For a timed automaton $A = (L, \ell_0, F, C, \Sigma, I, E)$, the maximal constant mapping is $c_A : C \rightarrow \mathbb{Z}$ defined by

$$c_A(x) = \max \{ c_{\max}(x)(I(\ell)), c_{\max}(x)(\varphi) \mid \ell \in L, \ell \xrightarrow{\varphi, a, r} \ell' \in E \}.$$

Notation 2.15. For $d \in \mathbb{R}_{\geq 0}$ we write $\lfloor d \rfloor$ and $\langle d \rangle$ for the integral, respectively fractional, part of d , so that $d = \lfloor d \rfloor + \langle d \rangle$.

Definition 2.16. For a timed automaton $A = (L, \ell_0, F, C, \Sigma, I, E)$, valuations $v, v' : C \rightarrow \mathbb{R}_{\geq 0}$ are said to be *region equivalent*, denoted $v \cong v'$, if

- $\lfloor v(x) \rfloor = \lfloor v'(x) \rfloor$ or $v(x), v'(x) > c_A(x)$, for all $x \in C$, and
- $\langle v(x) \rangle = 0$ iff $\langle v'(x) \rangle = 0$, for all $x \in C$, and
- $\langle v(x) \rangle \leq \langle v(y) \rangle$ iff $\langle v'(x) \rangle \leq \langle v'(y) \rangle$ for all $x, y \in C$.

Proposition 2.17 ([5]). For a timed automaton $A = (L, \ell_0, F, C, \Sigma, I, E)$, the equivalence relation $\cong$ defined on states of $\llbracket A \rrbracket$ by $(\ell, v) \cong (\ell', v')$ if $\ell = \ell'$ and $v \cong v'$ is an F -sensitive time-abstracted bisimulation. The quotient $\llbracket A \rrbracket_{\cong}$ is finite.

The equivalence classes of valuations of A with respect to $\cong$ are called *regions*, and the quotient $\llbracket A \rrbracket_{\cong}$ is called the *region automaton* associated with A .

Proposition 2.18 ([9]). The number of regions for a timed automaton A with a set C of n clocks is bounded above by

$$n! \cdot 2^n \cdot \prod_{x \in C} (2c_A(x) + 2).$$

Example 2.2 (continued). The 69 regions of the timed automaton A from Figure 2 are depicted in Figure 4.

Propositions 2.13 and 2.17 together now give the decidability part of the theorem below; for PSPACE-completeness see [7, 44].

Theorem 2.19. The reachability problem for timed automata is PSPACE-complete.

Figure 4.: Clock regions for the timed automaton from Example 2.2.

2.4. Behavioural refinement relations

We have already introduced time-abstracted simulations and bisimulations in Definition 2.11. As a corollary of Proposition 2.17, these are decidable:

Theorem 2.20. *Time-abstracted simulation and bisimulation are decidable for timed automata.*

Proof: One only needs to see that time-abstracted (bi)simulation in the timed automaton is the same as ordinary (bi)simulation in the associated region automaton; indeed, any state in $\llbracket A \rrbracket$ is untimed bisimilar to its image in $\llbracket A \rrbracket_{\cong}$. The result follows by finiteness of the region automaton. $\square$

The following provides a *time-sensitive* variant of (bi)simulation.

Definition 2.21. Let $A = (L, \ell_0, F, C, \Sigma, I, E)$ be a timed automaton. A relation $R \subseteq L \times \mathbb{R}_{\geq 0}^C \times L \times \mathbb{R}_{\geq 0}^C$ is a *timed simulation* provided that for all $(\ell_1, v_1) R (\ell_2, v_2)$,

- for all $(\ell_1, v_1) \xrightarrow{d} (\ell'_1, v'_1)$, $d \in \mathbb{R}_{\geq 0}$, there exists some (ℓ'_2, v'_2) such that $(\ell'_1, v'_1) R (\ell'_2, v'_2)$ and $(\ell_2, v_2) \xrightarrow{d} (\ell'_2, v'_2)$, and
- for all $(\ell_1, v_1) \xrightarrow{a} (\ell'_1, v'_1)$, $a \in \Sigma$, there exists some (ℓ'_2, v'_2) such that $(\ell'_1, v'_1) R (\ell'_2, v'_2)$ and $(\ell_2, v_2) \xrightarrow{a} (\ell'_2, v'_2)$.

A *timed bisimulation* is a timed simulation which is also symmetric, and two states $(\ell_1, v_1), (\ell_2, v_2) \in \llbracket A \rrbracket$ are said to be *timed bisimilar*, written $(\ell_1, v_1) \sim (\ell_2, v_2)$, if there exists a timed bisimulation R for which $(\ell_1, v_1) R (\ell_2, v_2)$.

Note that $\sim$ is itself a timed bisimulation on A , which is easily shown to be an equivalence relation and hence transitive, reflexive and symmetric.

Definition 2.22. Two timed automata $A = (L^A, \ell_0^A, F^A, C^A, \Sigma^A, I^A, E^A)$ and $B = (L^B, \ell_0^B, F^B, C^B, \Sigma^B, I^B, E^B)$ are said to be *timed bisimilar*, denoted $A \sim B$, if $(\ell_0^A, v_0) \sim (\ell_0^B, v_0)$ in the disjoint-union transition system $\llbracket A \rrbracket \sqcup \llbracket B \rrbracket$.

Timed simulation of timed automata can be analogously defined. The following decidability result was established for *parallel timed processes* in [43]; below we give a version of the proof which has been adapted for timed automata.

Theorem 2.23. *Timed similarity and bisimilarity are decidable for timed automata.*

Before the proof, we need a few auxiliary definitions and lemmas. The first is a product of timed transition systems which synchronizes on time, but not on actions:

Definition 2.24. The *independent product* of the timed transition systems $\llbracket A \rrbracket = (S^A, s_0^A, \Sigma^A \cup \mathbb{R}_{\geq 0}, T^A)$, $\llbracket B \rrbracket = (S^B, s_0^B, \Sigma^B \cup \mathbb{R}_{\geq 0}, T^B)$ associated with timed automata A, B is $\llbracket A \rrbracket \times \llbracket B \rrbracket = (S, s_0, \Sigma^A \cup \Sigma^B \cup \mathbb{R}_{\geq 0}, T)$ given by

$$\begin{aligned} S &= S^A \times S^B & s_0 &= (s_0^A, s_0^B) \\ T &= \{(p, q) \xrightarrow{a} (p', q') \mid a \in \Sigma, p \xrightarrow{a} p' \in T^A\} \\ &\cup \{(p, q) \xrightarrow{b} (p, q') \mid b \in \Sigma, q \xrightarrow{b} q' \in T^B\} \\ &\cup \{(p, q) \xrightarrow{d} (p', q') \mid d \in \mathbb{R}_{\geq 0}, p \xrightarrow{d} p' \in T^A, q \xrightarrow{d} q' \in T^B\} \end{aligned}$$

We need to extend region equivalence $\cong$ to the independent product. Below, $\oplus$ denotes vector concatenation (direct sum); note that $(p_1, q_1) \cong (p_2, q_2)$ is not the same as $p_1 \cong p_2$ and $q_1 \cong q_2$, as fractional orderings $\langle x^A \rangle \bowtie \langle x^B \rangle$, for $x^A \in C^A$, $x^B \in C^B$, have to be accounted for in the former, but not in the latter. Hence $(p_1, q_1) \cong (p_2, q_2)$ implies $p_1 \cong p_2$ and $q_1 \cong q_2$, but not vice-versa.

Definition 2.25. For states $p_i = (\ell^{p_i}, v^{p_i})$ in $\llbracket A \rrbracket$ and $q_i = (\ell^{q_i}, v^{q_i})$ in $\llbracket B \rrbracket$ for $i = 1, 2$, we say that $(p_1, q_1) \cong (p_2, q_2)$ iff $\ell^{p_1} = \ell^{p_2} \wedge \ell^{q_1} = \ell^{q_2}$ and $v^{p_1} \oplus v^{q_1} \cong v^{p_2} \oplus v^{q_2}$.

Note that the number of states in $(\llbracket A \rrbracket \times \llbracket B \rrbracket)_{\cong}$ is finite, with an upper bound given by Proposition 2.18. Next we define transitions in $(\llbracket A \rrbracket \times \llbracket B \rrbracket)_{\cong}$:

Notation 2.26. Regions in $(\llbracket A \rrbracket \times \llbracket B \rrbracket)_{\cong}$ will be denoted X, X' . The equivalence class of a pair $(p, q) \in \llbracket A \rrbracket \times \llbracket B \rrbracket$ is denoted $[p, q]$.

Definition 2.27. For $X, X' \in (\llbracket A \rrbracket \times \llbracket B \rrbracket)_{\cong}$ we say that

- $X \xrightarrow{a}_{\ell} X'$ for $a \in \Sigma$ if for all $(p, q) \in X$ there exists $(p', q) \in X'$ such that $(p, q) \xrightarrow{a} (p', q)$ in $\llbracket A \rrbracket \times \llbracket B \rrbracket$,
- $X \xrightarrow{b}_r X'$ for $b \in \Sigma$ if for all $(p, q) \in X$ there exists $(p, q') \in X'$ such that $(p, q) \xrightarrow{b} (p, q')$ in $\llbracket A \rrbracket \times \llbracket B \rrbracket$, and
- $X \xrightarrow{\delta} X'$ if for all $(p, q) \in X$ there exists $d \in \mathbb{R}_{\geq 0}$ and $(p', q') \in X'$ such that $(p, q) \xrightarrow{d} (p', q')$.

Definition 2.28. A subset $\mathcal{B} \subseteq ([A] \times [B])_{\cong}$ is a *symbolic bisimulation* provided that for all $X \in \mathcal{B}$,

- whenever $X \xrightarrow{a}_{\ell} X'$ for some $X' \in ([A] \times [B])_{\cong}$, then $X' \xrightarrow{a}_r X''$ for some $X'' \in \mathcal{B}$,
- whenever $X \xrightarrow{a}_r X'$ for some $X' \in ([A] \times [B])_{\cong}$, then $X' \xrightarrow{a}_{\ell} X''$ for some $X'' \in \mathcal{B}$, and
- whenever $X \xrightarrow{\delta} X'$ for some $X' \in ([A] \times [B])_{\cong}$, then $X' \in \mathcal{B}$.

Note that it is decidable whether $([A] \times [B])_{\cong}$ admits a symbolic bisimulation. The following proposition finishes the proof of Theorem 2.23.

Proposition 2.29. *The quotient $([A] \times [B])_{\cong}$ admits a symbolic bisimulation if and only if $A \sim B$.*

Proof (cf. [43]): For a given symbolic bisimulation $\mathcal{B} \subseteq ([A] \times [B])_{\cong}$, the set $R_{\mathcal{B}} = \{(p, q) \mid [p, q] \in \mathcal{B}\} \subseteq [A] \times [B]$ is a timed bisimulation. For the other direction, one can construct a symbolic bisimulation from a timed bisimulation $R \subseteq [A] \times [B]$ by $\mathcal{B}_R = \{[p, q] \mid (p, q) \in R\}$. $\square$

2.5. Language inclusion and equivalence

Similarly to the untimed setting, there is also a notion of language inclusion and equivalence for timed automata. We need to introduce the notion of *timed trace* first. Note that we restrict to *finite* timed traces here; similar results are available for infinite traces in timed automata with Büchi or Muller acceptance conditions, see [9].

Definition 2.30. A *timed trace* over a finite set of actions Σ is a finite sequence $((t_1, a_1), (t_2, a_2), \dots, (t_k, a_k))$, where $a_i \in \Sigma$ and $t_i \in \mathbb{R}_{\geq 0}$ for $i = 1, \dots, k$, and $t_i < t_{i+1}$ for $i = 1, \dots, k-1$. The set of all timed traces over Σ is denoted $T\Sigma^*$.

In a pair (t_i, a_i) , the number t_i is called the *time stamp* of the action a_i , *i.e.* the time at which event a_i occurs.

Remark 2.31. Timed traces as defined above are also known as *strongly monotonic* timed traces, because of the assumption that no consecutive events occur at the same time. *Weakly* monotonic timed traces, *i.e.* with requirement $t_i \leq t_{i+1}$ instead of $t_i < t_{i+1}$, have also been considered, and there are some subtle differences between the two; see [85] for an important example.

