

# Model checking for probabilistic timed automata

Gethin Norman · David Parker · Jeremy Sproston

Published online: 12 October 2012  
© Springer Science+Business Media New York 2012

**Abstract** Probabilistic timed automata (PTAs) are a formalism for modelling systems whose behaviour incorporates both probabilistic and real-time characteristics. Applications include wireless communication protocols, automotive network protocols and randomised security protocols. This paper gives an introduction to PTAs and describes techniques for analysing a wide range of quantitative properties, such as “the maximum probability of the airbag failing to deploy within 0.02 seconds”, “the maximum expected time for the protocol to terminate” or “the minimum expected energy consumption required to complete all tasks”. We present a temporal logic for specifying such properties and then give a survey of available model-checking techniques for formulae specified in this logic. We then describe two case studies in which PTAs are used for modelling and analysis: a probabilistic non-repudiation protocol and a task-graph scheduling problem.

**Keywords** Probabilistic timed automata · Probabilistic model checking · Probabilistic real-time systems

## 1 Introduction

Automated verification techniques, such as model checking, provide powerful methods for rigorously analysing the correctness of systems. Increasingly, this analysis must also take

---

G. Norman  
School of Computing Science, University of Glasgow, 18 Lilybank Gardens, Glasgow, G12 8RZ,  
Scotland, UK  
e-mail: [gethin.norman@glasgow.ac.uk](mailto:gethin.norman@glasgow.ac.uk)

D. Parker (✉)  
School of Computer Science, University of Birmingham, Edgbaston, Birmingham, B15 2TT, UK  
e-mail: [d.a.parker@cs.bham.ac.uk](mailto:d.a.parker@cs.bham.ac.uk)

J. Sproston  
Dipartimento di Informatica, Università degli Studi di Torino, Corso Svizzera 185, 10149 Torino, Italy  
e-mail: [sproston@di.unito.it](mailto:sproston@di.unito.it)

into account *quantitative* aspects of the systems being verified, including both *real-time* characteristics and *probabilistic* behaviour. Embedded systems, for example, whether in communication and multimedia devices or in automotive and avionic control systems, often operate under timing constraints. The need for automated formal verification techniques in this domain is clear, as evidenced by the take-up of timed automata verification tools such as UPPAAL [13]. On the other hand, many real-life systems also exhibit stochastic behaviour, due, for example, to component failures, unreliable communication media or the use of randomisation. Probabilistic verification tools such as PRISM [51] and MRMC [45] have been widely used to analyse many systems with stochastic behaviour. Another vital ingredient in system modelling is *nondeterminism*, which is often used to capture concurrency between parallel components and to under-specify or abstract certain aspects of a system.

*Probabilistic timed automata* (PTAs) [11, 35, 54] are a modelling formalism for systems that exhibit probabilistic, nondeterministic *and* real-time characteristics. In many application domains, all three aspects need to be modelled; these include wireless communication protocols such as Bluetooth or Zigbee, automotive network protocols such as FlexRay, randomised security protocols, e.g. for anonymity or non-interference, and many others. The interplay between these different aspects can be subtle, making automated verification techniques and tool support essential. Verification of PTAs permits analysis of a wide range of quantitative properties, from reliability to performance, e.g.:

- “the maximum probability of an airbag failing to deploy within 0.02 seconds”;
- “the minimum probability that a packet is correctly delivered with 1 second”;
- “the maximum expected time for the protocol to terminate”.

PTAs can also be augmented with additional quantitative information in the form of *costs* or *rewards*. The resulting model is sometimes referred to as *priced probabilistic timed automata* and allows reasoning about a wide range of additional properties, e.g.:

- “the maximum expected number of lost packets within the first hour”;
- “the minimum expected energy consumption for completion of all tasks”;
- “the maximum number of queued requests after 10 seconds of operation”.

This paper provides an introduction to PTAs and the techniques that have been developed to specify and verify properties such as those listed above.

*Paper structure* This paper is organised as follows. After background material in Sect. 2, Sect. 3 introduces the model of PTAs and discusses various issues relating to their use and analysis. Then, Sect. 4 presents a probabilistic temporal logic to represent properties of PTAs and Sect. 5 surveys the various techniques that can be used to perform model checking of this logic. Section 6 describes two case studies illustrating the usage of PTAs and their associated model-checking algorithms: a probabilistic non-repudiation protocol and a task-graph scheduling problem. Throughout the paper we give pointers to the relevant literature on PTAs and describe related work.

## 2 Background

We use  $\mathbb{R}_{\geq 0}$  to denote the set of non-negative real numbers,  $\mathbb{Q}_{\geq 0}$  for the set of non-negative rationals and  $\mathbb{N}$  for the set of natural numbers. A discrete probability *distribution* over a countable set  $Q$  is a function  $\mu : Q \rightarrow [0, 1]$  such that  $\sum_{q \in Q} \mu(q) = 1$ . For a function  $\mu : Q \rightarrow \mathbb{R}_{\geq 0}$  we define  $\text{Support}(\mu) = \{q \in Q \mid \mu(q) > 0\}$ . Then, for an arbitrary set  $Q$ , we

define  $Dist(Q)$  to be the set of functions  $\mu : Q \rightarrow [0, 1]$  such that  $Support(\mu)$  is a countable set and  $\mu$  restricted to  $Support(\mu)$  is a distribution. For  $q \in Q$ , let  $\mu_q$  be the *point distribution* at  $q$  which assigns probability 1 to  $q$ . Let  $AP$  be a set of *atomic propositions*, which we assume to be fixed throughout the paper.

*Markov decision processes* (MDPs) are a widely used formalism for modelling systems that exhibit both nondeterministic and probabilistic behaviour. In this paper, we will use *timed probabilistic systems* (TPSs) [52, 56], an extension of MDPs in which transitions are labelled with either an action or a time duration.

**Definition 1** (TPS) A *timed probabilistic systems*  $T$  is a tuple  $(S, \bar{s}, Act, Steps_T, lab)$  where  $S$  is (possibly infinite) a set of *states*,  $\bar{s} \in S$  an *initial state*,  $Act$  a (finite) set of *actions*,  $Steps_T : S \times (Act \cup \mathbb{R}_{\geq 0}) \rightarrow Dist(S)$  a (partial) *probabilistic transition function* and  $lab : S \rightarrow 2^{AP}$  a *labelling function*.

A TPS  $T$  starts in the initial state  $\bar{s}$  and, when in state  $s \in S$ , there is a nondeterministic choice between one or more available actions or time durations  $a \in Act \cup \mathbb{R}_{\geq 0}$  (those for which  $Steps_T(s, a)$  is defined). After the choice of an available action or time duration  $a$ , a successor state  $s'$  is selected at random according to the probability distribution  $Steps_T(s, a)$ . We use the notation  $s \xrightarrow{a, \mu} s'$  for such a transition, i.e., to denote that  $Steps_T(s, a)$  is defined and equal to  $\mu$ , and that  $s' \in Support(\mu)$ . We assume, for each state  $s \in S$ , there exists at least one available action or time duration. An MDP  $M$  is a special case of a TPS where time is omitted from the transition function, i.e., the probabilistic transition function takes the form  $Steps_M : S \times Act \rightarrow Dist(S)$ .

A path of a TPS represents a particular resolution of both the nondeterminism and probability present in the system. Formally, a path is a finite or infinite sequence of probabilistic transitions alternating between time durations and actions, for example:

$$\omega = s_0 \xrightarrow{a_0, \mu_0} s_1 \xrightarrow{a_1, \mu_1} s_2 \xrightarrow{a_2, \mu_2} \dots$$

where  $a_{2i} \in \mathbb{R}_{\geq 0}$  and  $a_{2i+1} \in Act$  for  $i \in \mathbb{N}$ . We denote by  $\omega(i)$  the  $(i + 1)$ th state  $s_i$  of  $\omega$  and the accumulated duration up until this state  $\omega(i)$  is defined by:

$$dur_\omega(i) \stackrel{\text{def}}{=} \sum_{0 \leq j < i \wedge a_j \in \mathbb{R}_{\geq 0}} a_j.$$

A *position* of  $\omega$  is a pair  $(i, t) \in \mathbb{N} \times \mathbb{R}_{\geq 0}$  such that  $t \leq dur_\omega(i + 1) - dur_\omega(i)$ . We say that the position  $(j, t')$  precedes the position  $(i, t)$ , written  $(j, t') < (i, t)$ , when  $j < i$  or  $j = i$  and  $t' < t$ .

To reason about the probabilistic behaviour of a TPS  $T$ , we use the notion of an *adversary*, which is a possible resolution of nondeterminism only. Formally, an adversary is a function from finite paths with an even number of transitions to available time durations, and from finite paths with an odd number of transitions to available actions). For a fixed adversary  $\sigma$  and state  $s$ , we can define a probability measure  $Pr_{T,s}^\sigma$  over the set  $Path_{T,s}^\sigma$  of infinite paths starting in  $s$  corresponding to  $\sigma$  [47]. For a real-valued random variable  $f$  over  $Path_{T,s}^\sigma$ , we let  $\mathbb{E}_{T,s}^\sigma(f)$  denote the expected value of  $f$  with respect to  $Pr_{T,s}^\sigma$ .

We restrict our attention to *time-divergent* (or *non-Zeno*) adversaries, i.e., we do not consider executions in which time does not advance beyond a certain point. These can be ignored on the grounds that they do not correspond to actual, realisable behaviour of the system being modelled [2, 40]. Formally, an adversary  $\sigma$  of a TPS  $T$  is time divergent if:

$$Pr_{T,s}^\sigma(\{\omega \in Path_{T,s}^\sigma \mid \forall c \in \mathbb{N}. \exists i \in \mathbb{N}. dur_\omega(i) > c\}) = 1$$

for all states  $s$  of  $T$ . We denote by  $Adv_T$  the set of all time-divergent adversaries of  $T$ . This issue is discussed in more depth in Sect. 3.

We next introduce *rewards* (or, equivalently, *costs* or *prices*) for TPSs. These are used to represent additional information about the system that the TPS represents, e.g., the number of packets sent or requests lost, the time spent in a particular state or the energy consumed.

**Definition 2** (Reward structure) A *reward structure* for a TPS  $T = (S, \bar{s}, Act, Steps_T, lab)$  is a pair  $r = (r_S, r_{Act})$  where  $r_S : S \rightarrow \mathbb{R}_{\geq 0}$  is a *state reward function* and  $r_{Act} : (S \times Act) \rightarrow \mathbb{R}_{\geq 0}$  is an *action reward function*.

For a reward structure  $r = (r_S, r_{Act})$  and state  $s$ , the value  $r_S(s)$  defines the *rate* (per time unit) at which reward is accumulated when in state  $s$ . On the other hand, for state  $s$  and action  $a$ , the value  $r_{Act}(s, a)$  defines the reward acquired when the action  $a$  is taken in state  $s$ . More formally, for any infinite path  $\omega = s_0 \xrightarrow{a_0, \mu_0} s_1 \xrightarrow{a_1, \mu_1} \dots$ , the reward accumulated during the transition of  $\omega$  from state  $s_i$  to  $s_{i+1}$  is defined by:

$$r(\omega, i) \stackrel{\text{def}}{=} \begin{cases} r_S(s_i) \cdot a_i & \text{if } a_i \in \mathbb{R}_{\geq 0} \text{ (or, equivalently, if } i \bmod 2 = 0), \\ r_{Act}(s_i, a_i) & \text{otherwise.} \end{cases}$$

Alternatively, we can also interpret state rewards as defining a reward at a particular time instant. An example of this usage would be a reward structure that represents the number of messages stored in a queue at a particular time instant. When using this interpretation, action reward values are not considered.

### 3 Probabilistic timed automata

Probabilistic timed automata (PTAs) [11, 42, 54] model real-time behaviour in the same fashion as classical timed automata [4], using *clocks*. Clocks are variables whose values range over the non-negative reals and which increase at the same rate as time. Throughout this paper, we assume a finite set of clocks  $\mathcal{X}$ . A function  $v : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$  is referred to as a *clock valuation* and the set of all clock valuations is denoted by  $\mathbb{R}_{\geq 0}^{\mathcal{X}}$ . For any  $v \in \mathbb{R}_{\geq 0}^{\mathcal{X}}$ ,  $t \in \mathbb{R}_{\geq 0}$  and  $X \subseteq \mathcal{X}$ , we use  $v + t$  to denote the clock valuation which increments all clock values in  $v$  by  $t$  and  $v[X:=0]$  for the clock valuation in which clocks in  $X$  are reset to 0. We use  $\mathbf{0}$  to denote the clock valuation that assigns 0 to all clocks in  $\mathcal{X}$ .

The set of *clock constraints* over  $\mathcal{X}$ , denoted  $CC(\mathcal{X})$ , is defined by the syntax:

$$\chi ::= \text{true} \mid x \leq d \mid c \leq x \mid x + c \leq y + d \mid \neg \chi \mid \chi \wedge \chi$$

where  $x, y \in \mathcal{X}$  and  $c, d \in \mathbb{N}$ . A clock valuation  $v$  satisfies a clock constraint  $\chi$ , denoted by  $v \models \chi$ , if  $\chi$  resolves to  $\text{true}$  when substituting each occurrence of clock  $x$  with  $v(x)$ . The set of valuations satisfying a clock constraint is called a *zone*. Clock constraints will be used in the syntactic definition of PTAs and for the specification of properties.

**Definition 3** (PTA syntax) A *probabilistic timed automaton* (PTA) is defined by a tuple  $P = (L, \bar{l}, \mathcal{X}, Act, inv, enab, prob, \mathcal{L})$  where:

- $L$  is a finite set of *locations* and  $\bar{l} \in L$  is an *initial location*;
- $\mathcal{X}$  is a finite set of *clocks*;

- $Act$  is a finite set of actions;
- $inv : L \rightarrow CC(\mathcal{X})$  is an invariant condition;
- $enab : L \times Act \rightarrow CC(\mathcal{X})$  is an enabling condition;
- $prob : L \times Act \rightarrow Dist(2^{\mathcal{X}} \times L)$  is a (partial) probabilistic transition function;
- $\mathcal{L} : L \rightarrow 2^{AP}$  is a labelling function mapping each location to a set of atomic propositions.

A state of a PTA is a pair  $(l, v) \in L \times \mathbb{R}_{\geq 0}^{\mathcal{X}}$  such that  $v \models inv(l)$ . In any state  $(l, v)$ , either a certain amount of time  $t \in \mathbb{R}_{\geq 0}$  elapses, or an action  $a \in Act$  is performed. If time elapses, then the choice of  $t$  requires that the invariant  $inv(l)$  remains continuously satisfied while time passes. The resulting state after this transition is  $(l, v + t)$  and, to ease notation, we denote this state by  $(l, v) + t$ . In the case where an action is performed, an action  $a$  can only be chosen if it is enabled, i.e., if the clock constraint  $enab(l, a)$  is satisfied by  $v$ . Once an enabled action  $a$  is chosen, a set of clocks to reset and a successor location are selected at random, according to the distribution  $prob(l, a)$ . We call each element  $(X, l') \in 2^{\mathcal{X}} \times L$  in the support of  $prob(l, a)$  an edge and use  $edges(l, a)$  to denote the set of such edges.

We assume that PTAs are well-formed, meaning that, for each state  $(l, v)$  and action  $a$  such that  $v$  satisfies  $enab(l, a)$ , every edge  $(X, l') \in edges(l, a)$  results in a transition to a valid state, i.e., we have  $v[X:=0] \models inv(l')$ . A PTA can be transformed into one that is well-formed by incorporating the invariant associated with the target location into the enabling condition of each location-action pair (see [56]).

**Definition 4** (PTA semantics) Let  $P = (L, \bar{l}, \mathcal{X}, Act, inv, enab, prob, \mathcal{L})$  be a PTA. The semantics of  $P$  is defined as the (infinite-state) TPS  $\llbracket P \rrbracket = (S, \bar{s}, Act, Steps_P, lab)$  where:

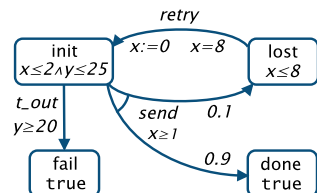
- $S = \{(l, v) \in L \times \mathbb{R}_{\geq 0}^{\mathcal{X}} \mid v \models inv(l)\}$  and  $\bar{s} = (\bar{l}, \mathbf{0})$ ;
- for any  $(l, v) \in S$  and  $a \in Act \cup \mathbb{R}_{\geq 0}$ , we have  $Steps_P((l, v), a) = \lambda$  if and only if either:
  - time transitions:  $a \in \mathbb{R}_{\geq 0}$ ,  $v + t' \models inv(l)$  for all  $0 \leq t' \leq a$ , and  $\lambda = \mu_{(l, v+a)}$ ,
  - action transitions:  $a \in Act$ ,  $v \models enab(l, a)$  and for each  $(l', v') \in S$ :

$$\lambda(l', v') = \sum \{ |prob(l, a)(X, l') \mid X \in 2^{\mathcal{X}} \wedge v' = v[X:=0] \};$$

- for any  $(l, v) \in S$  we have  $lab(l, v) = \mathcal{L}(l)$ .

*Example of a PTA* In Fig. 1, we present a PTA modelling a simple communication protocol. We adopt the standard conventions for the graphical representation of timed automata. Distributions are represented by an arc connecting edges at their source and by probability labels attached to edges (omitted for edges taken with probability 1). The PTA has two clocks  $x$  and  $y$ , which start with the value 0. In the location *init*, the system waits for at least 1 time unit (represented by the enabling condition  $x \geq 1$  on the outgoing distribution of action *send*) and at most 2 time units (represented by conjunct  $x \leq 2$  of the invariant

**Fig. 1** Example of a PTA



condition), before sending a message. With probability 0.9 the message is received correctly (edge to done); otherwise, with probability 0.1, the message is lost (edge to lost). In the latter case, once clock  $x$  reaches 8, the PTA returns to init where another attempt to send the message can be made. If, in total, at least 20 and at most 25 time units have elapsed since the start of system execution, the PTA performs a timeout and moves to location fail.

**Rewards for PTAs** We can also define *rewards* for a PTA. Often, these model usage of some resource and are equivalently referred to as *costs* or *prices*. A *reward structure* at the level of a PTA  $P$  is defined using a pair  $\mathbf{r} = (\mathbf{r}_L, \mathbf{r}_{Act})$ , where  $\mathbf{r}_L : L \rightarrow \mathbb{R}_{\geq 0}$  is a function assigning to each location the rate at which rewards are accumulated as time passes in that location and  $\mathbf{r}_{Act} : L \times Act \rightarrow \mathbb{R}_{\geq 0}$  is a function assigning the reward of executing each action in each location. The corresponding reward structure of the TPS  $\llbracket P \rrbracket$  is given by  $r = (r_S, r_{Act})$  where  $r_S(l, v) = \mathbf{r}_L(l)$  and  $r_{Act}((l, v), a) = \mathbf{r}_{Act}(l, a)$  for  $(l, v) \in L \times \mathbb{R}_{\geq 0}^X$  and  $a \in Act$ . PTAs equipped with reward structures are a probabilistic extension of linearly-priced timed automata (also known as weighted timed automata) [5, 14].

**Modelling with PTAs** We now summarise a variety of extensions to the standard definition of PTAs that facilitate high-level modelling using this formalism.

**Parallel composition.** It is often useful to define complex systems as the *parallel composition* of several interacting components. The definition of the parallel composition operator  $\parallel$  for PTAs [55] uses ideas from (untimed) probabilistic automata [62] and classical timed automata [4]. Let  $P_i = (L_i, \bar{l}_i, Act_i, \mathcal{X}_i, inv_i, enab_i, prob_i, \mathcal{L}_i)$  for  $i \in \{1, 2\}$  and assume that  $\mathcal{X}_1 \cap \mathcal{X}_2 = \emptyset$ . Given  $\mu_1 \in Dist(2^{\mathcal{X}_1} \times L_1)$  and  $\mu_2 \in Dist(2^{\mathcal{X}_2} \times L_2)$ , we let  $\mu_1 \otimes \mu_2 \in Dist(2^{\mathcal{X}_1 \cup \mathcal{X}_2} \times (L_1 \times L_2))$  be such that  $\mu_1 \otimes \mu_2(X_1 \cup X_2, (l_1, l_2)) = \mu_1(X_1, l_1) \cdot \mu_2(X_2, l_2)$  for  $X_i \subseteq \mathcal{X}_i, l_i \in L_i$  and  $i \in \{1, 2\}$ . The *parallel composition* of PTAs  $P_1$  and  $P_2$  is the PTA:

$$P_1 \parallel P_2 = (L_1 \times L_2, (\bar{l}_1, \bar{l}_2), \mathcal{X}_1 \cup \mathcal{X}_2, Act_1 \cup Act_2, inv, enab, prob, \mathcal{L})$$

such that, for each location pair  $(l_1, l_2) \in L_1 \times L_2$  and action  $a \in Act_1 \cup Act_2$ :

- the invariant condition is given by  $inv(l_1, l_2) = inv_1(l_1) \wedge inv_2(l_2)$ ;
- the enabling condition is given by:

$$enab((l_1, l_2), a) = \begin{cases} enab_1(l_1, a) \wedge enab_2(l_2, a) & \text{if } a \in Act_1 \cap Act_2, \\ enab_1(l_1, a) & \text{if } a \in Act_1 \setminus Act_2, \\ enab_2(l_2, a) & \text{if } a \in Act_2 \setminus Act_1; \end{cases}$$

- the probabilistic transition function is given by:

$$prob((l_1, l_2), a) = \begin{cases} prob_1(l_1, a) \otimes prob_2(l_2, a) & \text{if } a \in Act_1 \cap Act_2, \\ prob_1(l_1, a) \otimes \mu_{(\emptyset, l_2)} & \text{if } a \in Act_1 \setminus Act_2, \\ \mu_{(\emptyset, l_1)} \otimes prob_2(l_2, a) & \text{if } a \in Act_2 \setminus Act_1; \end{cases}$$

- the labelling function is given by  $\mathcal{L}(l_1, l_2) = \mathcal{L}_1(l_1) \cup \mathcal{L}_2(l_2)$ .

If PTA  $P_i$  has associated reward structure  $(\mathbf{r}_L^i, \mathbf{r}_{Act}^i)$ , then the reward structure  $\mathbf{r} = (\mathbf{r}_L, \mathbf{r}_{Act})$  for  $P_1 \parallel P_2$  is such that, for  $(l_1, l_2) \in L_1 \times L_2$  and  $a \in Act_1 \cup Act_2$ , we have  $\mathbf{r}_L(l_1, l_2) = \mathbf{r}_L^1(l_1) + \mathbf{r}_L^2(l_2)$ ,  $\mathbf{r}_{Act}((l_1, l_2), a) = \mathbf{r}_{Act}^1(l_1, a) + \mathbf{r}_{Act}^2(l_2, a)$  if  $a \in Act_1 \cap Act_2$ ,  $\mathbf{r}_{Act}((l_1, l_2), a) = \mathbf{r}_{Act}^1(l_1, a)$  if  $a \in Act_1 \setminus Act_2$  and  $\mathbf{r}_{Act}((l_1, l_2), a) = \mathbf{r}_{Act}^2(l_2, a)$  if  $a \in Act_2 \setminus Act_1$ .

*Discrete variables.* When modelling systems with (probabilistic) timed automata, it is often convenient to augment the model with discrete variables [13, 65]. We restrict ourselves to the case in which a finite number of variables, each with a finite domain, are added to the PTA framework; enabling conditions can then refer to the current value of the variables, and the probabilistic transition relation is extended to allow updating variable values. Such an extended PTA model can be represented in the standard PTA framework presented above in the following way. Let  $L$  be the set of locations of the extended PTA, and suppose the extended PTA has  $n$  bounded integer variables. The locations of the standard PTA are tuples comprising  $n + 1$  elements: the first element is a location from  $L$ , while the remaining  $n$  elements are values of the bounded integer variables. Enabling conditions are obtained by resolving partially the enabling conditions of the extended PTA, using the variable values corresponding to the locations and the probabilistic transition function is obtained by encoding variable updates into target locations of edges.

*Urgency.* When modelling real-time systems, it is often necessary to express the fact that a particular action should be taken immediately, without letting time pass. In this way, we can model, for example, an instantaneous system event comprising several atomic actions. A number of mechanisms for modelling such situations have been introduced for timed automata, for example in the system-description language of the UPPAAL model checker [13]; here, we describe how they are adapted to PTAs.

Firstly, an *urgent location* of a PTA is a location in which no time can pass. Urgent locations can be represented in the PTA framework by introducing an additional clock, which is reset on entry to an urgent location, and by including a conjunct in the invariant condition of the location to specify that the value of the clock should be equal to 0 in the location.

Secondly, a *committed location* of a PTA is, like an urgent location, a location in which no time can pass, but also must be left before any other system component makes a transition. We adapt to PTAs the method of [65] for encoding committed locations in the standard timed automata framework. First a global Boolean variable *atom* is added to the PTA. Now, consider a PTA which is to be composed in parallel with other PTAs. A committed location of the PTA is subject to the constructs added in the case of urgent locations and, in addition, *atom* is set to `true` on entry to the committed location, is set to `false` on exiting the committed location, and all enabling conditions of the PTA except those corresponding to committed locations have the conjunct requiring that *atom* is `false`.