Definition 2.32. A timed trace $((t_1, a_1), \dots, (t_k, a_k))$ is *accepted* by a timed automaton $A = (L, \ell_0, F, C, \Sigma, I, E)$ if there is an accepting run

$$\begin{aligned} (\ell_0, v_0) &\xrightarrow{t_1} (\ell_0, v_0 + t_1) \xrightarrow{a_1} (\ell_1, v_1) \xrightarrow{t_2 - t_1} \dots \\ &\dots \xrightarrow{a_{k-1}} (\ell_{k-1}, v_{k-1}) \xrightarrow{t_k - t_{k-1}} (\ell_{k-1}, v_{k-1} + t_k - t_{k-1}) \xrightarrow{a_k} (\ell_k, v_k) \end{aligned}$$

in A . The *timed language* of A is $L(A) = \{\tau \in T\Sigma^* \mid \tau \text{ accepted by } A\}$.

Figure 5.: Timed trace encoding an increment instruction I_{i+1} of a 2-counter machine.

It is clear that $L(A) = \emptyset$ if and only if none of the locations in F is reachable, hence Theorem 2.19 provides us with the decidability result in the following theorem. Undecidability of universality was established in [9]; we give an account of the proof below.

Theorem 2.33. *For a timed automaton $A = (L, \ell_0, F, C, \Sigma, I, E)$, deciding whether $L(A) = \emptyset$ is PSPACE-complete. It is undecidable whether $L(A) = T\Sigma^*$.*

Proof: We show that the universality problem for a timed automata is undecidable by reduction from the Σ_1^1 -hard problem of deciding whether a given 2-counter machine M has a recurring computation.

Let the timed language L_u be the set of timed traces encoding recurring computations of M . Observe that $L_u = \emptyset$ if and only if M does not have such a computation. We then construct a timed automaton A_u which accepts the complement of L_u , i.e. $L(A_u) = T\Sigma^* \setminus L_u$. Hence the language of A_u is universal if and only if M does not have a recurring computation.

Recall that a 2-counter, or Minsky, machine M is a finite sequence of labeled instructions $\{I_0, \dots, I_n\}$ and counters $\mathbf{x}_1$ and $\mathbf{x}_2$, with I_i for $0 \leq i \leq n-1$ on the form

$$I_i : \mathbf{x}_c := \mathbf{x}_c + 1; \text{ goto } I_j \quad \text{or} \quad I_i : \begin{cases} \text{if } \mathbf{x}_c = 0 \text{ then goto } I_j \\ \text{else } \mathbf{x}_c = \mathbf{x}_c - 1; \text{ goto } I_k \end{cases}$$

for $c \in 1, 2$, with a special $I_n : \text{Halt}$ instruction which stops the computation.

The language L_u is designed such that each I_i and the counters $\mathbf{x}_1$ and $\mathbf{x}_2$ are represented by actions in Σ . A correctly encoded computation is represented by a timed trace where “instruction actions” occur at discrete intervals, while the state (values of $\mathbf{x}_1$ and $\mathbf{x}_2$) is encoded by occurrences of “counter actions” in-between instruction actions (e.g. if $\mathbf{x}_i = 5$ after instruction I_j , then action x_i occurs 5 times within the succeeding interval of length 1).

When counters are incremented (or decremented), one more (or less) such action occurs through the next interval, and increments and decrements are always from the right. Additionally we require corresponding counter actions to occur exactly with a time difference of 1, such that if x_i occurs with time stamp a then also x_i occurs with time stamp $a + 1$, unless x_i is the rightmost x_i action and I_i at time stamp $\lfloor a \rfloor$ is a decrement of $\mathbf{x}_i$. Figure 5 shows an increment of $\mathbf{x}_1$ (from 4 to 5) using actions 1 and 2.

Figure 6.: Timed automaton which violates the encoding of the increment instruction.

We obtain A_u as a disjunction of timed automata $A^1, \dots, A^k$ where each A^i violates some property of a (correctly encoded) timed trace in L_u , either by accepting traces of incorrect format or inaccurate encodings of instructions.

Consider the instruction: $(p): x_1 := x_1 + 1 \text{ goto } (q)$, incrementing x_1 and jumping to q . A correct encoding would be similar to the one depicted in Figure 5 where all 1's and 2's are matched *one* time unit later, but with an additional 1 action occurring. In order to accept all traces except this encoding we must consider all possible violations, *i.e.*

- not incrementing the counter (no change),
- decrementing the counter,
- incrementing the counter more than once,
- jumping to the wrong instruction, or
- incrementing the wrong counter,

and construct a timed automaton having exactly such traces.

Figure 6 shows the timed automaton accepting traces in which instruction p yields no change of x_1 . $\square$

Turning our attention to timed trace inclusion and equivalence, we note the following.

Proposition 2.34. *Let A and B be timed automata. If A is timed simulated by B , then $L(A) \subseteq L(B)$. If A and B are timed bisimilar, then $L(A) = L(B)$.*

By a standard argument, Theorem 2.33 implies undecidability of timed trace inclusion and equivalence, a result first shown in [8].

Theorem 2.35. *Timed trace inclusion and equivalence are undecidable for timed automata.*

There is also a notion of *untimed* traces for timed automata.

Definition 2.36. The *untiming* of a set of timed traces $L \subseteq T\Sigma^*$ over a finite set of actions Σ is the set

$$UL = \{w = (a_1, \dots, a_k) \in \Sigma^* \mid \exists t_1, \dots, t_k \in \mathbb{R}_{\geq 0} : ((t_1, a_1), \dots, (t_k, a_k)) \in L\}.$$

Hence we have a notion of the set $UL(A)$ of *untimed language* of a timed automaton A . One can also define an untiming operation U for timed automata,

forgetting about the timing information of a timed automaton and thus converting it to a finite automaton; note however that $UL(A) \subsetneq L(UA)$ in general.

Lemma 2.37 ([9]). *For A a timed automaton, $UL(A) = L(\llbracket A \rrbracket_{\cong})$ provided that δ -transitions in $\llbracket A \rrbracket_{\cong}$ are taken as silent.*

As a corollary, sets of untimed traces accepted by timed automata are *regular*:

Theorem 2.38 ([9]). *For a timed automaton $A = (L, \ell_0, F, C, \Sigma, I, E)$, the set $UL(A) \subseteq \Sigma^*$ is regular. Accordingly, whether $UL(A) = \emptyset$ is decidable, and so is whether $UL(A) = \Sigma^*$. Also untimed trace inclusion and equivalence are decidable.*

2.6. Zones and difference-bound matrices

As shown in the above sections, regions provide a finite and elegant abstraction of the infinite state space of timed automata, enabling us to prove decidability of reachability, timed and untimed bisimilarity, untimed language equivalence and language emptiness.

Unfortunately, the number of states obtained from the region partitioning is extremely large. In particular, by Proposition 2.18 the number of regions is exponential in the number of clocks as well as in the maximal constants of the timed automaton. Efforts have been made in developing more efficient representations of the state space [23, 28, 74], using the notion of *zones* from Definition 2.5 on page 4 as a coarser and more compact representation of the state space.

An extended clock constraint over a finite set C may be represented using a directed weighted graph, where the nodes correspond to the elements of C together with an extra “zero” node x_0 , and an edge $x_i \xrightarrow{k} x_j$ corresponds to a constraint $x_i - x_j \leq k$ (if there is more than one upper bound on $x_i - x_j$, k is the minimum of all these constraints’ right-hand sides). The extra clock x_0 is fixed at value 0, so that a constraint $x_i \leq k$ can be represented as $x_i - x_0 \leq k$. Lower bounds on $x_i - x_j$ are represented as (possibly negative) upper bounds on $x_j - x_i$, and strict bounds $x_i - x_j < k$ are represented by adding a flag to the corresponding edge.

The weighted graph in turn may be represented by its adjacency matrix, which in this context is known as a *difference-bound matrix* or DBM. The above technique has been introduced in [50].

Example 2.3. Figure 7 gives an illustration of an extended clock constraint together with its representation as a difference-bound matrix. Note that the clock constraint contains superfluous information.

Zone-based reachability analysis of a timed automaton A uses symbolic states of the type (ℓ, Z) , where ℓ is a location of A and Z is a zone, instead of the region-based symbolic states of Proposition 2.17.

Definition 2.39. For a finite set C , $Z \subseteq \mathbb{R}_{\geq 0}^C$, and $r \subseteq C$, define

- the *delay* of Z by $Z^\uparrow = \{v + d \mid v \in Z, d \in \mathbb{R}_{\geq 0}\}$ and
- the *reset* of Z under r by $Z[r] = \{v[r] \mid v \in Z\}$.

Figure 7.: Graph representation of extended clock constraint.

Lemma 2.40 ([62, 95]). *If Z is a zone over C and $r \subseteq C$, then $Z^\uparrow$ and $Z[r]$ are also zones over C .*

Extended clock constraints representing $Z^\uparrow$ and $Z[r]$ may be computed efficiently (in time cubic in the number of clocks in C) by representing the zone Z in a canonical form obtained by computing the *shortest-path closure* of the directed graph representation of Z , see [72].

Example 2.3 (continued). Figure 8 shows two canonical representations of the difference-bound matrix for the zone Z of Figure 7. The left part illustrates the shortest-path closure of Z ; on the right is the *shortest-path reduction* [72] of Z , essentially obtained by removing redundant edges from the shortest-path closure. The latter is useful for checking zone inclusion, see below.

The *zone automaton* associated with a timed automaton is similar to the region automaton of Proposition 2.17, but uses zones for symbolic states instead of regions:

Definition 2.41. The *zone automaton* associated with a timed automaton $A = (L, \ell_0, F, C, \Sigma, I, E)$ is the transition system $\llbracket A \rrbracket_Z = (S, s_0, \Sigma \cup \{\delta\}, T)$ given as follows:

$$\begin{aligned} S &= \{(\ell, Z) \mid \ell \in L, Z \subseteq \mathbb{R}_{\geq 0}^C \text{ zone}\} & s_0 &= (\ell_0, \llbracket v_0 \rrbracket) \\ T &= \{(\ell, Z) \xrightarrow{\delta} (\ell, Z^\uparrow \wedge I(\ell))\} \\ &\cup \{(\ell, Z) \xrightarrow{a} (\ell', (Z \wedge \varphi)[r] \wedge I(\ell')) \mid \ell \xrightarrow{\varphi, a, r} \ell' \in E\} \end{aligned}$$

The analogue of Proposition 2.13 for zone automata is as follows:

Proposition 2.42 ([95]). *A state (ℓ, v) in a timed automaton $A = (L, \ell_0, F, C, \Sigma, I, E)$ is reachable if and only if there is a zone $Z \subseteq \mathbb{R}_{\geq 0}^C$ for which $v \in Z$ and such that (ℓ, Z) is reachable in $\llbracket A \rrbracket_Z$.*

The zone automaton associated with a given timed automaton is *infinite* and hence unsuitable for reachability analysis. Finiteness can be enforced by employ-

Figure 8.: Canonical representations. Left: shortest-path closure; right: shortest-path reduction.

ing *normalization*, using the fact that region equivalence $\cong$ has finitely many equivalence classes:

Definition 2.43. For a timed automaton A and a zone $Z \subseteq \mathbb{R}_{\geq 0}^C$, the *normalization* of Z is the set $\{v : C \rightarrow \mathbb{R}_{\geq 0} \mid \exists v' \in Z : v \cong v'\}$

The normalized zone automaton is defined in analogy to the zone automaton from above, and Proposition 2.42 also holds for the normalized zone automaton. Hence we can obtain a reachability algorithm by applying any search strategy (depth-first, breadth-first, or another) on the normalized zone automaton.

Remark 2.44. For timed automata on *extended* clock constraints, *i.e.* with diagonal constraints permitted, it can be shown [27,32] that normalization as defined above does *not* give rise to a sound and complete characterization of reachability. Instead, one can apply a refined normalization which depends on the difference constraints used in the timed automaton, see [27].

In addition to the efficient computation of symbolic successor states according to the $\rightsquigarrow$ relation, termination of reachability analysis requires that we can efficiently recognize whether the search algorithm has encountered a given symbolic state. Here it is crucial that there is an efficient way of deciding inclusion $Z_1 \subseteq Z_2$ between zones. Both the shortest-path-closure canonical form as well as the more space-economical shortest-path-reduced canonical form [72], *cf.* Example 2.3, allow for efficient inclusion checking.

In analogy to difference-bound matrices and overcoming some of their problems, the data structure called *clock difference diagram* has been proposed [74]. However, the design of efficient algorithms for delay and reset operations over that data structure is a challenging open problem; generally, the design of efficient data structures for computations with (unions of) zones is a field of active research, see [3,12,84] for some examples.

Figure 9.: A weighted timed automaton with two clocks.

3. Weighted timed automata

The notion of *weighted* — or *priced* — timed automata was introduced independently, at the very same conference, by Behrmann *et.al.* [21] and Alur *et.al.* [10]. In these models both edges and locations can be decorated with weights, or prices, giving the cost of taking an action transition or the cost per time unit of delaying in a given location. The total cost of a trace is then simply the accumulated (or total) weight of its discrete and delay transitions.

As a first result, the above two papers independently, and with quite different methods, showed that the problem of cost-optimal reachability is computable for weighted timed automata with non-negative weights. Later, optimal reachability for timed automata with several weight functions was considered in [77] as well as optimal infinite runs in [34, 53].

Definition 3.1. A *weighted timed automaton* is a tuple $A = (L, \ell_0, F, C, \Sigma, I, E, R, P)$, where $(L, \ell_0, F, C, \Sigma, I, E)$ is a timed automaton, $R : L \rightarrow \mathbb{Z}$ a location weight-rate mapping, and $P : E \rightarrow \mathbb{Z}$ an edge weight mapping.

The *semantics* of A is the weighted transition system $\llbracket A \rrbracket = (S, s_0, \Sigma \cup \mathbb{R}_{\geq 0}, T, w)$, where $(S, s_0, \Sigma \cup \mathbb{R}_{\geq 0}, T)$ is the semantics of the underlying timed automaton $(L, \ell_0, F, C, \Sigma, I, E)$, and the transition weights $w : T \rightarrow \mathbb{R}$ are given as follows:

$$\begin{aligned} w((\ell, v) \xrightarrow{d} (\ell, v + d)) &= d R(\ell) \\ w((\ell, v) \xrightarrow{a} (\ell', v')) &= P(\ell \xrightarrow{\varphi, a, r} \ell') \quad \text{with } v \models \varphi, v' = v[r] \end{aligned}$$

We shall denote weighted edges and transitions by symbols $\xrightarrow[e]{e}$ to illustrate an edge or a transition labeled e with weight w .

3.1. Optimal reachability

The objective of optimal reachability analysis is to find runs to a final location with the lowest *total weight* as defined below.

Example 3.1. Figure 9 shows a simple weighted timed automaton with final location ℓ_3 . Below we give a few examples of accepting runs, where we identify valuations $v : \{x, y\} \rightarrow \mathbb{R}_{\geq 0}$ with their values $(v(x), v(y))$. The total weights of

the runs given here are 17 and 11; actually the second run is *optimal* in the sense of Problem 3.1 below:

$$\begin{aligned}
(\ell_1, 0, 0) &\xrightarrow[12]{3} (\ell_1, 3, 3) \xrightarrow[1]{a} (\ell_2, 0, 3) \xrightarrow[4]{c} (\ell_3, 0, 3) \\
(\ell_1, 0, 0) &\xrightarrow[1]{a} (\ell_2, 0, 0) \xrightarrow[6]{3} (\ell_2, 3, 3) \xrightarrow[0]{b} (\ell_2, 0, 3) \xrightarrow[4]{c} (\ell_3, 0, 3)
\end{aligned}$$

Definition 3.2. The *total weight* of a finite run $\rho = s_0 \xrightarrow[w_1]{} s_1 \xrightarrow[w_2]{} \cdots \xrightarrow[w_k]{} s_k$ in a weighted transition system is $w(\rho) = \sum_{i=1}^k w_i$.

We are now in a position to state the problem with which we are concerned here: We want to find accepting runs with minimum total weight in a weighted timed automaton A . However due to the possible use of strict clock constraints on edges and in locations of A , the minimum total weight might not be realizable, *i.e.* there might be no run which achieves it. For this reason, one also needs to consider (infinite) *sets* of runs and the infimum of their members' total weights:

Problem 3.1 (Optimal reachability). Given a weighted timed automaton A , compute $W = \inf \{w(\rho) \mid \rho \text{ accepting run in } A\}$ and a set P of accepting runs for which $\inf_{\rho \in P} w(\rho) = W$.

The key ingredient in the proof of the following theorem is the introduction of *weighted regions* in [21]. A weighted region is a region as of Definition 2.16 enriched with an affine cost function describing in a finite manner the cost of reaching any point within it. This notion allows one to define the weighted region automaton associated with a weighted timed automaton, and one can then show that optimal reachability can be computed in the weighted region automaton. PSPACE-hardness in the below theorem follows from PSPACE-hardness of reachability for timed automata.

Theorem 3.3 ([21]). *The optimal reachability problem for weighted timed automata with non-negative weights is PSPACE-complete.*

Similar to the notion of regions for timed automata, the number of weighted regions is exponential in the number of clocks as well as in the maximal constants of the timed automaton. Hence a notion of *weighted zone* — a zone extended with an affine cost function — was introduced [71] together with an efficient, symbolic A^* -algorithm for searching for cost-optimal tracing using branch-and-bound techniques. In particular, efficient means of generalizing the notion of symbolic successor to incorporate the affine cost functions were given.

During the symbolic exploration, several small linear-programming problems in terms of determining the minimal value of the cost function over the given zone have to be dealt with. Given that the constraints of these problems are simple difference constraints, it turns out that substantial gain in performance may be achieved by solving the dual problem of minimum-cost flow [88]. The newly emerged branch UPPAAL-CORA provides an efficient tool for cost-optimal reachability analysis, applying the above data structures and algorithms and allowing the user to guide and heuristically prune the search.

Figure 10.: A doubly weighted timed automaton with two clocks.

3.2. Multi-weighted timed automata

The below formalism of doubly weighted timed automata is a generalization of weighted timed automata useful for modeling systems with several different resources.

Definition 3.4. A *doubly weighted timed automaton* is a tuple

$$A = (L, \ell_0, F, C, \Sigma, I, E, R, P)$$

where $(L, \ell_0, F, C, \Sigma, I, E)$ is a timed automaton, $R : L \rightarrow \mathbb{Z}^2$ a location weight-rate mapping, and $P : E \rightarrow \mathbb{Z}^2$ an edge weight mapping.

The semantics of a doubly weighted timed automaton is a doubly weighted transition system defined similarly to Definition 3.1, and the total weight of finite runs is defined accordingly as a pair; we shall refer to the total weights as w_1 and w_2 respectively. These definitions have natural generalizations to *multi-weighted* timed automata with more than two weight coordinates.

The objective of conditional reachability analysis is to find runs to a final location with the lowest total weight in the first weight coordinate while satisfying a constraint on the other weight coordinate.

Example 3.2. Figure 10 depicts a simple doubly weighted timed automaton with final location ℓ_3 . Under the constraint $w_2 \leq 3$, the optimal run of the automaton can be seen to be

$$(\ell_1, 0, 0) \xrightarrow[\left(\frac{1}{3}, \frac{4}{3}\right)]{\frac{1}{3}} (\ell_1, 1/3, 1/3) \xrightarrow{a} (\ell_2, 1/3, 0) \xrightarrow[\left(\frac{10}{3}, \frac{5}{3}\right)]{\frac{5}{3}} (\ell_2, 2, 5/3) \xrightarrow{b} (\ell_3, 2, 0)$$

with total weight $(\frac{11}{3}, 3)$.

The precise formulation of the conditional optimal reachability problem is as follows, where we again need to refer to (possibly infinite) sets of runs:

Problem 3.2 (Conditional optimal reachability). Given a doubly weighted timed automaton A and $M \in \mathbb{Z}$, compute $W = \inf \{w_1(\rho) \mid \rho \text{ accepting run in } A, w_2(\rho) \leq M\}$ and a set P of accepting runs such that $w_2(\rho) \leq M$ for all $\rho \in P$ and $\inf_{\rho \in P} w(\rho) = W$.

Theorem 3.5 ([76, 77]). *The conditional optimal reachability problem is computable for doubly weighted timed automata with non-negative weights and without weights on edges.*

Figure 11.: A weighted timed automaton modelling a simple production system.

The proof of the above theorem rests on a direct generalization of weighted to *doubly-weighted* zones. An extension can be found in [77], where it is shown that also the *Pareto frontier*, *i.e.* the set of cost vectors which cannot be improved in any cost variable, can be computed.

3.3. Optimal infinite runs

In this section we shall be concerned with computing optimal *infinite* runs in (doubly) weighted timed automata. We shall treat both the *limit ratio* viewpoint discussed in [34] and the *discounting* approach of [53, 54].

Example 3.3. Figure 11 shows a simple production system modelled as a weighted timed automaton. The system has three modes of production, High, Medium, and Low. The weights model the *cost* of production, so that the High production mode has a low cost, which is preferable to the high cost of the Low production mode. After operating in a High or Medium production mode for three time units, production automatically degrades (action d) to a lower mode. When in Medium or Low production mode, the system can be attended to (action a), which advances it to a higher mode.

The objective of *optimal-ratio analysis* is to find an infinite run in a doubly weighted timed automaton which minimizes the *ratio* between the two total weights. This will be formalized below.

Definition 3.6. The *total ratio* of a finite run $\rho = s_0 \xrightarrow[z_1]{w_1} s_1 \xrightarrow[z_2]{w_2} \dots \xrightarrow[z_k]{w_k} s_k$ in a doubly weighted transition system is

$$\Gamma(\rho) = \frac{\sum_{i=1}^k w_i}{\sum_{i=1}^k z_i}.$$

The total ratio of an infinite run $\rho = s_0 \xrightarrow[z_1]{w_1} s_1 \xrightarrow[z_2]{w_2} \dots$ is

$$\Gamma(\rho) = \liminf_{k \rightarrow \infty} \Gamma(s_0 \rightarrow \dots \rightarrow s_k).$$

A special case of optimal-ratio analysis is given by weight-per-time models, where the interest is in minimizing total weight per accumulated time. The example provided in this section is a case of this. In the setting of optimal-ratio analysis, these can be modelled as doubly weighted timed automata with $R_2(\ell) = 1$ and $P_2(e) = 0$ for all locations ℓ and edges e .

Example 3.3 (continued). In the timed automaton of Figure 11, the following cyclic behaviour provides an infinite run ρ :

$$\begin{aligned} (H, 0, 0) \xrightarrow{3} (H, 3, 3) \xrightarrow{d} (M, 0, 3) \xrightarrow{3} (M, 3, 6) \xrightarrow{d} (L, 3, 6) \xrightarrow{1} \\ (L, 4, 7) \xrightarrow{a} (M, 0, 0) \xrightarrow{3} (M, 3, 3) \xrightarrow{a} (H, 0, 0) \rightarrow \dots \end{aligned}$$

Taking the weight-per-time viewpoint, the total ratio of ρ is $\Gamma(\rho) = 4.8$.

Problem 3.3 (Minimum infinite ratio). Given a doubly weighted timed automaton A , compute $W = \inf \{ \Gamma(\rho) \mid \rho \text{ infinite run in } A \}$ and a set P of infinite runs for which $\inf_{\rho \in P} \Gamma(\rho) = W$.

The main tool in the proof of the following theorem is the introduction of the *corner-point abstraction* of a timed automaton in [34]. This is a finite refinement of the region automaton of Definition 2.16 in which one also keeps track of the corner points of regions. One can then show that any infinite run with minimum ratio must pass through corner points of regions, hence these can be found in the corner-point abstraction by an algorithm first proposed in [68].

The technical condition in the theorem that the second weight coordinate be *strongly diverging* means that any infinite run ρ in the closure of the timed automaton in question satisfies $w_2(\rho) = \infty$, see [34] for details.

Theorem 3.7 ([34]). *The minimum infinite ratio problem is computable for doubly weighted timed automata with non-negative and strongly diverging second weight coordinate.*

For *discount-optimal analysis*, the objective is to find an infinite run in a weighted timed automaton which minimizes the *discounted total weight* as defined below. The point of discounting is that the weight of actions is discounted with time, so that the impact of an event decreases, the further in the future it takes place.