Finally we mention *urgent actions* [13, 30]. Informally, an action is urgent if it must be chosen as soon as it is enabled. Urgent actions can be introduced to the PTA syntax simply by identifying the subset  $Act_u$  of actions which are interpreted as urgent. The presence of urgent actions necessitates the following modifications to the semantics of a PTA. For PTA  $P = (L, \bar{l}, Act, \mathcal{X}, inv, enab, prob, \mathcal{L})$  with urgent actions  $Act_u$ ,  $\llbracket P \rrbracket$  is the TPS  $(S, \bar{s}, Act, Steps_P, lab)$  where  $S, \bar{s}, lab$  and  $Steps_P((l, v), a)$  for  $(l, v) \in S$  and  $a \in Act$  are as in Definition 4, while for  $(l, v) \in L, t \in \mathbb{R}_{\geq 0}$ , we have  $Steps_P((l, v), t) = \mu_{(l, v+t)}$  if and only if:

- $v + t' \models inv(l)$  for all  $0 \leq t' \leq t$ ;
- for all  $0 \leq t' < t$  and  $a' \in Act$ , if  $v + t' \models enab(l, a')$  then  $a' \notin Act_u$ .

*Non-standard clock resets.* In the PTA framework, as in the standard definition of timed automata, clocks can only be reset to value 0 when taking a probabilistic transition. It may be useful, though, to also allow clocks to be reset to any non-negative integer value. If  $\mathbb{N}_\perp = \mathbb{N} \cup \{\perp\}$ ,  $v$  is a clock valuation and  $\theta : \mathcal{X} \rightarrow \mathbb{N}_\perp$ , then define the clock valuation  $v[\theta]$  such that  $v[\theta](x) = v(x)$  if  $\theta(x) = \perp$  and  $v[\theta](x) = \theta(x)$  otherwise. Then, a PTA with



*extended resets* is defined as a PTA in Definition 3, except that the probabilistic transition function is now defined as  $prob : L \times Act \rightarrow Dist((\mathcal{X} \rightarrow \mathbb{N}_\perp) \times L)$ . The semantics of a PTA with extended resets  $\mathbf{P}$  is defined as in Definition 4, except that for  $(l, v) \in S$  and  $a \in Act$ , we have  $Steps_{\mathbf{P}}((l, v), a) = \lambda$  if and only if  $v \models enab(l, a)$  and, for each  $(l', v') \in S$ :

$$\lambda(l', v') = \sum \{ |prob(l, a)(\theta, l')| \mid \theta \in \mathcal{X} \rightarrow \mathbb{N}_\perp \wedge v' = v[\theta] \}.$$

A PTA with extended resets can be translated into a PTA with the standard restriction of clock resets to 0, by adapting an analogous construction for (non-probabilistic) timed automata [26]. However, this construction can give an exponential blow-up in the size of the model, which is unavoidable [25]. Instead, model-checking algorithms for PTAs can be developed which incorporate the extended definition of resets directly.

*Channels.* The definitions of actions and parallel composition presented here can be extended to allow for channels and the sending and receiving (to either single or multiple recipients) of messages along them, as in UPPAAL [13]. Such behaviour can be encoded in the action names of a standard PTA.

*Time divergence* An important issue with regard to the verification of models of real-time systems is that of time divergence. As explained in Sect. 2, such behaviour does not correspond to that of the real system, and hence the verification technique used must be able to disregard such behaviour during analysis. We use the notion of time divergence of Sect. 2, i.e., we restrict our attention to the adversaries  $Adv_{\llbracket \mathbf{P} \rrbracket}$  (those adversaries for which the probability of time passing beyond any bound is 1).

Note that a PTA may feature states from which time cannot diverge for any adversary; such states correspond to a probabilistic generalisation of *timelocks* in the timed automata setting [40], and are considered to indicate modelling errors. The set of timelock states can be identified using (extensions of) the analysis methods that we present in Sect. 5, and removed from the model by modifying the invariant and enabling conditions.

For some PTAs, all adversaries will be time-divergent by construction. We give a syntactic and compositional condition, derived from analogous results on timed automata [66, 67], which guarantees that all adversaries are time-divergent: a PTA is *structurally divergent* if, for every sequence  $(l_0, a_0), (X_0, l_1), \dots, (l_n, a_n), (X_n, l_{n+1})$  such that  $(X_i, l_{i+1}) \in edges(l_i, a_i)$  for  $0 \leq i < n$  and  $l_{n+1} = l_0$ , there exists  $x \in \mathcal{X}$  and  $0 \leq i, j \leq n$  such that  $x \in X_i$  and  $enab(l_j, a_j) \Rightarrow (x \geq 1)$  (i.e.  $enab(l_j, a_j)$  contains a conjunct of the form  $x \geq c$  for  $c \geq 1$ ). If a PTA is not structurally divergent, verification algorithms can be adjusted to disregard non-divergent adversaries, as we will describe in Sect. 5.

Finally we note that a more restrictive notion of divergent adversary of a PTA, namely that of a *strictly divergent* adversary, has been presented in [63]. An adversary is strictly divergent if *all* of its paths entail time passing beyond any bound. In many contexts, strictly divergent adversaries are more realistic than divergent adversaries, e.g., if we regard the edge traversals of a PTA as corresponding to the change in a physical state of the system.

*Alternative models* We conclude this section with a brief discussion of some alternative probabilistic models that also incorporate nondeterministic, probabilistic and timed aspects. One example, which has recently attracted increased interest in the context of probabilistic verification, is *continuous-time Markov decision processes*, along with closely related models such as *interactive Markov chains* [41] and *Markov automata* [31]. These extend classical (discrete-time) Markov decision processes with real-time delays modelled by exponential



distributions. Thus, such models can alternatively be viewed as extensions of *continuous-time Markov chains*, which also model probabilistic real-time systems using exponentially-distributed delays, but which do not exhibit nondeterministic behaviour. Other proposed models in the literature include [1], which is based on generalised semi-Markov processes, and [53], which defines PTAs with continuously distributed random delays.

### 4 Property specification for PTAs

In this section, we present a temporal logic for the formal specification of quantitative properties of PTAs. The basis for this is the temporal logic PCTL [20, 37], a probabilistic extension of the logic CTL [29] which has been proposed for specifying properties of both MDPs [20] and discrete-time Markov chains [37]. We augment the logic with operators to reason about rewards (or costs or prices), in the style of the logic in [32] for MDPs (and as used in the probabilistic model checker PRISM [51]). We also discuss extensions to more expressive logics such as PTCTL and PCTL\*. In the next section, we will survey the available techniques for model checking of PTAs against properties specified in our logic.

**Definition 5** (Syntax) The syntax of our logic is given by the following grammar:

$$\begin{aligned} \phi &::= \text{true} \mid \mathbf{a} \mid \chi \mid \phi \wedge \phi \mid \neg\phi \mid P_{\triangleright p}[\psi] \mid R_{\triangleright q}^{\mathbf{r}}[\rho], \\ \psi &::= \phi \cup^{\leq k} \phi \mid \phi \cup \phi, \\ \rho &::= I^k \mid C^{\leq k} \mid F \phi \end{aligned}$$

where  $\mathbf{a} \in AP$  is an atomic proposition,  $\chi \in CC(\mathcal{X})$  is a clock constraint,  $\triangleright \in \{\leq, <, \geq, >\}$ ,  $p \in \mathbb{Q} \cap [0, 1]$ ,  $q \in \mathbb{Q}_{\geq 0}$ ,  $\mathbf{r}$  is a reward structure and  $k \in \mathbb{N}$ .

This logic extends propositional logic with a *probabilistic operator* ( $P$ ) and a *reward operator* ( $R$ ). Informally, a property of the form  $P_{\triangleright p}[\psi]$  states that the probability of *path formula*  $\psi$  being true always satisfies the bound  $\triangleright p$ . A property of the form  $R_{\triangleright q}^{\mathbf{r}}[\rho]$  means that the *expected* value of *reward function*  $\rho$  on *reward structure*  $\mathbf{r}$  meets the bound  $\triangleright q$ .

Formulae in the logic are always *state formulae*, i.e., those formed by the production  $\phi$  in the grammar above. These are evaluated over the states of a PTA  $P$  (or, more precisely, over the states of the TPS  $\llbracket P \rrbracket$  representing its semantics). For state  $s$  and formula  $\phi$ , we write  $s \models \phi$  to denote that  $\phi$  is satisfied in  $s$ . The syntax also includes *path formulae* ( $\psi$ ) and *reward operators* ( $\rho$ ), which appear only as subformulae of the  $P$  and  $R$  operators.

We include two types of path formulae: *time-bounded until* ( $\phi_1 \cup^{\leq k} \phi_2$ ) and (unbounded) *until* ( $\phi_1 \cup \phi_2$ ). Formula  $\phi_1 \cup \phi_2$  means that a state satisfying  $\phi_2$  is eventually reached and that, at every time-instant prior to that,  $\phi_1$  is satisfied. The time-bounded variant has the same meaning, with the additional constraint that the occurrence of  $\phi_2$  must occur within time  $k$ . We can derive several useful operators, such as  $F \phi \equiv \text{true} \cup \phi$ , which means that  $\phi$  is eventually satisfied, and  $F^{\leq k} \phi \equiv \text{true} \cup^{\leq k} \phi$ , which means that  $\phi$  is satisfied within time  $k$ . We also have  $G \phi \equiv \neg(F \neg\phi)$ , which means that  $\phi$  is always satisfied, and  $G^{\leq k} \phi \equiv \neg(F^{\leq k} \neg\phi)$  which means that  $\phi$  is continuously satisfied for time  $k$ . Although the  $G$  and  $G^{\leq k}$  operators cannot be derived from the basic syntax of the logic since there is no negation of path formulae, it can be shown that  $P_{\leq p}[\neg\psi] \equiv P_{\geq 1-p}[\psi]$ ,  $P_{< p}[\neg\psi] \equiv P_{> 1-p}[\psi]$ ,  $P_{\geq p}[\neg\psi] \equiv P_{\leq 1-p}[\psi]$  and  $P_{> p}[\neg\psi] \equiv P_{< 1-p}[\psi]$ . PCTL-style logics often include a *next* ( $X$ ) operator but this is of less use in timed models and omitted.

The reward operator  $\mathbb{I}^k$  refers to the reward of the current state at time instant  $k$ ,  $\mathbb{C}^{\leq k}$  to the total reward accumulated up until time point  $k$ , and  $\mathbb{F} \phi$  to the total reward accumulated until a state satisfying  $\phi$  is reached. Formally, we define the semantics of the logic as follows.

**Definition 6** (Semantics) Let  $\mathbb{P}$  be a PTA,  $\llbracket \mathbb{P} \rrbracket = (S, \bar{s}, Act, \mathbb{T}, Steps_{\mathbb{P}}, lab)$  be its semantics, and let  $r$  denote the reward structure over  $\llbracket \mathbb{P} \rrbracket$  corresponding to a reward structure  $\mathbf{r}$  over  $\mathbb{P}$ . For state  $s = (l, v) \in S$ , the satisfaction relation  $\models$  is defined inductively by:

$$\begin{aligned}
 s \models \text{true} & \quad \text{always} \\
 s \models a & \quad \iff a \in lab(s) \\
 s \models \chi & \quad \iff v \models \chi \\
 s \models \phi_1 \wedge \phi_2 & \quad \iff s \models \phi_1 \wedge s \models \phi_2 \\
 s \models \neg \phi & \quad \iff s \not\models \phi \\
 s \models \mathbb{P}_{\triangleright p}[\psi] & \quad \iff Pr_{\llbracket \mathbb{P} \rrbracket, s}^{\sigma}(\{\omega \in Path_{\llbracket \mathbb{P} \rrbracket, s}^{\sigma} \mid \omega \models \psi\}) \triangleright p \text{ for all } \sigma \in Adv_{\llbracket \mathbb{P} \rrbracket} \\
 s \models \mathbb{R}_{\triangleright q}^r[\rho] & \quad \iff \mathbb{E}_{\llbracket \mathbb{P} \rrbracket, s}^{\sigma}(rew(r, \rho)) \triangleright q \text{ for all } \sigma \in Adv_{\llbracket \mathbb{P} \rrbracket}
 \end{aligned}$$

where, for any infinite path  $\omega$  of  $\llbracket \mathbb{P} \rrbracket$ :

$$\begin{aligned}
 \omega \models \phi_1 \cup^{\leq k} \phi_2 & \quad \iff \text{there exists a position } (i, t) \text{ of } \omega \text{ such that } \omega(i) + t \models \phi_2 \\
 & \quad \text{and } dur_{\omega}(i) + t \leq k \text{ and } \omega(j) + t' \models \phi_1 \vee \phi_2 \text{ for all positions} \\
 & \quad (j, t') < (i, t) \text{ of } \omega, \\
 \omega \models \phi_1 \cup \phi_2 & \quad \iff \text{there exists a position } (i, t) \text{ of } \omega \text{ such that } \omega(i) + t \models \phi_2 \\
 & \quad \text{and } \omega(j) + t' \models \phi_1 \vee \phi_2 \text{ for all positions } (j, t') < (i, t) \text{ of } \omega
 \end{aligned}$$

and, for reward structure  $r = (r_S, r_{Act})$  over  $\llbracket \mathbb{P} \rrbracket$ , the random variable  $rew(r, \rho)$  over infinite paths of  $\llbracket \mathbb{P} \rrbracket$  is defined as follows:

$$\begin{aligned}
 rew(r, \mathbb{I}^k)(\omega) & = r_S(\omega(j_k)), \\
 rew(r, \mathbb{C}^{\leq k})(\omega) & = \sum_{i=0}^{j_k-1} r(\omega, i) + (k - dur_{\omega}(j_k)) \cdot r_S(\omega(j_k)), \\
 rew(r, \mathbb{F} \phi)(\omega) & = \begin{cases} \sum_{i=0}^{j_{\phi}-1} r(\omega, i) + t_{\phi} \cdot r_S(\omega(j_{\phi})) & \text{if } (j_{\phi}, t_{\phi}) \text{ exists,} \\ \infty & \text{otherwise} \end{cases}
 \end{aligned}$$

where  $j_0 = 0$ ,  $j_k = \max\{i \mid dur_{\omega}(i) < k\}$  for  $k > 0$  and, when it exists,  $(j_{\phi}, t_{\phi})$  is the minimum position under the ordering  $<$  such that  $\omega(j_{\phi}) + t_{\phi} \models \phi$ .

In addition to the basic syntax of Definition 6, we allow, in the style of the PRISM model checker, *quantitative* (numerical) queries yielding the minimum or maximum probability or expected reward value from a state  $s$ . We use  $\mathbb{P}_{\min=?}[\psi]$ ,  $\mathbb{P}_{\max=?}[\psi]$ ,  $\mathbb{R}_{\min=?}^r[\rho]$  and  $\mathbb{R}_{\max=?}^r[\rho]$ , which give, respectively:

$$\begin{aligned}
 Pr_{\mathbb{P}, s}^{\min}(\psi) & \stackrel{\text{def}}{=} \inf_{\sigma \in Adv_{\llbracket \mathbb{P} \rrbracket}} Pr_{\llbracket \mathbb{P} \rrbracket, s}^{\sigma}(\{\omega \in Path_{\llbracket \mathbb{P} \rrbracket, s}^{\sigma} \mid \omega \models \psi\}), \\
 Pr_{\mathbb{P}, s}^{\max}(\psi) & \stackrel{\text{def}}{=} \sup_{\sigma \in Adv_{\llbracket \mathbb{P} \rrbracket}} Pr_{\llbracket \mathbb{P} \rrbracket, s}^{\sigma}(\{\omega \in Path_{\llbracket \mathbb{P} \rrbracket, s}^{\sigma} \mid \omega \models \psi\}),
 \end{aligned}$$

$$\mathbb{E}_{P,s}^{r,\min}(\rho) \stackrel{\text{def}}{=} \inf_{\sigma \in \text{Adv}_{\|P\|}} \mathbb{E}_{\|P\|,s}^{\sigma}(rew(r, \rho)),$$

$$\mathbb{E}_{P,s}^{r,\max}(\rho) \stackrel{\text{def}}{=} \sup_{\sigma \in \text{Adv}_{\|P\|}} \mathbb{E}_{\|P\|,s}^{\sigma}(rew(r, \rho)).$$

Some typical examples of PTA properties, specified in this logic are:

- $P_{\geq 0.8}[\mathbb{F}^{\leq k} \text{ack}_n]$ —“the probability that the sender has received  $n$  acknowledgements within  $k$  clock-ticks is at least 0.8”;
- $\text{trigger} \rightarrow P_{< 0.0001}[\mathbb{G}^{\leq 20} \neg \text{deploy}]$ —“the probability of the airbag failing to deploy within 20 milliseconds of being triggered is strictly less than 0.0001”;
- $P_{\max=?}[\neg \text{sent} \cup \text{fail}]$ —“what is the maximum probability of a failure occurring before message transmission is complete?”;
- $R_{\max=?}^{\text{time}}[\mathbb{F} \text{end}]$ —“what is the maximum expected time for the protocol to terminate?”;
- $R_{< q}^{\text{PWT}}[\mathbb{C}^{\leq 60}]$  – “the expected energy consumption during the first 60 seconds is  $< q$ ”.

*Property reductions* We now describe how model checking for several of the operators included in our logic can be reduced to checking satisfaction of a simpler formula on a modified PTA. Consider first a time-bounded until property  $P_{\triangleright p}[\phi_1 \cup^{\leq k} \phi_2]$  on a PTA  $P$ . If we augment  $P$  with an additional clock  $z$ , then it follows that a state  $(l, v)$  of  $P$  satisfies the formula  $P_{\triangleright p}[\phi_1 \cup^{\leq k} \phi_2]$  if and only if the state  $(l, v')$  of the augmented PTA where  $v'(x) = v(x)$  for all clocks  $x$  of  $P$  and  $v'(z) = 0$  satisfies  $P_{\triangleright p}[\phi_1 \cup (\phi_2 \wedge (z \leq k))]$  (see [56]). Second, given an until property  $P_{\triangleright p}[\phi_1 \cup \phi_2]$ , if we modify the PTA such that, upon reaching a state not satisfying  $\phi_1$ , only a transition to a sink state is possible, then a state of the PTA satisfies  $P_{\triangleright p}[\phi_1 \cup \phi_2]$  if and only if a state of the modified PTA satisfies  $P_{\triangleright p}[\mathbb{F} \phi_2]$ .

Next, we show how to reduce checking a property of the form  $P_{\geq p}[\mathbb{F}^{\leq k} \phi]$ , to a formula of the form  $P_{\leq 1-p}[\mathbb{F} \text{a}_{\text{exc}}]$  (the former requires computation of minimum probabilities, whereas the latter needs maximum probabilities, which are sometimes easier to compute). The reduction is correct only for states satisfying  $P_{\geq 1}[\mathbb{F} \phi]$ , i.e.  $\phi$  is always reachable with probability 1. We augment  $P$  with an extra clock  $z$  and then modify it such that states satisfying  $\phi$  are forced to make a transition to a sink-location and, in all other states, we add a transition to a different, sink-location *exceeded*, enabled when  $z > k$ . Then  $P_{\geq p}[\mathbb{F}^{\leq k} \phi]$  is true in a state of the original model if and only if the corresponding state of the modified PTA with  $z = 0$  satisfies  $P_{\leq 1-p}[\mathbb{F} \text{a}_{\text{exc}}]$ , where  $\text{a}_{\text{exc}}$  is true only in location *exceeded* (see [55]).

Finally, we show how to reduce properties of the form  $R_{\triangleright q}^r[\mathbb{C}^{\leq k}]$  or  $R_{\triangleright q}^r[\mathbb{I}^{\leq k}]$  to  $R_{\triangleright q}^r[\mathbb{F} \phi]$ . In both cases, we add an extra clock  $z$  to  $P$ . For  $R_{\triangleright q}^r[\mathbb{C}^{\leq k}]$ , it suffices to check  $R_{\triangleright q}^r[\mathbb{F} (z = k)]$  on the augmented PTA. For  $R_{\triangleright q}^r[\mathbb{I}^{\leq k}]$ , we add the conjunct  $z \leq k$  to all invariants and a transition to a new sink location *done* (labelled  $\text{a}_{\text{done}}$ ) with enabling condition  $z = k$  to all locations, while changing the enabling conditions of all other transitions so that they are not enabled when  $z = k$ . It then suffices to check  $R_{\triangleright q}^r[\mathbb{F} \text{a}_{\text{done}}]$ , where the only non-zero rewards are action rewards on the new transitions, set to the location reward of the source location.

*More expressive logics* The next section of this paper will discuss techniques for model checking the properties expressible in the logic given above. A variety of more expressive logics have also been considered for PTAs. PTCTL [54] is a probabilistic extension of the timed temporal logic TCTL. In particular, it includes a *freeze quantifier* (or *reset quantifier*)  $z.\phi$ , which introduces a formula clock  $z$ , reset to zero, that can be referred to in the subformula  $\phi$ . Although not considered further in this paper, PTCTL can be model checked

using the region graph construction and backwards reachability method discussed in Sect. 5. Formulae of the logics LTL and PCTL\*, originally proposed for discrete-time probabilistic systems, can be verified on PTAs using a Rabin automaton product construction and the model checking algorithm of [64].

## 5 Model checking for PTAs

We now consider the problem of model checking a PTA  $\mathcal{P}$  with respect to a property  $\phi$  of the logic presented in Sect. 4, i.e., determining  $Sat(\phi) \stackrel{\text{def}}{=} \{s \in S \mid s \models \phi\}$ , where  $S$  is the set of states of  $\llbracket \mathcal{P} \rrbracket$ . We will survey the various PTA model checking techniques that have been proposed in the literature, which support different fragments of the logic. We will cover:

- the region graph construction [54];
- the boundary region graph [43];
- the digital clocks method [52];
- backwards reachability [56];
- abstraction refinement with stochastic games [50].

The first two approaches, which are based on the concept of the *region graph* [2, 4], are used primarily to establish the decidability and complexity of model checking, rather than for practical implementations. The others provide efficient methods for model checking particular fragments of the logic. Table 1 provides a summary of the methods and their applicability. We omit the *forwards reachability* algorithm described in [54] since it only computes bounds on probabilities of system behaviour, rather than the exact values.

Unless otherwise stated, we also assume that  $\mathcal{P}$  has no timelocks and is structurally divergent (structurally non-Zeno). In Sect. 5.1, we *do* explain how to treat PTAs that are not structurally divergent. Similar adaptations can also be applied to the other methods that we discuss in this section.

### 5.1 The region graph

We first discuss the *region graph* construction for a PTA [54], which is based on the classic construction for timed automata [2, 4]. This approach provides a way to model check the fragment of the logic from Sect. 4 that excludes  $\mathbb{R}_{>q}^r[\rho]$  formulae (in the next section, we will relax this restriction). The region graph of a PTA  $\mathcal{P}$  and formula  $\phi$  takes the form of a finite-state MDP whose states are *regions* of the form  $(l, \alpha)$ , where  $l$  is a location and  $\alpha$  is an equivalence class of clock valuations according to the equivalence defined below. Let  $c$  be the maximal constant to which any clock is compared in the clock constraints of  $\mathcal{P}$  and  $\phi$ . Then clock valuations  $v$  and  $v'$  are equivalent if and only if they satisfy the following conditions:

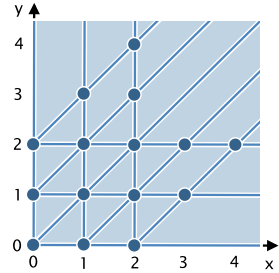
- for any  $x \in \mathcal{X}$ , either  $v(x) > c$  and  $v'(x) > c$ , or  $v(x)$  and  $v'(x)$  agree on their integer parts;
- for any  $x, x' \in \mathcal{X}$ , either  $v(x) - v(x') > c$  and  $v'(x) - v'(x') > c$ , or  $v(x) - v(x')$  and  $v'(x) - v'(x')$  agree on their integer parts,

where two values  $q, q' \in \mathbb{R}_{\geq 0}$  agree on their integer parts when  $\lfloor q \rfloor = \lfloor q' \rfloor$  and  $\lfloor q \rfloor - q = 0$  if and only if  $\lfloor q' \rfloor - q' = 0$ . Figure 2 shows the set of possible equivalence classes of clock valuations for clocks  $x, y$  when  $c = 2$ . Note that equivalent clock valuations satisfy the same clock constraints of the PTA. The set of regions needed for the region graph, denoted  $R$ , is

**Table 1** Summary of PTA model checking techniques and their applicability

|  | Region graph | Boundary region graph | Digital clocks (for closed PTAs) | Backwards reachability | Stochastic games |
|--|--------------|-----------------------|----------------------------------|------------------------|------------------|
| Single $P_{\triangleright q}[\psi]$ operator   | ✓            | ✓                     | ✓                                | ✓                      | ✓                |
| Single $R_{\triangleright q}^r[\rho]$ operator | ×            | ✓                     | ✓                                | Open                   | Open             |
| Logic without $R_{\triangleright q}^r[\rho]$   | ✓            | ✓                     | ×                                | ✓                      | Open             |
| Full logic                                     | ×            | ×                     | ×                                | Open                   | Open             |

**Fig. 2** Clock equivalence classes for two clocks  $x$  and  $y$  ( $c = 2$ )



the set of regions  $(l, \alpha)$  such that there exists  $v \in \alpha$  such that  $v \models \text{inv}(l)$ . The size of  $R$  is bounded by  $|L| \cdot (2c + 2)^{(|X|+1)^2}$  (see [23]).