In the definition below, ε is the empty run, and $(\ell, v) \rightarrow \rho$ denotes the concatenation of the transition $(\ell, v) \rightarrow$ with the run ρ .

Definition 3.8. The *discounted total weight* of finite runs in a weighted timed automaton under discounting factor $\lambda \in [0, 1[$ is given inductively as follows:

$$\begin{aligned} w_\lambda(\varepsilon) &= 0 \\ w_\lambda((\ell, v) \xrightarrow{P} \rho) &= P + w_\lambda(\rho) \\ w_\lambda((\ell, v) \xrightarrow{d} \rho) &= R(\ell) \int_0^d \lambda^\tau d\tau + \lambda^d w_\lambda(\rho) \end{aligned}$$

The discounted total weight of an infinite run $\rho = (\ell_0, v_0) \xrightarrow{d_1} (\ell_0, v_0 + d_1) \xrightarrow{P_1} (\ell_1, v_1) \rightarrow \dots$ is

$$w_\lambda(\rho) = \lim_{k \rightarrow \infty} w_\lambda((\ell_0, v_0) \rightarrow \cdots \xrightarrow{P_k} (\ell_k, v_k))$$

provided that the limit exists.

Example 3.3 (continued). The discounted total weight of the infinite run ρ in the timed automaton of Figure 11 satisfies the following equality, where $I_t = \int_0^t \lambda^\tau d\tau = -\frac{1}{\ln \lambda}(1 - \lambda^t)$:

$$w_\lambda(\rho) = 2I_3 + \lambda^3(5I_3 + \lambda^3(9I_1 + \lambda(1 + 5I_3 + \lambda^3(2 + w_\lambda(\rho)))))$$

With a discounting factor of $\lambda = .9$ for example, the discounted total weight of ρ would hence be $w_\lambda(\rho) \approx 40.5$.

Problem 3.4 (Minimum discounted weight). Given a weighted timed automaton A and $\lambda \in [0, 1[$, compute $W = \inf \{w_\lambda(\rho) \mid \rho \text{ infinite run in } A\}$ and a set P of infinite runs for which $\inf_{\rho \in P} w_\lambda(\rho) = W$.

The proof of the following theorem rests again on the corner-point abstraction, and on a result in [11]. The technical condition that the timed automaton be time-divergent is analogous to the condition on the second weight coordinate in Theorem 3.7.

Theorem 3.9 ([53]). *The minimum discounted weight problem is computable for time-divergent weighted timed automata with non-negative weights and rational λ .*

4. Timed games

Recently, substantial effort has been made towards the synthesis of winning strategies for timed games with respect to *safety* and *reachability control objectives*. From known region-based decidability results, efficient on-the-fly algorithms have been developed [41, 92] and implemented in the newest branch UPPAAL-TIGA.

For timed games, as for untimed ones, transitions are either controllable or uncontrollable (*i.e.* under the control of an environment), and the problem is to synthesize a strategy for *when* to take *which* (enabled) controllable transitions in order that a given objective is guaranteed regardless of the behaviour of the environment.

Definition 4.1. A *timed game* is a tuple $(L, \ell_0, F, C, \Sigma_c, \Sigma_u, I, E)$ with $\Sigma_c \cap \Sigma_u = \emptyset$ and for which the tuple $(L, \ell_0, F, C, \Sigma = \Sigma_c \cup \Sigma_u, I, E)$ is a timed automaton.

Edges with actions in Σ_c are said to be *controllable*, those with actions in Σ_u are *uncontrollable*.

Example 4.1. Figure 12 provides a simple example of a timed game. Here, $\Sigma_c = \{c_1, c_2, c_4\}$ and $\Sigma_u = \{u_1, u_2, u_3\}$, and the controllable edges are drawn with solid lines, the uncontrollable ones with dashed lines.

Figure 12.: A timed game with one clock. Controllable edges (with actions from Σ_c) are solid, uncontrollable edges (with actions from Σ_u) are dashed.

We need the notion of *strategy*; essentially, a strategy provides instructions for which controllable edge to take, or whether to wait, in a given state:

Definition 4.2. A *strategy* for a timed game $A = (L, \ell_0, F, C, \Sigma_c, \Sigma_u, I, E)$ is a mapping σ from finite runs of A to $\Sigma_c \cup \{\delta\}$, where $\delta \notin \Sigma$, such that for any run $\rho = (\ell_0, v_0) \rightarrow \dots \rightarrow (\ell_k, v_k)$,

- if $\sigma(\rho) = \delta$, then $(\ell, v) \xrightarrow{d} (\ell, v + d)$ in $\llbracket A \rrbracket$ for some $d > 0$, and
- if $\sigma(\rho) = a$, then $(\ell, v) \xrightarrow{a} (\ell', v')$ in $\llbracket A \rrbracket$.

A strategy σ is said to be *memoryless* if $\sigma(\rho)$ only depends on the last state of ρ , i.e. if $\rho_1 = (\ell_0, v_0) \xrightarrow{d_1} (\ell_0, v_0 + d_1) \rightarrow \dots \rightarrow (\ell_k, v_k)$, $\rho_2 = (\ell_0, v_0) \xrightarrow{d'_1} (\ell_0, v_0 + d'_1) \rightarrow \dots \rightarrow (\ell_k, v_k)$ imply $\sigma(\rho_1) = \sigma(\rho_2)$.

An *outcome* of a strategy is any run which adheres to its instructions in the obvious manner:

Definition 4.3. A run $(\ell_0, v_0) \xrightarrow{d_1} (\ell_0, v_0 + d_1) \rightarrow \dots \rightarrow (\ell_k, v_k)$ in a timed game $A = (L, \ell_0, F, C, \Sigma_c, \Sigma_u, I, E)$ is said to be an *outcome* of a strategy σ provided that

- for all $(\ell_i, v_i) \xrightarrow{d} (\ell_i, v_i + d)$ and for all $d' < d$, we have $\sigma((\ell_0, v_0) \rightarrow \dots \rightarrow (\ell_i, v_i + d')) = \delta$, and
- for all $(\ell_i, v_i + d) \xrightarrow{a} (\ell_{i+1}, v_{i+1})$ for which $a \in \Sigma_c$, we have $\sigma((\ell_0, v_0) \rightarrow \dots \rightarrow (\ell_i, v_i)) = a$.

An outcome is said to be *maximal* if $\ell_k \in F$, or if $(\ell_k, v_k) \xrightarrow{a} (\ell_{k+1}, v_{k+1})$ implies $a \in \Sigma_u$.

Hence an outcome is maximal if it stops in a final state, or if no controllable actions are available at its end. An underlying assumption is that uncontrollable actions cannot be forced, hence a maximal outcome which does not end in a final state may “get stuck” in a non-final state. The aim of reachability games is to find strategies all of whose maximal outcomes end in a final state; the aim of safety

games is to find strategies all of whose (not necessarily maximal) outcomes avoid final states:

Definition 4.4. A strategy is said to be *winning for the reachability game* if any of its maximal outcomes is an accepting run. It is said to be *winning for the safety game* if none of its outcomes are accepting.

Example 4.1 (continued). The following memoryless strategy is winning for the reachability game on the timed game from Figure 12:

$$\begin{aligned}\sigma(\ell_1, v) &= \begin{cases} \delta & \text{if } v(x) \neq 1 \\ c_1 & \text{if } v(x) = 1 \end{cases} & \sigma(\ell_2, v) &= \begin{cases} \delta & \text{if } v(x) < 2 \\ c_2 & \text{if } v(x) \geq 2 \end{cases} \\ \sigma(\ell_3, v) &= \begin{cases} \delta & \text{if } v(x) < 1 \\ c_3 & \text{if } v(x) \geq 1 \end{cases} & \sigma(\ell_4, v) &= \begin{cases} \delta & \text{if } v(x) \neq 1 \\ c_4 & \text{if } v(x) = 1 \end{cases}\end{aligned}$$

Problem 4.1 (Reachability and safety games). Given a timed game A , does there exist a winning strategy for the reachability game on A ? Does there exist a winning strategy for the safety game on A ?

An important ingredient in the proof of the following theorem is the fact that for reachability as well as safety games, it is sufficient to consider *memoryless* strategies. This is not the case for other, more subtle, control objectives (e.g. counting properties modulo some N) as well as for the synthesis of winning strategies under *partial observability*.

Theorem 4.5 ([13, 82]). *The reachability and safety games are decidable for timed games.*

In [42] the on-the-fly algorithm applied in UPPAAL-TIGA has been extended to timed games under partial observability.

The field of timed games is a very active research area. Research has been conducted towards the synthesis of *optimal* winning strategies for reachability games on *weighted timed games*. In [6, 35] computability of optimal strategies is shown under a certain condition of *strong cost non-zenoness*, requiring that the total weight diverges with a given minimum rate per time. Later undecidability results [33, 38] show that for weighted timed games with three or more clocks this condition (or a similar one) is necessary. Lately [36] proves that optimal reachability strategies are computable for one-clock weighted timed games, though there is an unsettled (large) gap between the known lower bound complexity P and an upper bound of $3EXPTIME$.

We conclude this section by reestablishing the connection between the notion of games and bisimulation [90] in the presence of time:

Proposition 4.6. *Timed bisimilarity polynomial-time reduces to timed safety games.*

Observe that this provides an alternative proof of the decidability of timed bisimilarity in Theorem 2.23 on page 10.

Figure 13.: A timed game constructed for bisimilarity checking of two simple timed automata

Proof: Given timed automata $A_i = (L_i, \ell_0^i, F, C_i, \Sigma, I_i, E_i)$, for $i \in \{1, 2\}$, with $C_1 \cap C_2 = \emptyset$, we consider the timed game with locations $L = \{\perp\} \cup (L_1 \times L_2) \cup (L_1 \times L_2 \times \Sigma \times \{1, 2\})$, where $F = \{\perp\}$ is a designated final location. We set $C = C_1 \cup C_2 \cup \{z\}$, where $z \notin C_1 \cup C_2$ is a fresh clock, $\Sigma_c = \Sigma \cup \{\perp\}$ and $\Sigma_u = \{a' \mid a \in \Sigma\}$, and E is defined by

$$\begin{aligned}
(p, q) &\xrightarrow{\varphi, a', \tilde{r}} (p', q, a)_1 \in E \quad \text{if } p \xrightarrow{\varphi, a, r} p' \in E_1, \\
(p, q) &\xrightarrow{\varphi, a', \tilde{r}} (p, q', a)_2 \in E \quad \text{if } q \xrightarrow{\varphi, a, r} q' \in E_2, \\
(p', q, a)_1 &\xrightarrow{\tilde{\varphi}, a, r} (p', q') \in E \quad \text{if } q \xrightarrow{\varphi, a, r} q' \in E_2, \\
(p, q', a)_2 &\xrightarrow{\tilde{\varphi}, a, r} (p', q') \in E \quad \text{if } p \xrightarrow{\varphi, a, r} p' \in E_1, \text{ and} \\
(p, q, a)_i &\xrightarrow{\Phi \wedge z=0, \perp, \emptyset} \perp \quad \text{for all } i \in \{1, 2\}.
\end{aligned}$$

Here we denote $\tilde{r} = r \cup \{z\}$ and $\tilde{\varphi} = \varphi \wedge z = 0$, and $\Phi = \bigvee_j \neg \varphi_j$ when $q \xrightarrow{\varphi_j, a, r} q'$ and $i = 1$ and symmetrically for $i = 2$. Location invariants are defined by $I(p, q) = I_1(p) \wedge I_2(q)$ for $(p, q) \in L_1 \times L_2$ and $I(p, q, a)_i = (z = 0)$ for all $(p, q, a)_i \in L_1 \times L_2 \times \Sigma \times \{1, 2\}$. See Figure 13 for a simple example of this construction.

It remains to be seen that A_1 and A_2 are timed bisimilar if and only if a strategy σ exists for which any outcome $\rho = (\ell_0, v_0) \xrightarrow{d_1} (\ell_0, v_0 + d_1) \rightarrow \dots \rightarrow (\ell_k, v_k)$ satisfies $\ell_k \neq \perp$ (i.e. it avoids the final location $\perp$).