A region  $(l, \alpha) \in R$  may have a *time-successor*, defined as follows. If  $v + t \in \alpha$  for all  $v \in \alpha$  and  $t \in \mathbb{R}_{\geq 0}$ , then the time-successor of  $(l, \alpha)$  is  $(l, \alpha)$  itself. Otherwise, there exists the unique region  $(l, \beta) \neq (l, \alpha)$  for which there are  $v \in \alpha$  and  $t \in \mathbb{R}_{\geq 0}$  such that  $v + t \in \beta$  and  $v + t' \in \alpha \cup \beta$  for all  $0 \leq t' \leq t$ . If additionally we have  $v + t' \models \text{inv}(l)$  for all  $0 \leq t' \leq t$ , then  $(l, \beta)$  is the time-successor of  $(l, \alpha)$ , otherwise  $(l, \alpha)$  has no time-successor.

The region graph for  $P$  and  $\phi$  is the finite-state MDP  $\text{Reg}[P, \phi] = (R, (\vec{l}, [0]), \text{Act} \cup \{\tau\}, \text{Steps}, \text{lab})$ , where for each  $(l, \alpha) \in R$  and  $a \in \text{Act} \cup \{\tau\}$ , we have  $\text{Steps}((l, \alpha), a) = \lambda$  if and only if either:

- *time transitions*:  $a = \tau$ ,  $\lambda = \mu_{(l, \beta)}$  and  $(l, \beta)$  is the time-successor of  $(l, \alpha)$ ;
- *action transitions*:  $a \in \text{Act}$ , there is a  $v \in \alpha$  with  $v \models \text{enab}(l, a)$  and, for each  $(l', \beta) \in R$

$$\lambda(l', \beta) = \sum \{ | \text{prob}(l, a)(X, l') \mid X \in 2^X \wedge \beta = \alpha[X := 0] \};$$

and  $\text{lab}(l, \alpha) = \mathcal{L}(l)$  for all  $(l, \alpha) \in R$ .

We now consider how  $\text{Reg}[P, \phi]$  can be used to verify  $P$  against  $\phi$  (recall that we consider the fragment of the logic without the  $R_{\triangleright q}^r[\cdot]$  operator). For simplicity, we first assume that  $P$  is structurally divergent. Model checking proceeds in standard fashion (for a branching-time logic), recursing over subformulae  $\phi'$  of  $\phi$  and computing  $\text{Sat}(\phi')$ . Identifying states that satisfy atomic propositions, clock constraints or Boolean connectives is straightforward and time-bounded properties are dealt with using the reduction given in Sect. 4. Hence, we focus on formulae of the form  $P_{\triangleright q}[\phi_1 \cup \phi_2]$ . If the state sets satisfying  $\phi_1$  and  $\phi_2$  have already been computed and the regions of  $\text{Reg}[P, \phi]$  corresponding to these sets are labelled  $a_1$  and  $a_2$ , respectively, then the states in  $\llbracket P \rrbracket$  satisfying  $P_{\triangleright q}[\phi_1 \cup \phi_2]$  are simply those in the regions corresponding to states in the MDP  $\text{Reg}[P, \phi]$  that satisfy  $P_{\triangleright q}[a_1 \cup a_2]$ .

Thus, the region graph yields an algorithm for model checking structurally divergent PTAs against properties of the logic without the  $R_{\triangleright q}^r[\cdot]$  operator. The algorithm runs in

exponential time, because verifying properties of the form  $P_{\bowtie p}[a_1 \cup a_2]$  on MDPs can be done in polynomial time [10, 20], and the size of the region graph  $\text{Reg}[P, \phi]$  is exponential in the size of  $P$  (the size of  $P$  is the sum of the number of locations and clocks, the size of the binary encoding of the constants used in invariant and enabling conditions, and the size of the encoding of its transition probabilities, which are expressed as a ratio between two natural numbers, each in binary). The problem is EXPTIME-complete, where an EXPTIME lower bound can be obtained even for the restricted case in which the PTA has only two clocks (for PTAs with one clock, the model-checking problem for certain restricted classes of properties, such as PCTL properties or time-bounded until properties with probability thresholds 0 or 1 only, is PTIME-complete) [44].

We now consider the case where  $P$  is not structurally divergent. For  $P_{\leq p}[\phi_1 \cup \phi_2]$  or  $P_{< p}[\phi_1 \cup \phi_2]$  (i.e., where maximum probabilities are needed), the method given above can be applied without further modification [63]. The case of  $P_{\geq p}[\phi_1 \cup \phi_2]$  or  $P_{> p}[\phi_1 \cup \phi_2]$  (which needs minimum probabilities) is more involved. Intuitively, adversaries which can avoid reaching states satisfying  $\phi_2$  only because they perform some non-divergent behaviour can result in the minimum probability of satisfying  $\phi_1 \cup \phi_2$  computed over all adversaries being lower than the probability computed over divergent adversaries only. An algorithm that resolves this problem, based on the computation of the maximum probability of satisfying the dual path property  $\neg(\phi_1 \cup \phi_2)$ , is presented in [63]. The model-checking problem remains EXPTIME-complete in this case. Details on model-checking algorithms for the strictly divergent adversaries described in Sect. 3 are also given in [63].

Finally, we note that region equivalence is an example of a time-abstracting bisimulation, a relation which combines time-abstracting bisimulation [3, 59] and (probabilistic) bisimulation [58, 62]. As with region equivalence, time-abstracting bisimilar PTA states satisfy the same formulae (without  $R_{\bowtie q}^r[\rho]$  formulae). An algorithm for computing a time-abstracting equivalence relation of a PTA, which may be coarser than region equivalence, has been presented in [28]. This approach is described as being applicable to formulae of the logic without  $R_{\bowtie q}^r[\rho]$  or  $P_{\bowtie p}[\phi_1 \cup \phi_2]$ ; however, as explained in Sect. 4, time-bounded until properties can be reduced to unbounded until properties on a modified PTA.

## 5.2 The boundary region graph

The region graph construction presented above is not sufficient for verifying reward properties. In particular, we note that, for a particular region  $(l, \alpha)$ , the values  $\mathbb{E}_{P,s}^{\text{r},\min}(\rho)$  and  $\mathbb{E}_{P,s}^{\text{r},\max}(\rho)$  will generally not be uniform on states  $s \in (l, \alpha)$ . We now briefly describe a generalisation of the region graph, called the *boundary region graph* [43], which is a finite MDP equipped with a reward structure on which we can decide whether a particular state  $s$  of a PTA satisfies an  $R_{\bowtie q}^r[F \phi]$  property, under the restriction that there is no nesting of the  $R_{\bowtie q}^r[\cdot]$  operator and the bound  $\bowtie$  is non-strict. For  $R_{\bowtie q}^r[C^{\leq k}]$  and  $R_{\bowtie q}^r[I^{\leq k}]$  properties, the reduction given in Sect. 4 can be used under the same restrictions.

The boundary region graph construction is an extension of the *corner point abstraction* for timed automata [24]. The underlying idea is that optimal behaviour (resulting in minimum or maximum expected rewards) corresponds to the case in which edge transitions of the PTA are taken either in a clock equivalence class in which the value of at least one clock equals an integer, or close to a boundary of a clock equivalence class in which the fractional parts of all clocks are positive: in this case, for computing the accumulation of rewards over time, it is necessary to distinguish which of the class's boundaries is considered. Furthermore, the exact position on the clock equivalence boundaries is determined by the values of the clocks in the state  $s$ , and must also be encoded in the boundary region graph. The reward structure for the boundary region graph is derived directly from that of the PTA. From

the boundary region graph (a finite-state MDP), we can compute the minimum or maximum expected accumulated reward to a target set (see, e..g [32]). We can then obtain either  $\mathbb{E}_{P,s}^{r,\min}(F \phi)$  or  $\mathbb{E}_{P,s}^{r,\max}(F \phi)$  and hence decide whether  $s$  satisfies  $R_{\triangleright q}^r[F \phi]$ .

### 5.3 Digital clocks

Region graph-based approaches are not usually practically applicable since the region graphs are generally of a prohibitive size. Thus, various other PTA model checking approaches have been developed. We first describe the *digital clocks* method [52], which restricts the standard continuous-time semantics of a PTA so that only time transitions of duration 1 occur. This means that clocks take only integer, rather than real, values. Using this fact, and knowing there is a maximal constant  $c_x$  to which each clock  $x$  is compared in PTA  $P$  and property  $\phi$ , we can again build and analyse a finite-state MDP. This approach builds on the use of digital clocks for (non-probabilistic) timed automata verification [8, 19, 39].

The digital clocks method is applicable to properties of the form  $P_{\triangleright p}[\psi]$  and  $R_{\triangleright q}^r[\rho]$  without nesting of further  $P_{\triangleright p}[\cdot]$  and  $R_{\triangleright q}^r[\cdot]$  operators within the subformulae  $\psi$  and  $\rho$  (see [52] for an explanation as to this limitation). It can only be used to determine satisfaction in states where all clocks take natural-numbered values, and we will restrict our attention to checking satisfaction in the initial state. The correctness of the digital clocks method also relies on the assumption that  $P$  and  $\phi$  are *closed*, meaning that all clock constraints of the form  $x \leq d$  or  $d \leq x$  are contained within an even number of negations. Furthermore, all invariant and enabling conditions of  $P$  are assumed to be *diagonal-free*, meaning that constraints of the form  $x + c \leq y + d$  are not permitted. Any PTA can be transformed into one containing only diagonal-free constraints by applying the construction of [15] (the construction is presented for timed automata, and requires minor modifications for PTAs).

For a digital clock valuation  $v \in \mathbb{N}^{\mathcal{X}}$ , let  $v \oplus 1$  be the clock valuation such that  $(v \oplus 1)(x) = \min\{v(x) + 1, c_x + 1\}$  for all  $x \in \mathcal{X}$ . The *digital clock semantics* of  $P$  and  $\phi$  is defined as for the standard semantics, except that the rule for time transitions restricts durations to 1, and each clock  $x$  can increase to at most  $c_x + 1$ . Formally, the digital clock semantics of a closed PTA  $P$  is defined as the finite-state MDP  $\text{Dgt}(P, \phi) = (S, (\vec{l}, \mathbf{0}), \text{Act} \cup \{1\}, \text{Steps}, \text{lab})$  where:

- $S = \{(l, v) \in L \times \mathbb{N}^{\mathcal{X}} \mid v \models \text{inv}(l) \wedge (\forall x \in \mathcal{X}. v(x) \leq c_x + 1)\}$ ;
- $\text{Steps}((l, v), a) = \lambda$  if and only if either:
  - *time transitions*:  $a = 1, v \oplus 1 \models \text{inv}(l)$  and  $\lambda = \mu_{(l, v \oplus 1)}$ ;
  - *action transitions*:  $a \in \text{Act}, v \models \text{enab}(l, a)$ , and, for any  $(l', v') \in S$ :

$$\lambda(l', v') = \sum \{|\text{prob}(l, a)(X, l') \mid X \in 2^{\mathcal{X}} \wedge v' = v[X:=0]\}|;$$

- $\text{lab}(l, v) = \mathcal{L}(l)$  for each  $(l, v) \in S$ .

The number of states of the digital clock semantics of  $P$  is bounded by  $|L| \cdot \prod_{x \in \mathcal{X}} (c_x + 1)$ .

Model checking for formulae of the form  $P_{\triangleright p}[\psi]$  and  $R_{\triangleright q}^r[\rho]$  without nesting can then be carried out directly on the finite MDP from the digital clock semantics. For  $P_{\triangleright p}[\psi]$  formulae, we proceed as in the case of the region graph in Sect. 5.1. For a formula  $R_{\triangleright q}^r[F \phi]$  formulae and PTA reward structure  $\mathbf{r} = (\mathbf{r}_{Act}, \mathbf{r}_L)$ , we proceed as follows. We construct the reward structure  $r = (r_S, r_{Act})$  where  $r_S(l, v) = 0, r_{Act}((l, v), 1) = \mathbf{r}_L(l)$  and  $r_{Act}((l, v), a) = \mathbf{r}_{Act}(l, a)$  for all  $(l, v) \in S$  and  $a \in \text{Act}$ . We then use standard algorithms for MDPs to compute the minimum or maximum expected reward to reach the set of states  $\text{Sat}(\phi)$ . The cases for  $\rho = C^{\leq k}$  and  $\rho = I^{\leq k}$  use the reductions presented in Sect. 4.



## 5.4 Backwards reachability

The next method we consider is *backwards reachability* [56], which provides model checking for properties without  $R_{\bowtie q}^f[\rho]$  operators. This is based on the repeated application of a predecessor operation that, given a set of states  $S'$ , returns the set of states that can reach  $S'$  by performing an action and then letting time pass. Sets of states are represented by *symbolic states*, pairs  $z = (l, \zeta)$  comprising a location  $l$  and a clock constraint  $\zeta$  over  $\mathcal{X}$ , representing the set of states  $\{(l, v) \mid v \models \zeta\}$ .

This approach is an adaptation of the algorithm in [40] for model checking timed automata. Whereas the latter just requires iteration of a generic predecessor operation, for PTAs it is necessary to retain information about the probabilities of the PTA edges used along paths. First, the predecessor operation is parameterised by actions and edges of the PTA. Then, as the predecessor iteration proceeds, a graph is constructed, where the nodes are the generated symbolic states, and an edge is added from symbolic state  $z$  to symbolic state  $z'$  if  $z$  was generated from  $z'$  by a predecessor operation. The edge  $(z, z')$  is labelled by the corresponding action and PTA edge. The symbolic states generated in this manner do *not* form a partition of the state space of the PTA, unlike the region graph or time-abstracting bisimulation approaches described in Sect. 5.1.

After the iteration of predecessor operations terminates (which is guaranteed because a symbolic state corresponds to a union of regions), the obtained graph can then be used as a basis for the construction of a finite MDP. To build the probabilistic transition function, information has to be combined from different symbolic states in order to obtain the exact combination of PTA edges (corresponding to a particular action) available in states of the PTA. This is done by computing the conjunctions of symbolic states which have at least one outgoing edge labelled with the same action, and adding the corresponding graph edges to the newly generated symbolic states. For more information, see [56].

The approach outlined above only applies to  $P_{\bowtie p}[\psi]$  operators where  $\bowtie \in \{\leq, <\}$ , i.e. it computes only maximum probabilities for the PTA. For the case  $\bowtie \in \{\geq, >\}$ , which needs minimum probabilities, the method needs to be adapted [56]. Like the reduction described in Sect. 5.1 for the region graph approach on non-structurally-divergent PTAs, the solution taken in [56] works by considering the dual formula.

## 5.5 Abstraction refinement using stochastic games

The final model checking technique we discuss for PTAs is *abstraction refinement* using *stochastic games*, which can verify  $P_{\bowtie p}[\psi]$  properties via computation of reachability probabilities. This approach uses the game-based notion of abstraction put forward for MDPs in [49] and the corresponding refinement techniques proposed in [46]. The method can be applied to PTAs containing diagonal-free constraints only, because one of its procedures is a (forwards) reachability exploration and, by the results of Bouyer [22], such a procedure is only correct if the considered PTA is diagonal-free.

The idea of [49] is to build an abstraction of a large or infinite-state MDP based on a finite partition of its state space. Abstractions take the form of stochastic two-player games, a generalisation of MDPs in which there are two distinct types of nondeterminism, each controlled by a separate player. In this case, player 1 controls nondeterminism introduced by the abstraction, and player 2 controls nondeterminism from the original MDP. An analysis of the *optimal probabilities* in the stochastic game (e.g. the maximum probability that player 1 can achieve for some objective, assuming that player 2 aims to minimise it) yields lower and upper bounds on reachability probabilities for the original MDP.

The abstraction-refinement framework of [46] provides a way to automatically construct stochastic game abstractions, by iteratively refining abstractions until the difference between the lower and upper bounds produced is below some desired level of precision  $\varepsilon$ . In [50], this technique is adapted to an algorithm to compute exact (minimum or maximum) reachability probabilities for PTAs. First, a *reachability graph* is constructed, based on a successor operation that returns the set of states that can be reached by performing an action and then letting time pass. This works in a similar style to the classic approach to verifying (non-probabilistic) timed automata. From this, a stochastic game abstraction is created, over *symbolic states* of the same form as in Sect. 5.4. The abstraction is repeatedly analysed and refined until the exact required probabilities are obtained (i.e.  $\varepsilon = 0$ ). The iterative refinement process is guaranteed to terminate thanks to the fact that there is an underlying finite time-abstracting bisimulation quotient, namely the region graph.

## 5.6 A comparison of the methods

We briefly summarise the relative merits of the three practical approaches to PTA model checking discussed above. In terms of applicability, digital clocks is currently the only method for computing expected rewards, but handles only closed (non-strict) and diagonal-free clock constraints. The other limitations of the remaining two methods are that abstraction refinement only applies to PTAs with diagonal-free clock constraints, and backwards reachability requires (non-trivial) adaptation to compute minimum probabilities.

In terms of efficiency and scalability, the digital clocks approach has proved to work well in practice, but performance suffers for large numbers of clocks or when large constants appear in the clock constraints; in these cases, the other two methods described have been shown to work better. The techniques based on parameter synthesis of [7] can be used to reduce the size of the constants of clock constraints for subclasses of PTAs and for probabilistic properties without time bounds. Another good approach to improving performance is the use of symbolic (binary decision diagram based) implementations. The original implementation of backward reachability [56] showed that it yielded relatively small MDPs, but that the algorithm can be expensive to implement. Experimental results for abstraction refinement [50] later demonstrated better performance in all cases. However, subsequent optimisations presented for backwards reachability [18] have led to much better performance, improving on abstraction refinement in many cases.

## 5.7 Implementations and tool support

Thanks to increasing interest in the verification of probabilistic real-time systems, a variety of related software tools have recently been developed. The probabilistic model checker PRISM [51], for example, which provides verification of Markov chains and MDPs, now also supports PTAs, via the digital clocks, backwards reachability and abstraction-refinement methods. A second tool is *mcpta* [38], which applies the digital clocks method to translate a subset of the modelling language Modest [21] directly into the PRISM modelling language.

Fortuna [18] is a tool that focuses on PTAs augmented with prices (called rewards in this paper). In particular, it implements the semi-algorithm of [17] for computing the maximum probability of reaching a target while accumulating a reward below a given threshold (this problem is shown to be undecidable in [16]). Since this algorithm generalises the backwards reachability method of [56], computation of (maximum) reachability probabilities is also supported. Several optimisations of the basic algorithm are also implemented. Finally,

UPPAAL PRO [68] is an extension of the popular timed automaton verifier UPPAAL. It computes the maximum probability of reaching a set of target states of a PTA, by progressively partitioning the state space, constructing and solving a finite MDP at each step.

## 6 Case studies

PTAs have been used for the modelling and analysis of a wide variety of systems, including communication protocols [33, 55], aviation security systems [34], streaming download protocols [69] and service level agreements [48]. In this section, we give an illustration of the PTA model checking techniques described in this paper by presenting two case studies: a non-repudiation protocol and a task-graph scheduling problem.

### 6.1 Markowitch & Roggeman’s non-repudiation protocol

This case study analyses Markowitch & Roggeman’s non-repudiation protocol for information transfer [61]. Our models extend those presented previously in [57]. One party, the *originator*, sends information to a second party, the *recipient*. *Repudiation* is defined as the denial of either party of having participated in all or part of the information transfer. For example, in electronic commerce, if the information represents the transfer of a service, then *non-repudiation* ensures the client (the recipient) cannot deny receiving the service as a reason for non-payment.

The protocol of Markowitch & Roggeman is probabilistic and does not require a trusted third party. It achieves the following non-repudiation properties:

- “ $\epsilon$ -fair”: at each step of the protocol run, either both parties receive their expected items, or the probability that a cheating party gains any valuable information about its expected items and the other party gains nothing is at most  $\epsilon$ ;
- “time-bounded”: if at least one party behaves correctly, then the protocol will complete within a finite amount of time (with probability 1);
- “viable”: if both parties behave correctly and finish the protocol, then they both receive their expected items at the end of the protocol (with probability 1).

The steps of the protocol are outlined in Fig. 3. First, to prevent replay attacks, the recipient  $R$  selects a date  $D$  which it sends, along with a request, to the originator  $O$ . Next,  $O$  randomly selects an integer  $n$  representing the number of steps of the protocol (which is never revealed to  $R$  during the execution) and computes functions  $f_i$  such that their composition satisfies the following:

$$f_n(\text{message}) \circ (f_{n-1}(\text{message}) \circ (\dots \circ (f_2(\text{message}) \circ f_1(\text{message})) \dots)) = \text{message}.$$

The composition operator  $\circ$  needs to be non-commutative to ensure that the recipient cannot start to compute the message until the final message  $f_1(\text{message})$  has been received.

**Fig. 3** The steps of Markowitch & Roggeman’s non-repudiation protocol

|  |   |
|--|---|
| The recipient chooses the date $D$   |   |
| 1.   | $R \rightarrow O : \text{Sign}_R(\text{request}, R, O, D)$      |
| The originator checks $D$ , chooses $n$ and computes functions $f_1, \dots, f_n$ |   |
| 2.   | $O \rightarrow R : \text{Sign}_O(f_n(\text{message}), O, R, D)$ |
| 3.   | $R \rightarrow O : \text{Sign}_R(\text{ack}_1)$                 |
| ...  | ...   |
| $2n$ .   | $O \rightarrow R : \text{Sign}_O(f_1(\text{message}), O, R, D)$ |
| $2n + 1$ .   | $R \rightarrow O : \text{Sign}_R(\text{ack}_n)$                 |

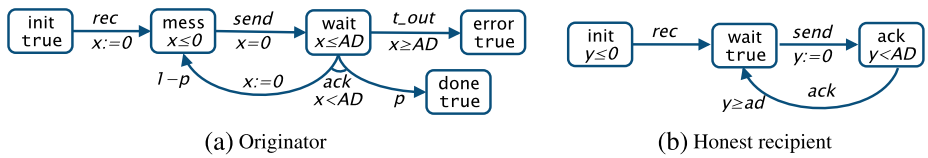


Fig. 4 PTAs used to model the non-repudiation protocol

To prevent the recipient from gaining an advantage, if the originator does not receive an acknowledgement within a certain time bound (denoted  $AD$ ), the protocol is stopped and the originator states that the recipient is trying to cheat. The time bound is chosen such that it is greater than the time it takes for a recipient to send a reply, but is not sufficient for the recipient to be able to compute the composition  $f_n(message) \circ (\dots \circ f_i(message) \dots)$  for any  $i < n$ , i.e., the recipient does not have time to check if it has received the complete message before sending an acknowledgement.

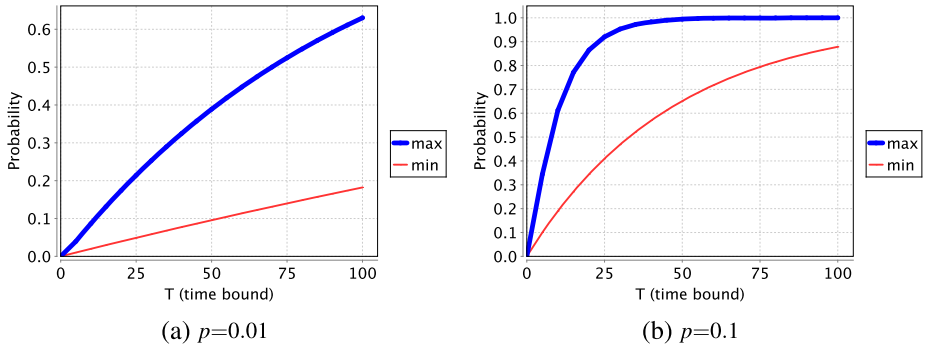
We consider two different versions of the protocol. In the first, both the originator and recipient act *honestly*, while in the second the recipient can act *maliciously* (i.e., stop early by not returning an acknowledgement). For each, we assume that the choice of  $n$  is made by the recipient according to a geometric distribution with parameter  $p$  and the minimum time for the recipient to send an acknowledgement is  $ad$ . For the malicious version, we consider two variants. The first corresponds to the one described in [61], where the only malicious behaviour corresponds to stopping early. In the second, we introduce a more powerful malicious recipient, which has access to a method that takes time less than  $AD$  and with probability  $\frac{1}{4}$  correctly computes the composition while with probability  $\frac{3}{4}$  fails to compute the composition.

Each model is the parallel composition of two PTAs, one representing the originator and one the recipient. These synchronise on the actions  $rec$ ,  $send$  and  $ack$ , corresponding to the recipient initiating the transaction, the originator sending messages to the recipient and the recipient sending acknowledgements back. The PTAs for the originator and (honest) recipient are shown in Fig. 4. The probabilistic choice in the originator correctly models selection of  $n$  according to a geometric distribution with parameter  $p$ , since the probability of each message being the last one to be sent is  $p$ . Notice that, in the PTA for the honest recipient, an acknowledgement is sent after between  $ad$  and  $AD$  time units.