Assume A_1 and A_2 are timed bisimilar, then we can prove something stronger than the above, namely that *any* strategy will avoid $\perp$. Indeed, if $(\ell_0, v_0) \xrightarrow{d_1}$

$(\ell_0, v_0 + d_1) \rightarrow \dots \rightarrow (\ell_k, v_k)$ is a run in the timed game, and the transition $(\ell_{k-2}, v_{k-2}) \xrightarrow{a} (\ell_{k-1}, v_{k-1})$ exists due to the first component (p_{k-2}, v_{k-2}) of (ℓ_{k-2}, v_{k-2}) , then the corresponding $\xrightarrow{a}$ transition in $\llbracket A_1 \rrbracket$ has a matching $\xrightarrow{a}$ transition in $\llbracket A_2 \rrbracket$. Hence the state (ℓ_{k-1}, v_{k-1}) has an enabled a -labeled edge, which by definition of Φ implies that the edge to $\perp$ is disabled, thus $\ell_k \neq \perp$. A symmetric argument applies in the other case.

Now assume σ is a strategy which ensures avoidance of $\perp$, then we shall show that any $(\ell_j, v) = ((p_j, q_j), v)$ for $j \geq 0$, (*i.e.* of type $S_1 \times S_2$) occurring in an outcome of σ satisfies $(p_j, v) \sim (q_j, v)$. Assume to the contrary that $(p_j, v) \xrightarrow{a} (p'_j, v') \in \llbracket A_1 \rrbracket$ and $(q_j, v) \not\xrightarrow{a} (q'_j, v'') \in \llbracket A_2 \rrbracket$ for some $a \in \Sigma$, then we may extend the run to (ℓ_j, v) by $(\ell_j, v) \xrightarrow{a} (\ell_{j+1}, v') \xrightarrow{\perp} \perp$, moreover $\sigma((\ell_0, v_0) \rightarrow \dots \rightarrow (\ell_{j+1}, v')) = \perp$ is the only choice for σ as by construction neither δ (due to the invariant $z = 0$) nor any $b \neq a$ is available. $\square$

5. Statistical Model Checking for Networks of Price Timed Automata

A weak point of model checking is undoubtedly the state-space explosion, *i.e.* the exponential growth in the analysis effort measured in the number of model-components. Another limitation of real-time model checking is that it merely provides – admittedly most important – hard quantitative guarantees, *e.g.* the worst case response time of a recurrent task under a certain scheduling principle, the worst case execution time of a piece of code running on a particular execution platform, or the worst case time before consensus is reached by a real-time network protocol. In addition to these hard guarantees, it would be desirable in several situations to obtain refined performance information concerning likely or expected behaviors in terms of timing and resource consumption. In particular, this would allow to distinguish and select between systems that perform identically from a worst-case perspective.

In a series of recent works [48], we proposed a stochastic semantics for Priced Timed Automata (PTA), whose clocks can evolve with different rates, while² being used with no restrictions in guards and invariants. Networks of PTAs (NPTA) are created by composing PTAs via input and output actions. More precisely, we define a natural stochastic semantics for networks of NPTAs based on races between components being composed. We shall observe that such race can generate arbitrarily complex stochastic behaviors from simple assumptions on individual components. We shall see that our semantics cannot be emulated by applying the existing stochastic semantic of [14, 30] to the product of components. Other related work includes the very rich framework of stochastic timed systems of MoDeST [31]. Here, however, general hybrid variables are not considered and parallel composition does not yield fully stochastic models. For the notion of probabilistic hybrid systems considered in [91] the choice of time is resolved non-deterministically rather than stochastically as in our case. Moreover, based on

²in contrast to the usual restriction of priced timed automata [10]

Figure 14.: An NPTA, $(A|B|T)$.

the stochastic semantics, we are able to express refined performance properties, e.g. in terms of probabilistic guarantees of time- and cost-bounded properties³.

To allow for the efficient analysis of probabilistic performance properties we propose to work with Statistical Model Checking (SMC) [89, 96], an approach that has been proposed as an alternative to avoid an exhaustive exploration of the state-space of the model. The core idea of SMC is to monitor some simulations of the system, and then use results from the statistic area (including sequential hypothesis testing or Monte Carlo simulation) in order to decide whether the system satisfies the property with some degree of confidence.

In this section, we first give insights on the model and on the stochastic semantic, then on the use of statistical model checking. Finally, we conclude with a brief discussion and some applications;

5.1. Networks of Stochastic Automata

Networks of Price Timed Automata We consider the analysis of Priced Timed Automata (PTAs) that are timed automata whose clocks can evolve with different rates in different locations. In fact, the expressive power (up to timed bisimilarity) of NPTA equals that of general linear hybrid automata (LHA) [4], rendering most problems – including that of reachability – undecidable. We also assume PTAs are input-enabled, deterministic (with a probability measure defined on the sets of successors), and non-zeno. PTAs communicate via broadcast channels and shared variables to generate Networks of Price Timed Automata (NPTA).

Fig. 14 provides an NPTA with three components A , B , and T as specified using the UPPAAL GUI. One can easily see that the composite system $(A|B|T)$ has the transition sequence:

$$\begin{aligned} ((A_0, B_0, T_0), [x = 0, y = 0, C = 0]) &\xrightarrow{1} \xrightarrow{a!} \\ ((A_1, B_0, T_1), [x = 1, y = 1, C = 4]) &\xrightarrow{1} \xrightarrow{b!} \\ ((A_1, B_1, T_2), [x = 2, y = 2, C = 6]), \end{aligned}$$

demonstrating that the final location T_3 of T is reachable. In fact, location T_3 is reachable within cost 0 to 6 and within total time 0 and 2 in $(A|B|T)$ depending on when (and in which order) A and B choose to perform the output actions $a!$ and $b!$. Assuming that the choice of these time-delays is governed by probability

³Clocks with different rates can be used to model costs.

Figure 15.: Cumulative probabilities for **time** and **Cost**-bounded reachability of T_3 .

distributions, a measure on sets of runs of NPTAs is induced, according to which quantitative properties such as “the probability of T_3 being reached within a total cost-bound of 4.3” become well-defined.

Probabilistic Semantics of NPTA Components In our early works [48], we provide a natural stochastic semantics, where PTA components associate probability distributions to both the time-delays spent in a given state as well as to the transition between states. In UPPAAL-SMC uniform distributions are applied for bounded delays and exponential distributions for the case where a component can remain indefinitely in a state. In a network of PTAs the components repeatedly race against each other, i.e. they independently and stochastically decide on their own how much to delay before outputting, with the “winner” being the component that chooses the minimum delay. For instance, in the NPTA of Fig. 14, A wins the initial race over B with probability 0.75.

In contrast to the probabilistic semantics of timed automata in [14, 30] our semantics deals with networks and thus with races between components. Let $\mathcal{A}^j = (L^j, X^j, \Sigma, E^j, R^j, I^j)$ ($j = 1 \dots n$) be a collection of composable NPTAs. Under the assumption of input-enabledness, disjointness of clock sets and output actions, states of the the composite NPTA $\mathcal{A} = (\mathcal{A}_1 | \dots | \mathcal{A}_n)$ may be seen as tuples $\mathbf{s} = (s_1, \dots, s_n)$ where s_j is a state of $\mathcal{A}^j$, i.e. of the form (ℓ, ν) where $\ell \in L^j$ and $\nu \in \mathbb{R}_{\geq 0}^{X^j}$. Our probabilistic semantics is based on the principle of independency between components. Repeatedly each component decides on its own – based on a given delay density function and output probability function – how much to delay before outputting and what output to broadcast at that moment. Obviously, in such a race between components the outcome will be determined by the component that has chosen to output after the minimum delay: the output is broadcast and all other components may consequently change state.

Let us first consider a component $\mathcal{A}^j$ and let S^j denote the corresponding set of states. For each state $s = (\ell, \nu)$ of $\mathcal{A}^j$ we shall provide probability distributions for both delays and outputs. In this presentation, we restrict to uniform and universal distributions, but arbitrary distributions can be considered.

The *delay density function* μ_s over delays in $\mathbb{R}_{\geq 0}$ will be either a uniform or an exponential distribution depending on the invariant of ℓ . Denote by E_ℓ the disjunction of guards g such that $(\ell, g, o, -, -) \in E^j$ for some output o . Denote by $d(\ell, \nu)$ the infimum delay before enabling an output, i.e. $d(\ell, \nu) = \inf\{d \in \mathbb{R}_{\geq 0} : \nu + R^j \cdot d \models E_\ell\}$, and denote by $D(\ell, \nu)$ the supremum delay,

i.e. $D(\ell, \nu) = \sup\{d \in \mathbb{R}_{\geq 0} : \nu + R^j \cdot d \models I^j(\ell)\}$. If $D(\ell, \nu) < \infty$ then the delay density function μ_s is a uniform distribution on $[d(\ell, \nu), D(\ell, \nu)]$. Otherwise – that is $I^j(\ell)$ does not put an upper bound on the possible delays out of s – the delay density function μ_s is an exponential distribution with a rate $P(\ell)$, where $P : L^j \rightarrow \mathbb{R}_{\geq 0}$ is an *additional* distribution rate component added to the NPTA $\mathcal{A}^j$. For every state $s = (\ell, \nu)$, the *output probability function* γ_s over Σ_o^j is the uniform distribution over the set $\{o : (\ell, g, o, -, -) \in E^j \wedge \nu \models g\}$ whenever this set is non-empty⁴. We denote by s^o the state after the output of o . Similarly, for every state s and any input action ι , we denote by s^ι the state after having received the input ι .

Probabilistic Semantics of Networks of NPTA We shall now see that while the stochastic semantics of each PTA is rather simple (but quite realistic), arbitrarily complex stochastic behavior can be obtained by their composition.

Reconsider the closed network $\mathcal{A} = (\mathcal{A}_1 \mid \dots \mid \mathcal{A}_n)$ with a state space $St = St_1 \times \dots \times St_n$. For $\mathbf{s} = (s_1, \dots, s_n) \in St$ and $a_1 a_2 \dots a_k \in \Sigma^*$ we denote by $\pi(\mathbf{s}, a_1 a_2 \dots a_k)$ the set of all maximal runs from $\mathbf{s}$ with a prefix $t_1 a_1 t_2 a_2 \dots t_k a_k$ for some $t_1, \dots, t_n \in \mathbb{R}_{\geq 0}$, that is runs where the i 'th action a_i has been outputted by the component $A_{c(a_i)}$. We now inductively define the following measure for such sets of runs:

$$P_A(\pi(\mathbf{s}, a_1 \dots a_n)) = \int_{t \geq 0} \mu_{s_c}(t) \cdot \left(\prod_{j \neq c} \int_{\tau > t} \mu_{s_j}(\tau) d\tau \right) \cdot \gamma_{s_c}^t(a_1) \cdot P_A(\pi(\mathbf{s}^t)^{a_1, a_2 \dots a_n}) dt$$

where $c = c(a_1)$, and as base case we take $P_A(\pi(\mathbf{s}), \varepsilon) = 1$.

This definition requires a few words of explanation: at the outermost level we integrate over all possible initial delays t . For a given delay t , the outputting component $c = c(a_1)$ will choose to make the broadcast at time t with the stated density. Independently, the other components will choose to a delay amount, which – in order for c to be the winner – must be larger than t ; hence the product of the probabilities that they each make such a choice. Having decided for making the broadcast at time t , the probability of actually outputting a_1 is included. Finally, in the global state resulting from all components having delayed t time-units and changed state according to the broadcasted action a_1 the probability of runs according to the remaining actions $a_2 \dots a_n$ is taken into account.

The Hammer Game To illustrate the stochastic semantics further consider the network of two priced timed automata in Fig. 16 modeling a competition between the two players Axel and Alex both having to hammer three nails down. As can be seen by the representing **Work**-locations the time (-interval) and rate of energy-consumption required for hammering a nail depends on the player and the nail-number. As expected Axel is initially quite fast and uses a lot of energy but becomes slow towards the last nail, somewhat in contrast to Alex. To make it an interesting competition, there is only *one* hammer illustrated by repeated competitions between the two players in the **Ready**-locations, where the slowest player has to wait in the **Idle**-location until the faster player has finished hammering the next nail. Interestingly, despite the somewhat different strategy applied, the best- and worst-case completion times are identical for Axel and Alex:

⁴otherwise a specific weight distribution can be specified and used instead.