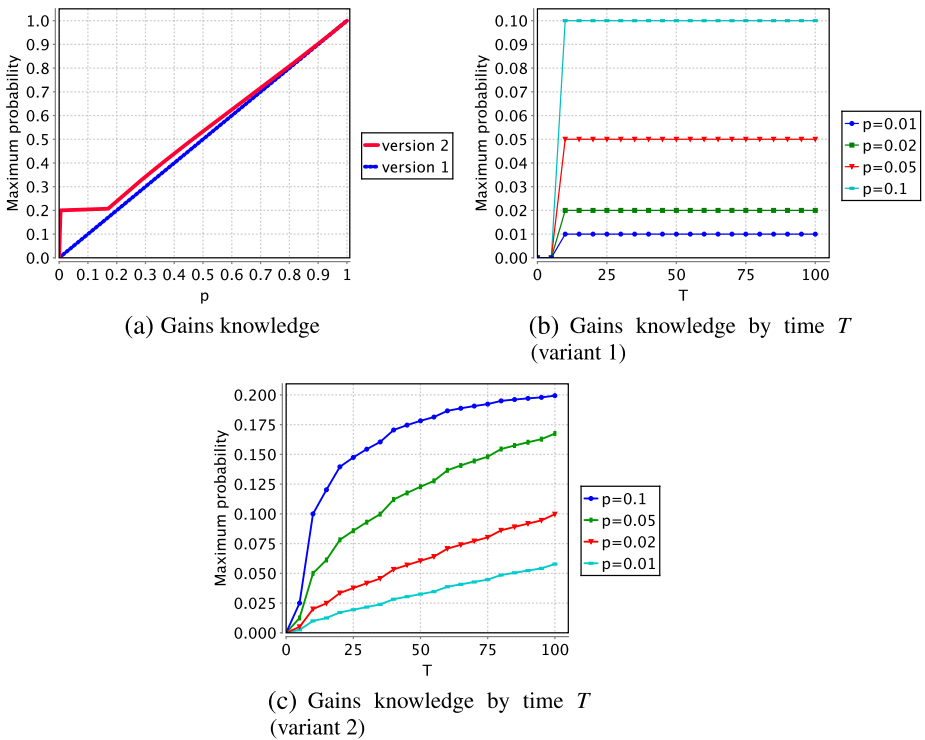
We analysed these PTA models in PRISM using the stochastic games technique (since the originator PTA contains strict inequalities, the digital clocks method is not applicable). For the analysis, we assume  $AD = 5$  and  $ad = 1$ .

For the honest version of the protocol, the first property we consider is “*time-bounded*”. More precisely, we check the formula  $P_{\geq 1}[F \text{ done}]$ , stating that the minimum probability of the protocol terminating correctly is 1. Next, we investigate the performance of the protocol with the quantitative properties  $P_{\min=?}[F \leq T \text{ done}]$  and  $P_{\max=?}[F \leq T \text{ done}]$ , i.e., the minimum and maximum probability of termination by time  $T$ . Figure 5 plots these values for  $T$  between 0 and 100, with  $p = 0.01$  and  $p = 0.1$ . We see that increasing the parameter  $p$  improves the performance of the protocol when the parties behave honestly. However, as we shall see below, when the recipient behaves maliciously, increasing this parameter comes at a cost since it also increases the likelihood that the recipient gains an advantage.

For the two models with a malicious recipient, we find that the minimum and maximum probability of the protocol terminating correctly ( $P_{\min=?}[F \text{ done}]$  and  $P_{\max=?}[F \text{ done}]$ ) are 0 and 1, respectively. The minimum probability is achieved when the (malicious) recipient returns no acknowledgements at all and the maximum when it acts honestly.

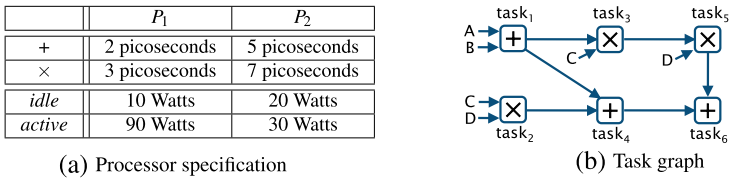


**Fig. 5** Probability that the protocol terminates successfully by time  $T$  (honest version)



**Fig. 6** Maximum probability that the recipient gains knowledge (malicious versions)

Figure 6 presents results for the properties  $P_{\max=?}[F \text{ unfair}]$  and  $P_{\max=?}[F \leq T \text{ unfair}]$ , for both variants of the malicious recipient and several values of parameter  $p$ . These correspond to the maximum probability that the recipient (eventually, or within time bound  $T$ ) gains an advantage. For the first variant (describing the scenario in [61]), Fig. 6(a) shows that the protocol is indeed  $\epsilon$ -fair with  $\epsilon$  equal to  $p$ . Essentially, all this recipient can do to gain knowledge is to correctly guess which message is the last (which, when a message arrives, is true with probability  $p$ ). As shown in Fig. 6(b), the probability of gaining an advantage



**Fig. 7** Scheduling problem for computing the term  $D \times (C \times (A + B)) + ((A + B) + (C \times D))$

over time remains constant after the arrival of the first message: since each message has an equal chance of being the last, there is nothing to be gained by waiting for a later one. The figures also show that the malicious recipient in the second variant of the model (which has a chance of correctly decoding the message before the deadline  $AD$ ) has a greater chance of gaining knowledge, thanks to its additional power.

### 6.2 Task-graph scheduling

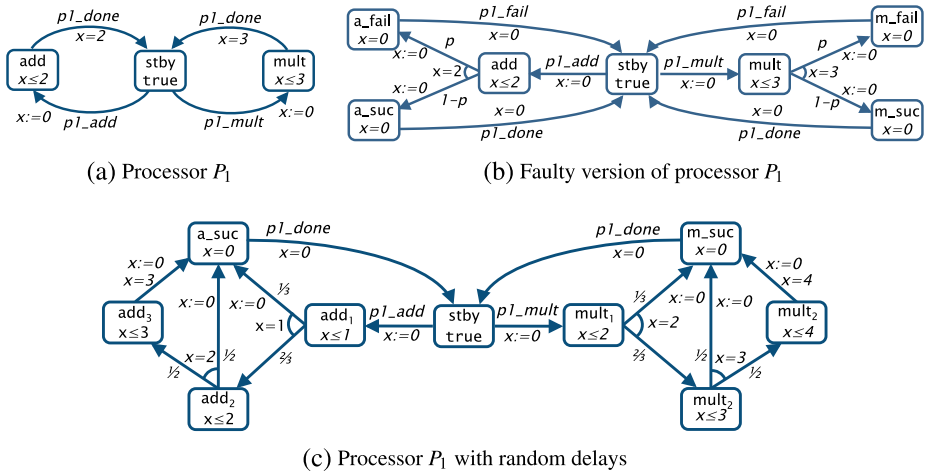
One common use of (non-probabilistic) timed automata is the formalisation and solution of scheduler optimisation problems. In this case study, we consider an extension of the *task-graph scheduling* problem described in [27]. We show how PTAs can be used to introduce uncertainty with regards to time delays and to consider the possibility of failures. We demonstrate how the digital clocks method can determine optimal schedulers for these problems, in terms of the (expected) time and energy consumption required to complete all tasks.

*The basic model* First, we introduce the basic problem formalisation and model presented in [27]. This considers the example of evaluating the expression  $D \times (C \times (A + B)) + ((A + B) + (C \times D))$ , where its subterms are evaluated on two processors,  $P_1$  and  $P_2$ . Processor  $P_1$  is faster than  $P_2$ , but also consumes more energy. Figure 7(a) shows the specifications of the processors, in terms of their processing times and energy usage for addition and multiplication operations. Figure 7(b) presents the *task graph* for this example, illustrating the set of tasks (corresponding to subterms of the expression) and the dependencies that exist between tasks (in terms of their required order of evaluation).

In [27], a (non-probabilistic) timed automaton model is built, consisting of the parallel composition of one automaton for each processor and one automaton for the scheduler which decides when tasks get performed and on which processors. The scheduler automaton includes integer-valued variables to indicate whether a task still has to be processed, is currently being processed or has been completed. To ensure that the restrictions of the task graph are met, there are conditions placed on enabling conditions such that a task cannot be scheduled until all of its dependencies have been computed.

The timed automaton for processor  $P_1$  is given in Fig. 8(a). The actions  $pl\_add$  and  $pl\_mult$  correspond to an addition and multiplication task being started on  $P_1$ , and  $pl\_done$  indicates the completion of a task. The clock  $x$  is used to record the time that a task has been running and is initialised when a task is started. The automaton for  $P_2$  is similar except that the action names change and the delays are modified to reflect the values in Fig. 7(a).

The time and energy consumed to complete all tasks are modelled as rewards (called prices in [27]). We introduce reward structures `time` and `energy`, respectively, to represent each of these quantities. For the case of time, only the scheduler automaton has a non-zero reward structure, where the location reward is 1 in each location and all action rewards are zero. For the case of energy consumption, the scheduler automaton has a zero



**Fig. 8** PTAs for the task-graph scheduling case study

reward structure, while each processor has a location reward equal to the current rate of energy usage (as shown in Fig. 7(a)) and has zero action rewards.

We built a PTA model for this case study using PRISM and, by applying the digital clocks method, calculated both the minimum (expected) time and energy consumption for completion of all tasks. For this, we used the two quantitative reward properties  $R_{\min=?}^{\text{time}}[F \text{ complete}]$  and  $R_{\min=?}^{\text{energy}}[F \text{ complete}]$ . We also used PRISM to generate the corresponding schedulers that achieve these optimal values. The results agree with those reported in [27]. A scheduler that minimises the elapsed time requires 12 picoseconds to complete all tasks and schedules the tasks as follows:

| time  | 1     | 2 | 3     | 4 | 5     | 6 | 7     | 8 | 9     | 10 | 11    | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 |
|-------|-------|---|-------|---|-------|---|-------|---|-------|----|-------|----|----|----|----|----|----|----|----|----|
| $P_1$ | task1 |   | task3 |   | task4 |   | task5 |   | task4 |    | task6 |    |    |    |    |    |    |    |    |    |
| $P_2$ |       |   | task2 |   |       |   |       |   |       |    |       |    |    |    |    |    |    |    |    |    |

On the other hand, a scheduler optimising the energy consumption requires 1.3200 nanojoules (and 19 picoseconds) and makes the following scheduling decisions:

| time  | 1     | 2 | 3     | 4 | 5     | 6 | 7 | 8 | 9 | 10    | 11 | 12 | 13 | 14 | 15 | 16 | 17    | 18 | 19 | 20 |
|-------|-------|---|-------|---|-------|---|---|---|---|-------|----|----|----|----|----|----|-------|----|----|----|
| $P_1$ | task1 |   | task3 |   | task4 |   |   |   |   |       |    |    |    |    |    |    |       |    |    |    |
| $P_2$ |       |   | task2 |   |       |   |   |   |   | task5 |    |    |    |    |    |    | task6 |    |    |    |

Due to the additional energy consumption of processor  $P_1$ , the first scheduler above, which optimises the time for task completion, requires 1.3900 nanojoules.

*Random task execution times* Now, we extend the formalisation of the task-graph problem, making the time required for each processor to perform a task probabilistic (in a more general setting, we can easily envisage situations where the exact time required to complete a task is unknown, but can be represented by some probability distribution). More precisely, we consider the following simple scenario. If, in the original problem the time for a processor to perform a task was  $k \in \mathbb{N}$ , we suppose now that the time taken is uniformly distributed between the delays  $k-1, k$  and  $k+1$ , e.g. the time for  $P_1$  to perform a multiplication operation is either 1, 2 or 3 and the probability of each execution time is  $\frac{1}{3}$ .

The PTA for processor  $P_1$  with random delays is presented in Fig. 8(c) where, to ease notation, we have omitted action labels if they do not synchronise. Additional locations



are added to encode the random delays. For example, in the case of multiplication, with probability  $\frac{1}{3}$  the task completes after 2 time units; with probability  $\frac{2}{3}$ , the PTA moves to a location where, with probability  $\frac{1}{2}$  the task completes after 1 additional time unit (i.e., of a total of 3 time units) or moves to a location where the task completes after 2 more time units (i.e., 4 time units in total). When the task completes, the PTA moves to a location where no time can pass (clock  $x$  is reset upon entering and the invariant of the location is  $x \leq 0$ ) and immediately notifies the scheduler the task is computed through action  $pI\_done$ . To prevent the scheduler from seeing into the future when making decisions, the probabilistic choice for task completion is made on completion rather than on initialisation.

Analysing this model, we find that the optimal expected time and energy consumption to complete all tasks equals 12.226 picoseconds and 1.3201 nanojoules, respectively. This improves on the results obtained using the optimal schedulers for the original model, where the expected time and energy consumption equal 13.1852 picoseconds and 1.3211 nanojoules. Examining the optimal schedulers, we find that they change their decision based upon the delays of previously completed tasks. For example, for elapsed time, the optimal scheduler starts as for the non-probabilistic case, first scheduling task1 followed by task3 on  $P_1$  and task2 on  $P_2$ . However, it is now possible for task2 to complete before task3 (if the execution times for task1, task2 and task3 are 3, 6 and 4 respectively), in which case the optimal scheduler now makes a different decision from the non-probabilistic case. Under one possible set of execution times for the remaining tasks, the optimal scheduling is as follows:

|       |       |   |       |   |   |       |   |   |       |    |    |    |    |    |    |    |    |    |    |    |  |
|-------|-------|---|-------|---|---|-------|---|---|-------|----|----|----|----|----|----|----|----|----|----|----|--|
| time  | 1     | 2 | 3     | 4 | 5 | 6     | 7 | 8 | 9     | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 |  |
| $P_1$ | task1 |   | task3 |   |   | task5 |   |   | task6 |    |    |    |    |    |    |    |    |    |    |    |  |
| $P_2$ | task2 |   |       |   |   | task4 |   |   |       |    |    |    |    |    |    |    |    |    |    |    |  |

*Adding a faulty processor* As a second extension of the scheduling problem, we add a third processor  $P_3$  which consumes the same energy as  $P_2$  but is faster (addition takes 3 picoseconds and multiplication 5 picoseconds). However, this comes at a cost: there is a chance (probability  $p$ ) that the processor fails and the computation must be rescheduled and performed again.