Figure 16.: 3-Nail Hammer Game between Axel and Alex.

59 seconds and 150 seconds. So, there is no difference between the two players and their strategy, or is there?

Assume now that a third person wants to bet on who is the more likely winner – Axel or Alex – given a refined semantics, where the time-delay before performing an output is chosen stochastically (e.g. by drawing from a uniform distribution) and independently by each player (component).

Under such a refined semantics there is a significant difference between the two players (Axel and Alex) in the Hammer Game. In Fig. 17a) the probability distributions for either of the two players winning before a certain time is given. Though it is clear that Axel has a higher probability of winning than Alex (59% versus 41%) given unbounded time, declaring the competition a draw if it has not finished before 50 seconds actually makes Alex the more likely winner. Similarly, Fig. 17b) illustrates the probability of either of the two players winning given an upper bound on energy. With an unlimited amount of energy, clearly Axel is the most likely winner, whereas limiting the consumption of energy to maximum 52 “energy-units” gives Alex an advantage.

Figure 17.: Time- and Cost-dependent Probability of winning the Hammer Game

5.2. Verifying Queries using Statistical Model Checking

Following [86], the measure P_A may be extended in a standard and unique way to the σ -algebra generated by the sets of runs (so-called cylinders) $\pi(s, a_1 a_2 \dots a_n)$. As we shall see this will allow us to give proper semantics to a range of probabilistic time- and cost-constrained temporal properties. Let $\mathcal{A}$ be a NPTA. Then we consider the following non-nested PWCTL properties:

Figure 18.: Cumulative probabilities for time and cost-bounded reachability of T_3 .

$$\psi ::= P(\Diamond_{C \leq c} \varphi) \sim p \mid P(\Box_{C \leq c} \varphi) \sim p$$

where C is an observer clock (of $\mathcal{A}$), φ a state-property (wrt. $\mathcal{A}$), $\sim \in \{<, \leq, =, \geq, >\}$, and $p \in [0, 1]$. This logic is a stochastic extension of the classical WCTL logic for non-stochastic systems, where the existential quantifier is replaced by a probability operator. For the semantics let $\mathcal{A}^*$ be the modification of $\mathcal{A}$, where the guard $C \leq c$ has been conjoined to the invariant of all locations and an edge $(\ell, \varphi, o_\varphi, \emptyset, \ell)$ has been added to all locations ℓ , where o_φ is a new output action. Then:

$$\mathcal{A} \models P(\Diamond_{C \leq c} \varphi) \sim p \text{ iff } P_{\mathcal{A}^*} \left(\bigcup_{\sigma \in \Sigma^*} \pi(s_0, \sigma o_\varphi) \right) \sim p$$

which is well-defined since the σ -algebra on which $P_{\mathcal{A}^*}$ is defined is closed under countable unions and finite intersections. To complete the semantics, we note that $P(\Box_{C \leq c} \varphi) \sim p$ is equivalent to $(1 - p) \sim P(\Diamond_{C \leq c} \neg \varphi)$.⁵

Compared with previous stochastic semantics of timed automata (see e.g., [14, 30]), we emphasize the novelty of the semantics of NPTA in terms of RACES between components, truthfully reflecting their independencies. In particular our stochastic semantics of a network $(A_1 | \dots | A_n)$ is significantly different from that obtained by applying the stochastic semantics of [14, 30] to a product construction $A_1 A_2 \dots A_n$, as information about independencies are lost. So though $(A_1 | \dots | A_n)$ and $A_1 A_2 \dots A_n$ are timed bisimilar they are in general not probabilistic timed bisimilar, and hence distinguishable by PWCTL. The situation is illustrated with the following example.

Example 5.1. Reconsider the Example of Fig. 14. Then it can be shown that $(A|B|T) \models P(\Diamond_{t \leq 2} T_3) = 0.75$ and $(A|B|T) \models P(\Diamond_{C \leq 6} T_3) = 0.75$, whereas $(AB|T) \models P(\Diamond_{t \leq 2} T_3) = 0.50$ and $(AB|T) \models P(\Diamond_{C \leq 6} T_3) = 0.50$. Fig. 18 gives a time- and cost-bounded reachability probabilities for $(A|B|T)$ and $(AB|T)$ for a range of bounds. Thus, though the two NPTAs satisfy the same WCTL properties, they are obviously quite different with respect to PWCTL. The NPTA B_r of Fig. 14 is a variant of B , with the uniform delay distribution enforced by the

⁵We also note that the above (stochastic) interpretation of PWCTL is a conservative extension of the classical (non-stochastic) interpretation of WCTL, in the sense that $\mathcal{A} \models P(\Diamond_{C \leq c} \varphi) > 0$ implies $\mathcal{A}_n \models E\Diamond_{C \leq c} \varphi$, where $\mathcal{A}_n$ refers to the standard non-stochastic semantics of $\mathcal{A}$.

invariant $y \leq 2$ being replaced by an exponential distribution with rate $\frac{1}{2}$. Here $(A|B_r|T)$ satisfies $P(\Diamond_{t \leq 2} T_3) \approx 0.41$ and $P(\Diamond_{C \leq 6} T_3) \approx 0.49$.

The problem of checking $P_M(\psi) \geq p$ ($p \in [0, 1]$) for a PWCTL property ψ is unfortunately undecidable in general ⁶. Our solution is to approximate the answer using simulation-based algorithms known under the name of statistical model checking algorithms. We briefly recap statistical algorithms permitting to answer the following three types of questions:

1. *Hypothesis Testing*: Is the probability $P_M(\psi)$ for a given NPTA M greater or equal to a certain threshold $p \in [0, 1]$?
2. *Probability evaluation*: What is the probability $P_M(\psi)$ for a given NPTA M ?
3. *Probability comparison*: Is the probability $P_M(\psi_1)$ greater than the probability $P_M(\psi_2)$?

From a conceptual point of view solving the above questions using SMC is simple. First, each run of the system is encoded as a Bernoulli random variable that is true if the run satisfies the property and false otherwise. Then a statistical algorithm groups the observations to answer the three questions. For the qualitative questions (1 and 3), we shall use sequential hypothesis testing, while for the quantitative question (2) we will use an estimation algorithm that resemble the classical Monte Carlo simulation. The two solutions are detailed hereafter.

Hypothesis Testing This approach reduces the qualitative question to testing the hypothesis $H : p = P_M(\psi) \geq \theta$ against $K : p < \theta$. To bound the probability of making errors, we use strength parameters α and β and we test the hypothesis $H_0 : p \geq p_0$ and $H_1 : p \leq p_1$ with $p_0 = \theta + \delta_0$ and $p_1 = \theta - \delta_1$. The interval $p_0 - p_1$ defines an indifference region, and p_0 and p_1 are used as thresholds in the algorithm. The parameter α is the probability of accepting H_0 when H_1 holds (false positives) and the parameter β is the probability of accepting H_1 when H_0 holds (false negatives). The above test can be solved by using Wald's *sequential hypothesis testing* [94]. This test computes a proportion r among those runs that satisfy the property. With probability 1, the value of the proportion will eventually cross $\log(\beta/(1-\alpha))$ or $\log((1-\beta)/\alpha)$ and one of the two hypothesis will be selected.

Probability Estimation This algorithm [63] computes the number of runs needed in order to produce an approximation interval $[p - \varepsilon, p + \varepsilon]$ for $p = Pr(\psi)$ with a confidence $1 - \alpha$. The values of ε and α are chosen by the user and the number of runs relies on the Chernoff-Hoeffding bound.

Probability Comparison This algorithm, which is detailed in [48], exploits an extended Wald testing.

5.3. UPPAAL-SMC

We have implemented the above model and algorithms in a statistical extension of UPPAAL called UPPAAL-SMC. In addition to the features exposed above, the

⁶Exceptions being PTA with 0 or 1 clocks.

tool also proposes a friendly-user interface to plot results of estimating distributions as well as a distributed engine to exploit computer grids. Details on UPPAAL-SMC can be found in [48, 49]. As an illustration, here is how we translate the SMC queries from previous section in UPPAAL-SMC.

- Hypothesis testing: $\Pr[bound](\varphi) \geq p_0$, where *bound* defines how to bound the runs. The three ways to bound them are 1) implicitly by time by specifying $\leq M$ (where M is a positive integer), 2) explicitly by cost with $x \leq M$ where x is a specific clock, or 3) by number of discrete steps with $\# \leq M$. In the case of hypothesis testing p_0 is the probability to test for. The formula φ is either $\langle \rangle q$ or $\Box q$ where q is a state predicate.
- Estimation: $\Pr[bound](\varphi)$
- Comparison: $\Pr[bound_1](\varphi_1) \geq \Pr[bound_2](\varphi_2)$.

5.4. Some Illustrations

We briefly survey some recent results obtained via UPPAAL-SMC.

Robot Control In [39] we considered a case – explored in [15] – of a robot moving on a two-dimensional grid. We are interested in the probability that the robot reaches its goal location without staying on consecutive fire fields for more than one time unit and on consecutive ice fields for more than two time units. We applied UPPAAL-SMC to compute the probability of the robot reaching this, without staying more than x time units in some fixed position.

Bluetooth [83] is a wireless telecommunication protocol using frequency-hopping to cope with interference between the devices in the wireless network. In paper [49] we adopted the model from [51], annotated the model to record the power utilization and evaluated the probability distributions of likely response times and energy consumption.

Lightweight Medium Access Protocol (LMAC) [93] LMAC is a communication scheduling protocol based on time slot distribution for nodes sharing the same medium. The protocol is designed having wireless sensor networks in mind: it is simple enough to fit on a modest hardware and at the same time robust against topology reconfiguration, minimizing collisions and power consumption. In [48] we showed how collisions can be analyzed and power consumption estimated using statistical model checking techniques.

Computing Nash Equilibrium in Wireless Ad Hoc Networks One of the important aspects in designing wireless ad-hoc networks is to make sure that a network is robust to the selfish behavior of its participants, i.e. that its configuration satisfies Nash equilibrium (NE). In [40] we proposed an SMC-based algorithm for computing NE for the case when network nodes are modeled by SPTA and an utility function of a single node is equal to a probability that the node will reach its goal.

Energy aware Buildings In [47], we considered energy aware buildings. We refer to a recently developed framework including components for layout of buildings, availability of heaters, climate and user behaviour allowing to evaluate different

strategies for distributing heaters among rooms in terms of the resulting comfort and energy consumption. To indicate central parts of this framework and the clear advantages of modeling the evolution of room temperatures with ODEs, we illustrated in [47] the framework with a small instance comprising two rooms with a single shared heater.

Systems Biology In [46, 47], we extended our model in order to incorporate ODEs. We then showed how the combination of ODEs and SMC allows us to reason on biological oscillations – a problem that is beyond the scope of most existing formal verification techniques. We model a genetic circadian oscillator, which is used to distill the essence of several real circadian oscillators.

Duration Probabilistic Automata In [48] we compared UPPAAL-SMC to Prism [69] in the context of Duration Probabilistic Automata (DPA) [81]. A Duration Probabilistic Automaton (DPA) is a composition of Simple Duration Probabilistic Automata (SDPA). An SDPA is a linear sequence of tasks that must be performed in a sequential order. Each task is associated with a duration interval which gives the possible durations of the task. The actual duration of the tasks is given by a uniform choice from this interval. The comparison with Prism was made by randomly generating models with a specific number of SDPAs and a specific number of tasks per SDPA and translate these into Prism and UPPAAL models. The queries to the models were *What is the probability of all SDPAs ending within t time units* (Estimation) and *Is the probability that all SDPAs end within t time units greater than 40%* (Hypothesis testing). The value of t is different for each model as it was computed by simulating the system 369 times and represent the value for which at least 60% of the runs finished all their tasks.

References

- [1] Yasmina Abdeddaïm, Abdelkarim Kerbaa, and Oded Maler. Task graph scheduling using timed automata. In *IPDPS*, page 237. IEEE Computer Society, 2003.
- [2] Luca Aceto, Anna Ingólfssdóttir, Kim G. Larsen, and Jiří Srba. *Reactive Systems: Modeling, Specification and Verification*. Cambridge University Press, 2007.
- [3] Xavier Allamigeon, Stephane Gaubert, and Eric Goubault. Inferring min and max invariants using max-plus polyhedra. In María Alpuente and Germán Vidal, editors, *SAS*, volume 5079 of *Lecture Notes in Computer Science*, pages 189–204. Springer-Verlag, 2008.
- [4] R. Alur, C. Courcoubetis, N. Halbwachs, T. 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.
- [5] Rajeev Alur. Timed automata. In Halbwachs and Peled [59], pages 8–22.
- [6] Rajeev Alur, Mikhail Bernadsky, and P. Madhusudan. Optimal reachability for weighted timed games. In Josep Díaz, Juhani Karhumäki, Arto Lepistö, and Donald Sannella, editors, *ICALP*, volume 3142 of *Lecture Notes in Computer Science*, pages 122–133. Springer-Verlag, 2004.
- [7] Rajeev Alur, Costas Courcoubetis, and David L. Dill. Model-checking for real-time systems. In *LICS*, pages 414–425. IEEE Computer Society, 1990.
- [8] Rajeev Alur and David L. Dill. Automata for modeling real-time systems. In Mike Paterson, editor, *ICALP*, volume 443 of *Lecture Notes in Computer Science*, pages 322–335. Springer-Verlag, 1990.
- [9] Rajeev Alur and David L. Dill. A theory of timed automata. *Theor. Comput. Sci.*, 126(2):183–235, 1994.

- [10] Rajeev Alur, Salvatore La Torre, and George J. Pappas. Optimal paths in weighted timed automata. In Benedetto and Sangiovanni-Vincentelli [26], pages 49–62.
- [11] D. Andersson. Improved combinatorial algorithms for discounted payoff games. Master's thesis, Uppsala University, Department of Information Technology, 2006.
- [12] Eugene Asarin, Marius Bozga, Alain Kerbrat, Oded Maler, Amir Pnueli, and Anne Rasse. Data-structures for the verification of timed automata. In Oded Maler, editor, *HART*, volume 1201 of *Lecture Notes in Computer Science*, pages 346–360. Springer-Verlag, 1997.
- [13] Eugene Asarin, Oded Maler, and Amir Pnueli. Symbolic controller synthesis for discrete and timed systems. In Panos J. Antsaklis, Wolf Kohn, Anil Nerode, and Shankar Sastry, editors, *Hybrid Systems*, volume 999 of *Lecture Notes in Computer Science*, pages 1–20. Springer-Verlag, 1994.
- [14] Christel Baier, Nathalie Bertrand, Patricia Bouyer, Thomas Brihaye, and Marcus Größer. Probabilistic and topological semantics for timed automata. In *FSTTCS*, volume 4855 of *LNCS*, pages 179–191. Springer, 2007.
- [15] Benoît Barbot, Taolue Chen, Tingting Han, Joost-Pieter Katoen, and Alexandru Mereacre. Efficient CTMC model checking of linear real-time objectives. In Parosh A. Abdulla and K. Rustan M. Leino, editors, *TACAS*, volume 6605 of *Lecture Notes in Computer Science*, pages 128–142. Springer, 2011.
- [16] Gerd Behrmann, Johan Bengtsson, Alexandre David, Kim G. Larsen, Paul Pettersson, and Wang Yi. UPPAAL implementation secrets. In Werner Damm and Ernst-Rüdiger Olderog, editors, *FTRTFT*, volume 2469 of *Lecture Notes in Computer Science*, pages 3–22. Springer-Verlag, 2002.
- [17] Gerd Behrmann, Patricia Bouyer, Kim G. Larsen, and Radek Pelánek. Lower and upper bounds in zone based abstractions of timed automata. In Jensen and Podolski [66], pages 312–326.
- [18] Gerd Behrmann, Ed Brinksma, Martijn Hendriks, and Angelika Mader. Production scheduling by reachability analysis - a case study. In *IPDPS*. IEEE Computer Society, 2005.
- [19] Gerd Behrmann, Agnès Cougnard, Alexandre David, Emmanuel Fleury, Kim G. Larsen, and Didier Lime. UPPAAL-TIGA: Time for playing games! In Werner Damm and Holger Hermanns, editors, *CAV*, volume 4590 of *Lecture Notes in Computer Science*, pages 121–125. Springer-Verlag, 2007.
- [20] Gerd Behrmann, Alexandre David, and Kim G. Larsen. A tutorial on UPPAAL. In Marco Bernardo and Flavio Corradini, editors, *SFM*, volume 3185 of *Lecture Notes in Computer Science*, pages 200–236. Springer-Verlag, 2004.
- [21] Gerd Behrmann, Ansgar Fehnker, Thomas Hune, Kim G. Larsen, Paul Pettersson, Judi Romijn, and Frits W. Vaandrager. Minimum-cost reachability for priced timed automata. In Benedetto and Sangiovanni-Vincentelli [26], pages 147–161.
- [22] Gerd Behrmann, Thomas Hune, and Frits W. Vaandrager. Distributing timed model checking - how the search order matters. In E. Allen Emerson and A. Prasad Sistla, editors, *CAV*, volume 1855 of *Lecture Notes in Computer Science*, pages 216–231. Springer-Verlag, 2000.
- [23] Gerd Behrmann, Kim G. Larsen, Justin Pearson, Carsten Weise, and Wang Yi. Efficient timed reachability analysis using clock difference diagrams. In Halbwachs and Peled [59], pages 341–353.
- [24] Gerd Behrmann, Kim G. Larsen, and Radek Pelánek. To store or not to store. In Warren A. Hunt Jr. and Fabio Somenzi, editors, *CAV*, volume 2725 of *Lecture Notes in Computer Science*, pages 433–445. Springer-Verlag, 2003.
- [25] Gerd Behrmann, Kim G. Larsen, and Jacob Illum Rasmussen. Optimal scheduling using priced timed automata. *SIGMETRICS Performance Evaluation Review*, 32(4):34–40, 2005.
- [26] Maria Domenica Di Benedetto and Alberto L. Sangiovanni-Vincentelli, editors. *Hybrid Systems: Computation and Control, 4th International Workshop, HSCC 2001, Rome, Italy, March 28-30, 2001, Proceedings*, volume 2034 of *Lecture Notes in Computer Science*. Springer-Verlag, 2001.
- [27] Johan Bengtsson and Wang Yi. On clock difference constraints and termination in reacha-

- bility analysis of timed automata. In Jin Song Dong and Jim Woodcock, editors, *ICFEM*, volume 2885 of *Lecture Notes in Computer Science*, pages 491–503. Springer-Verlag, 2003.
- [28] Johan Bengtsson and Wang Yi. Timed automata: Semantics, algorithms and tools. In Jörg Desel, Wolfgang Reisig, and Grzegorz Rozenberg, editors, *Lectures on Concurrency and Petri Nets*, volume 3098 of *Lecture Notes in Computer Science*, pages 87–124. Springer-Verlag, 2003.
- [29] Béatrice Bérard, Antoine Petit, Volker Diekert, and Paul Gastin. Characterization of the expressive power of silent transitions in timed automata. *Fundam. Inform.*, 36(2–3):145–182, 1998.
- [30] Nathalie Bertrand, Patricia Bouyer, Thomas Brihaye, and Nicolas Markey. Quantitative model-checking of one-clock timed automata under probabilistic semantics. In *QEST*, pages 55–64. IEEE Computer Society, 2008.
- [31] H. Bohnenkamp, P.R. D’Argenio, H. Hermanns, and J.-P. Katoen. Modest: A compositional modeling formalism for real-time and stochastic systems. Technical Report CTIT 04-46, University of Twente, 2004.
- [32] Patricia Bouyer. Untameable timed automata! In Helmut Alt and Michel Habib, editors, *STACS*, volume 2607 of *Lecture Notes in Computer Science*, pages 620–631. Springer-Verlag, 2003.
- [33] Patricia Bouyer, Thomas Brihaye, and Nicolas Markey. Improved undecidability results on weighted timed automata. *Inf. Process. Lett.*, 98(5):188–194, 2006.
- [34] Patricia Bouyer, Ed Brinksma, and Kim G. Larsen. Staying alive as cheaply as possible. In Rajeev Alur and George J. Pappas, editors, *HSCC*, volume 2993 of *Lecture Notes in Computer Science*, pages 203–218. Springer-Verlag, 2004.
- [35] Patricia Bouyer, Franck Cassez, Emmanuel Fleury, and Kim G. Larsen. Optimal strategies in priced timed game automata. In Kamal Lodaya and Meena Mahajan, editors, *FSTTCS*, volume 3328 of *Lecture Notes in Computer Science*, pages 148–160. Springer-Verlag, 2004.
- [36] Patricia Bouyer, Kim G. Larsen, Nicolas Markey, and Jacob Iltis Rasmussen. Almost optimal strategies in one clock priced timed games. In S. Arun-Kumar and Naveen Garg, editors, *FSTTCS*, volume 4337 of *Lecture Notes in Computer Science*, pages 345–356. Springer-Verlag, 2006.
- [37] Marius Bozga, Conrado Daws, Oded Maler, Alfredo Olivero, Stavros Tripakis, and Sergio Yovine. Kronos: A model-checking tool for real-time systems. In Alan J. Hu and Moshe Y. Vardi, editors, *CAV*, volume 1427 of *Lecture Notes in Computer Science*, pages 546–550. Springer-Verlag, 1998.
- [38] Thomas Brihaye, Véronique Bruyère, and Jean-François Raskin. On optimal timed strategies. In Pettersson and Yi [87], pages 49–64.
- [39] Peter E. Bulychyev, Alexandre David, Kim G. Larsen, Axel Legay, Guangyuan Li, Danny Bøgsted Poulsen, and Amélie Stainer. Monitor-based statistical model checking for weighted metric temporal logic. In Nikolaj Bjørner and Andrei Voronkov, editors, *LPAR*, volume 7180 of *Lecture Notes in Computer Science*, pages 168–182. Springer, 2012.
- [40] Peter E. Bulychyev, Alexandre David, Kim G. Larsen, Axel Legay, and Marius Mikučionis. Computing Nash equilibrium in wireless ad hoc networks: A simulation-based approach. In Johannes Reich and Bernd Finkbeiner, editors, *IWIGP*, volume 78 of *EPTCS*, pages 1–14, 2012.
- [41] Franck Cassez, Alexandre David, Emmanuel Fleury, Kim G. Larsen, and Didier Lime. Efficient on-the-fly algorithms for the analysis of timed games. In Martín Abadi and Luca de Alfaro, editors, *CONCUR*, volume 3653 of *Lecture Notes in Computer Science*, pages 66–80. Springer-Verlag, 2005.
- [42] Franck Cassez, Alexandre David, Kim G. Larsen, Didier Lime, and Jean-François Raskin. Timed control with observation based and stuttering invariant strategies. In Kedar S. Namjoshi, Tomohiro Yoneda, Teruo Higashino, and Yoshio Okamura, editors, *ATVA*, volume 4762 of *Lecture Notes in Computer Science*, pages 192–206. Springer-Verlag, 2007.
- [43] Kārlis Čerāns. Decidability of bisimulation equivalences for parallel timer processes. In Gregor von Bochmann and David K. Probst, editors, *CAV*, volume 663 of *Lecture Notes in Computer Science*, pages 302–315. Springer-Verlag, 1992.
- [44] Costas Courcoubetis and Mihalis Yannakakis. Minimum and maximum delay problems in

- real-time systems. In Larsen and Skou [78], pages 399–409.
- [45] Pedro R. D’Argenio, Joost-Pieter Katoen, Theo C. Ruys, and Jan Tretmans. The bounded retransmission protocol must be on time! In Ed Brinksma, editor, *TACAS*, volume 1217 of *Lecture Notes in Computer Science*, pages 416–431. Springer-Verlag, 1997.
 - [46] Alexandre David, Dehui Du, Kim G. Larsen, Axel Legay, Marius Mikucionis, Danny Bøgsteds Poulsen, and Sean Sedwards. Statistical model checking for stochastic hybrid systems. In *HSB*, volume 92 of *EPTCS*, pages 122–136, 2012.
 - [47] Alexandre David, Kim G. Larsen, Axel Legay, Marius Mikucionis, Danny Bøgsteds Poulsen, and Sean Sedwards. Runtime verification of biological systems. In *ISoLA (1)*, volume 7609 of *Lecture Notes in Computer Science*, pages 388–404. Springer, 2012.
 - [48] Alexandre David, Kim G. Larsen, Axel Legay, Marius Mikucionis, Danny Bøgsteds Poulsen, Jonas van Vliet, and Zheng Wang. Statistical model checking for networks of priced timed automata. In Uli Fahrenberg and Stavros Tripakis, editors, *FORMATS*, volume 6919 of *Lecture Notes in Computer Science*, pages 80–96. Springer, 2011.
 - [49] Alexandre David, Kim G. Larsen, Axel Legay, Marius Mikucionis, and Zheng Wang. Time for statistical model checking of real-time systems. In Ganesh Gopalakrishnan and Shaz Qadeer, editors, *Computer Aided Verification*, volume 6806 of *Lecture Notes in Computer Science*, pages 349–355. Springer, 2011.
 - [50] David L. Dill. Timing assumptions and verification of finite-state concurrent systems. In Joseph Sifakis, editor, *Automatic Verification Methods for Finite State Systems*, volume 407 of *Lecture Notes in Computer Science*, pages 197–212. Springer-Verlag, 1989.
 - [51] Marie Duflot, Marta Kwiatkowska, Gethin Norman, and David Parker. A formal analysis of bluetooth device discovery. *International Journal on Software Tools for Technology Transfer (STTT)*, 8:621–632, 2006.
 - [52] Juhan P. Ernits. Memory arbiter synthesis and verification for a radar memory interface card. *Nord. J. Comput.*, 12(2):68–88, 2005.
 - [53] Uli Fahrenberg and Kim G. Larsen. Discount-optimal infinite runs in priced timed automata. *Electr. Notes Theor. Comput. Sci.*, 239:179–191, 2009.
 - [54] Uli Fahrenberg and Kim G. Larsen. Discounting in time. *Electr. Notes Theor. Comput. Sci.*, 253(3):25–31, 2009.
 - [55] Uli Fahrenberg, Kim G. Larsen, and Claus Thrane. Verification, performance analysis and controller synthesis for real-time systems. In Manfred Broy, Wassiou Sitou, and Tony Hoare, editors, *ASI 08*, volume 22 of *NATO Science for Peace and Security Series - D: Information and Communication Security*. IOS Press BV, 2008.
 - [56] Uli Fahrenberg, Kim G. Larsen, and Claus Thrane. Verification, performance analysis and controller synthesis for real-time systems. In Farhad Arbab and Marjan Sirjani, editors, *FSEN*, volume 5961 of *Lecture Notes in Computer Science*, pages 34–61. Springer, 2009.
 - [57] Uli Fahrenberg, Kim G. Larsen, and Claus Thrane. Model-based verification and analysis for real-time systems. In Manfred Broy, Christian Leuxner, and Tony Hoare, editors, *Software and Systems Safety - Specification and Verification*, volume 30 of *NATO Science for Peace and Security Series - D: Information and Communication Security*, pages 231–259. IOS Press, 2011.
 - [58] Ansgar Fehnker. Scheduling a steel plant with timed automata. In *RTCSA*, pages 280–286. IEEE Computer Society, 1999.
 - [59] Nicolas Halbwachs and Doron Peled, editors. *Computer Aided Verification, 11th International Conference, CAV ’99, Trento, Italy, July 6-10, 1999, Proceedings*, volume 1633 of *Lecture Notes in Computer Science*. Springer-Verlag, 1999.
 - [60] Michael R. Hansen, Jan Madsen, and Aske Wiid Brekling. Semantics and verification of a language for modelling hardware architectures. In Cliff B. Jones, Zhiming Liu, and Jim Woodcock, editors, *Formal Methods and Hybrid Real-Time Systems*, volume 4700 of *Lecture Notes in Computer Science*, pages 300–319. Springer-Verlag, 2007.
 - [61] Martijn Hendriks. Model checking the time to reach agreement. In Pettersson and Yi [87], pages 98–111.
 - [62] Thomas A. Henzinger, Xavier Nicollin, Joseph Sifakis, and Sergio Yovine. Symbolic model checking for real-time systems. *Inf. Comput.*, 111(2):193–244, 1994.
 - [63] Thomas Hérault, Richard Lassaigne, Frédéric Magniette, and Sylvain Peyronnet. Ap-

- proximate probabilistic model checking. In *VMCAI*, volume 2937 of *LNCS*, pages 73–84. Springer, 2004.
- [64] Thomas Hune, Kim G. Larsen, and Paul Pettersson. Guided synthesis of control programs using UPPAAL. *Nord. J. Comput.*, 8(1):43–64, 2001.
 - [65] Henrik Ejersbo Jensen, Kim G. Larsen, and Arne Skou. Scaling up UPPAAL automatic verification of real-time systems using compositionality and abstraction. In Mathai Joseph, editor, *FTRTFT*, volume 1926 of *Lecture Notes in Computer Science*, pages 19–30. Springer-Verlag, 2000.
 - [66] Kurt Jensen and Andreas Podelski, editors. *Tools and Algorithms for the Construction and Analysis of Systems, 10th International Conference, TACAS 2004, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2004, Barcelona, Spain, March 29 - April 2, 2004, Proceedings*, volume 2988 of *Lecture Notes in Computer Science*. Springer-Verlag, 2004.
 - [67] Jan Jakob Jensen, Jacob Illum Rasmussen, Kim G. Larsen, and Alexandre David. Guided controller synthesis for climate controller using UPPAAL-TIGA. In Jean-François Raskin and P. S. Thiagarajan, editors, *FORMATS*, volume 4763 of *Lecture Notes in Computer Science*, pages 227–240. Springer-Verlag, 2007.
 - [68] Richard M. Karp. A characterization of the minimum cycle mean in a digraph. *Disc. Math.*, 23(3):309–311, 1978.
 - [69] Marta Z. Kwiatkowska, Gethin Norman, and David Parker. PRISM 2.0: A tool for probabilistic model checking. In *QEST*, pages 322–323. IEEE, 2004.
 - [70] Leslie Lamport. Real-time model checking is really simple. In Dominique Borriore and Wolfgang J. Paul, editors, *CHARME*, volume 3725 of *Lecture Notes in Computer Science*, pages 162–175. Springer-Verlag, 2005.
 - [71] Kim G. Larsen, Gerd Behrmann, Ed Brinksma, Ansgar Fehnker, Thomas Hune, Paul Pettersson, and Judi Romijn. As cheap as possible: Efficient cost-optimal reachability for priced timed automata. In Gérard Berry, Hubert Comon, and Alain Finkel, editors, *CAV*, volume 2102 of *Lecture Notes in Computer Science*, pages 493–505. Springer-Verlag, 2001.
 - [72] Kim G. Larsen, Fredrik Larsson, Paul Pettersson, and Wang Yi. Efficient verification of real-time systems: compact data structure and state-space reduction. In *IEEE Real-Time Systems Symposium*, pages 14–24. IEEE Computer Society, 1997.
 - [73] Kim G. Larsen, Marius Mikucionis, Brian Nielsen, and Arne Skou. Testing real-time embedded software using UPPAAL-TRON: an industrial case study. In Wayne Wolf, editor, *EMSOFT*, pages 299–306. ACM, 2005.
 - [74] Kim G. Larsen, Justin Pearson, Carsten Weise, and Wang Yi. Clock difference diagrams. *Nord. J. Comput.*, 6(3):271–298, 1999.
 - [75] Kim G. Larsen, Paul Pettersson, and Wang Yi. UPPAAL in a nutshell. *STTT*, 1(1–2):134–152, 1997.
 - [76] Kim G. Larsen and Jacob Illum Rasmussen. Optimal conditional reachability for multi-priced timed automata. In Vladimiro Sassone, editor, *FoSSaCS*, volume 3441 of *Lecture Notes in Computer Science*, pages 234–249. Springer-Verlag, 2005.
 - [77] Kim G. Larsen and Jacob Illum Rasmussen. Optimal reachability for multi-priced timed automata. *Theor. Comput. Sci.*, 390(2-3):197–213, 2008.
 - [78] Kim G. Larsen and Arne Skou, editors. *Computer Aided Verification, 3rd International Workshop, CAV '91, Aalborg, Denmark, July, 1-4, 1991, Proceedings*, volume 575 of *Lecture Notes in Computer Science*. Springer-Verlag, 1992.
 - [79] Magnus Lindahl, Paul Pettersson, and Wang Yi. Formal design and analysis of a gear controller. In Bernhard Steffen, editor, *TACAS*, volume 1384 of *Lecture Notes in Computer Science*, pages 281–297. Springer-Verlag, 1998.
 - [80] Oded Maler. Timed automata as an underlying model for planning and scheduling. In Maria Fox and Alexandra M. Coddington, editors, *AIPS Workshop on Planning for Temporal Domains*, pages 67–70, 2002.
 - [81] Oded Maler, Kim G. Larsen, and Bruce H. Krogh. On zone-based analysis of duration probabilistic automata. In *INFINITY*, volume 39 of *EPTCS*, pages 33–46, 2010.
 - [82] Oded Maler, Amir Pnueli, and Joseph Sifakis. On the synthesis of discrete controllers for timed systems (an extended abstract). In *STACS*, pages 229–242, 1995.

- [83] P. McDermott-Wells. What is bluetooth? *Potentials, IEEE*, 23(5):33 – 35, 2004-jan. 2005.
- [84] Jesper B. Møller, Jakob Lichtenberg, Henrik Reif Andersen, and Henrik Hulgaard. Difference decision diagrams. In Jörg Flum and Mario Rodríguez-Artalejo, editors, *CSL*, volume 1683 of *Lecture Notes in Computer Science*, pages 111–125. Springer-Verlag, 1999.
- [85] Joël Ouaknine and James Worrell. Universality and language inclusion for open and closed timed automata. In Oded Maler and Amir Pnueli, editors, *HSCC*, volume 2623 of *Lecture Notes in Computer Science*, pages 375–388. Springer-Verlag, 2003.
- [86] P. Panangaden. *Labelled Markov Processes*. Imperial College Press, 2010.
- [87] Paul Pettersson and Wang Yi, editors. *Formal Modeling and Analysis of Timed Systems, Third International Conference, FORMATS 2005, Uppsala, Sweden, September 26-28, 2005, Proceedings*, volume 3829 of *Lecture Notes in Computer Science*. Springer-Verlag, 2005.
- [88] Jacob Illum Rasmussen, Kim G. Larsen, and K. Subramani. Resource-optimal scheduling using priced timed automata. In Jensen and Podelski [66], pages 220–235.
- [89] Koushik Sen, Mahesh Viswanathan, and Gul Agha. Statistical model checking of black-box probabilistic systems. In *CAV*, LNCS 3114, pages 202–215. Springer, 2004.
- [90] Colin Stirling. Modal and temporal logics for processes. In *Proc. Banff Higher Order Workshop*, volume 1043 of *Lecture Notes in Computer Science*, pages 149–237. Springer-Verlag, 1995.
- [91] Tino Teige, Andreas Eggers, and Martin Fränzle. Constraint-based analysis of concurrent probabilistic hybrid systems: An application to networked automation systems. *Nonlinear Analysis: Hybrid Systems*, 2011.
- [92] Stavros Tripakis and Karine Altisen. On-the-fly controller synthesis for discrete and dense-time systems. In Jeannette M. Wing, Jim Woodcock, and Jim Davies, editors, *World Congress on Formal Methods*, volume 1708 of *Lecture Notes in Computer Science*, pages 233–252. Springer-Verlag, 1999.
- [93] Lodewijk F.W. van Hoesel and Paul J.M. Havinga. A lightweight medium access protocol (LMAC) for wireless sensor networks: Reducing preamble transmissions and transceiver state switches. In *1st International Workshop on Networked Sensing Systems (INSS'04)*, pages 205–208, Tokyo, Japan, 2004. Society of Instrument and Control Engineers (SICE).
- [94] Abraham Wald. *Sequential Analysis*. Dover Publications, 2004.
- [95] Wang Yi, Paul Pettersson, and Mats Daniels. Automatic verification of real-time communicating systems by constraint-solving. In Dieter Hogrefe and Stefan Leue, editors, *FORTE*, volume 6 of *IFIP Conference Proceedings*, pages 243–258. Chapman & Hall, 1994.
- [96] Håkan L. S. Younes. *Verification and Planning for Stochastic Processes with Asynchronous Events*. PhD thesis, Carnegie Mellon, 2005.