In Fig. 8(b), we show the PTA for the faulty version of processor  $P_1$ . In this PTA, when a task completes, there is a probabilistic choice between moving to a location corresponding to successful completion and one to failure. In both cases, we move to a location where no time can pass and immediate notify the scheduler of either the success or failure of the computation. The automaton for the scheduler also changes for this model since it must react to the failure signals from the processors. In addition, the reward structure *energy* is extended to include the energy consumed by the additional processor.

The graphs in Fig. 9 plot the optimal expected time and energy consumption for this extended model as the failure probability  $p$  varies. The dashed lines show the optimal results for the original model, i.e., when not using the processor  $P_3$ . As can be seen, once the probability of failure becomes sufficiently large, there is no gain in using the processor  $P_3$  but, while when the probability of failure is small, it uses offers considerable gains in performance. To illustrate this fact, below we give a scheduler that optimises (minimises) the expected energy consumption when  $p = 0.5$ . Dark boxes for tasks are used to denote processor  $P_3$  failing to complete a task correctly, meaning that the task needs to be rescheduled.

|       |       |   |   |       |   |       |   |   |   |    |    |    |    |    |       |    |    |    |    |    |  |
|-------|-------|---|---|-------|---|-------|---|---|---|----|----|----|----|----|-------|----|----|----|----|----|--|
| time  | 1     | 2 | 3 | 4     | 5 | 6     | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15    | 16 | 17 | 18 | 19 | 20 |  |
| $P_1$ |       |   |   | task3 |   |       |   |   |   |    |    |    |    |    | task6 |    |    |    |    |    |  |
| $P_2$ | task2 |   |   |       |   | task5 |   |   |   |    |    |    |    |    |       |    |    |    |    |    |  |
| $P_3$ | task1 |   |   |       |   | task4 |   |   |   |    |    |    |    |    |       |    |    |    |    |    |  |

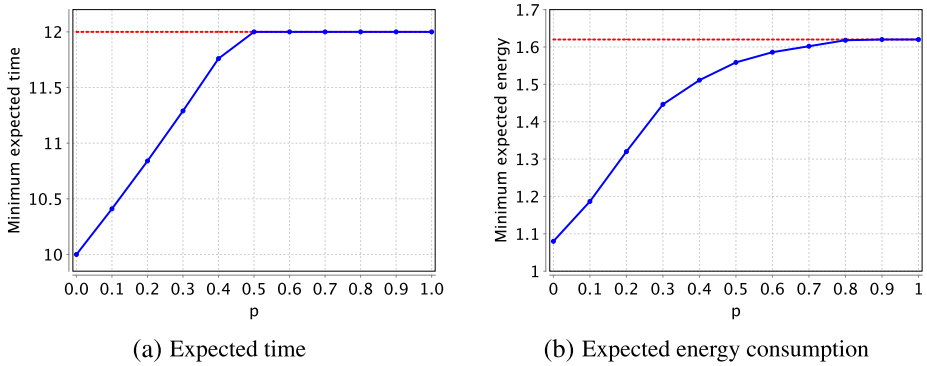


Fig. 9 Optimal expected time and energy consumption as the failure probability of processor  $P_3$  varies

|       |       |   |       |       |   |       |   |   |       |    |    |    |       |    |    |    |    |    |    |    |
|-------|-------|---|-------|-------|---|-------|---|---|-------|----|----|----|-------|----|----|----|----|----|----|----|
| time  | 1     | 2 | 3     | 4     | 5 | 6     | 7 | 8 | 9     | 10 | 11 | 12 | 13    | 14 | 15 | 16 | 17 | 18 | 19 | 20 |
| $P_1$ |       |   |       | task1 |   | task3 |   |   | task5 |    |    |    | task6 |    |    |    |    |    |    |    |
| $P_2$ |       |   | task2 |       |   |       |   |   | task4 |    |    |    |       |    |    |    |    |    |    |    |
| $P_3$ | task1 |   |       |       |   |       |   |   |       |    |    |    |       |    |    |    |    |    |    |    |

|       |       |   |       |       |   |   |   |       |   |       |    |    |    |       |    |    |    |    |    |    |
|-------|-------|---|-------|-------|---|---|---|-------|---|-------|----|----|----|-------|----|----|----|----|----|----|
| time  | 1     | 2 | 3     | 4     | 5 | 6 | 7 | 8     | 9 | 10    | 11 | 12 | 13 | 14    | 15 | 16 | 17 | 18 | 19 | 20 |
| $P_1$ |       |   |       | task3 |   |   |   |       |   | task4 |    |    |    | task6 |    |    |    |    |    |    |
| $P_2$ |       |   | task2 |       |   |   |   |       |   | task5 |    |    |    |       |    |    |    |    |    |    |
| $P_3$ | task1 |   |       |       |   |   |   | task4 |   |       |    |    |    |       |    |    |    |    |    |    |

Notice that the scheduler uses the processor  $P_3$  for task1 and, if this task is completed successfully, it later uses  $P_3$  for task4. However, if task1 fails to complete,  $P_3$  is not used again.

### 7 Conclusions

In this paper, we have presented an introduction to the model of probabilistic timed automata and summarised the various techniques developed to perform probabilistic model checking. Verification of probabilistic real-time systems is an active field of research and further progress is required in several important directions. Examples include the development of verification techniques for *probabilistic timed games* [6, 43] and for *probabilistic hybrid automata* [9, 36, 64]. The former have proved, in the non-probabilistic setting, to being applicable to a variety of useful synthesis problems [12]. The latter provide essential modelling capabilities for domains such as embedded systems and cyber-physical systems; they represent a useful, but more tractable, subclass of the model of stochastic hybrid automata. Other important issues to investigate in the context of PTAs include robustness [7] and continuously-distributed time delays [1, 53, 60].

**Acknowledgements** David Parker was partially supported by ERC Advanced Grant VERIWARE. Jeremy Sproston was partially supported by the project AMALFI (Advanced Methodologies for the AnaLysis and management of the Future Internet, Università di Torino/Compagnia di San Paolo). We thank the anonymous referees for various useful comments.

## References

1. Alur R, Courcoubetis C, Dill D (1991) Model-checking for probabilistic real-time systems. In: Proc of 19th international colloquium on automata, languages and programming (ICALP'91). LNCS, vol 510. Springer, Berlin, pp 115–136
2. Alur R, Courcoubetis C, Dill D (1993) Model checking in dense real time. *Inf Comput* 104(1):2–34
3. Alur R, Courcoubetis C, Halbwachs N, Dill D, Wong-Toi H (1992) Minimization of timed transition systems. In: Cleaveland R (ed) Proc of 3rd int conf on concurrency theory (CONCUR'92). LNCS, vol 630. Springer, Berlin, pp 340–354
4. Alur R, Dill D (1994) A theory of timed automata. *Theor Comput Sci* 126:183–235
5. Alur R, La Torre S, Pappas G (2004) Optimal paths in weighted timed automata. *Theor Comput Sci* 318(3):297–322
6. Alur R, Trivedi A (2011) Relating average and discounted costs for quantitative analysis of timed systems. In: Proc of 11th int conf on embedded software (EMSOFT'11). ACM, New York, pp 165–174
7. André E, Fribourg L, Sproston J (2012) An extension of the inverse method to probabilistic timed automata. *Form Methods Syst Des*. doi:[10.1007/s10703-012-0169-x](https://doi.org/10.1007/s10703-012-0169-x)
8. Asarin E, Maler O, Pnueli A (1998) On discretization of delays in timed automata and digital circuits. In: Sangiorgi D, de Simone R (eds) Proc of 9th int conf on concurrency theory (CONCUR'98). LNCS, vol 1466. Springer, Berlin, pp 470–484
9. Assouramou J, Desharnais J (2011) Analysis of non-linear probabilistic hybrid systems. In: Massink M, Norman G (eds) Proc of 9th workshop on quantitative aspects of programming languages (QAPL'11), pp 104–119
10. Baier C, Kwiatkowska M (1998) Model checking for a probabilistic branching time logic with fairness. *Distrib Comput* 11(3):125–155
11. Beauquier D (2003) On probabilistic timed automata. *Theor Comput Sci* 292(1):65–84
12. Behrmann G, Cournard A, David A, Fleury E, Larsen K, Lime D (2007) UPPAAL-Tiga: time for playing games! In: Proc of 19th international conference on computer aided verification (CAV'07). LNCS, vol 4590. Springer, Berlin, pp 121–125
13. Behrmann G, David A, Larsen KG, Håkansson J, Pettersson P, Yi W, Hendriks M (2006) Uppaal 4.0. In: Proc of 3rd int conf on quantitative evaluation of systems (QEST'06). IEEE Press, New York, pp 125–126
14. Behrmann G, Fehnker A, Hune T, Larsen K, Pettersson P, Romijn J, Vaandrager F (2001) Minimum-cost reachability for linearly priced timed automata. In: Benedetto MD, Sangiovanni-Vincentelli A (eds) Proc of 4th int workshop on hybrid systems: computation and control (HSCC'01). LNCS, vol 2034. Springer, Berlin, pp 147–162
15. Bérard B, Petit A, Diekert V, Gastin P (1998) Characterization of the expressive power of silent transitions in timed automata. *Fundam Inform* 36(2–3):145–182
16. Berendsen J, Chen T, Jansen D (2009) Undecidability of cost-bounded reachability in priced probabilistic timed automata. In: Chen J, Cooper SB (eds) Proc of 6th conf on theory and applications of models of computation (TAMC'09). LNCS, vol 5532. Springer, Berlin, pp 128–137
17. Berendsen J, Jansen D, Katoen JP (2006) Probably on time and within budget: on reachability in priced probabilistic timed automata. In: Proc of 3rd int conf on quantitative evaluation of systems (QEST'06). IEEE Press, New York, pp 311–322
18. Berendsen J, Jansen D, Vaandrager F (2010) Fortuna: model checking priced probabilistic timed automata. In: Proc of 7th int conf on quantitative evaluation of systems (QEST'10). IEEE Press, New York, pp 273–281
19. Beyer D (2001) Improvements in BDD-based reachability analysis of timed automata. In: Oliveira J, Zave P (eds) Int symp of formal methods Europe, FME 2001: formal methods for increasing software productivity. LNCS, vol 2021. Springer, Berlin, pp 318–343
20. Bianco A, de Alfaro L (1995) Model checking of probabilistic and nondeterministic systems. In: Thiagarajan P (ed) Proc of 15th conf on foundations of software technology and theoretical computer science (FSTTCS'95). LNCS, vol 1026. Springer, Berlin, pp 499–513
21. Bohnenkamp H, D'Argenio P, Hermanns H, Katoen JP (2006) Modest: a compositional modeling formalism for hard and softly timed systems. *IEEE Trans Softw Eng* 32(10):812–830
22. Bouyer P (2003) Untameable timed automata! In: Alt H, Habib M (eds) Proc of 20th int symp on theoretical aspects of computer science (STACS'03). LNCS, vol 2607. Springer, Berlin, pp 620–631
23. Bouyer P (2009) From qualitative to quantitative analysis of timed systems. Mémoire d'habilitation. Université Paris 7, Paris, France
24. Bouyer P, Brinksma E, Larsen K (2008) Optimal infinite scheduling for multi-priced timed automata. *Form Methods Syst Des* 32(1):3–23

25. Bouyer P, Chevalier F (2005) On conciseness of extensions of timed automata. *J Autom Lang Comb* 10(4):393–405
26. Bouyer P, Dufourd C, Fleury E, Petit A (2004) Updatable timed automata. *Theor Comput Sci* 321(2–3):291–345
27. Bouyer P, Fahrenberg U, Larsen K, Markey N (2011) Quantitative analysis of real-time systems using priced timed automata. *Commun ACM* 54(9):78–87
28. Chen T, Han T, Katoen JP (2008) Time-abstracting bisimulation for probabilistic timed automata. In: *Proc of 2nd IEEE int symp on theoretical aspects of software engineering (TASE'08)*. IEEE Press, New York, pp 177–184
29. Clarke E, Emerson A (1981) Design and synthesis of synchronization skeletons using branching time temporal logic. In: Kozen D (ed) *Proc of workshop on logic of programs*. LNCS, vol 131. Springer, Berlin
30. Daws C, Yovine S (1995) Two examples of verification of multirate timed automata with KRONOS. In: *Proc of IEEE real-time systems symposium (RTSS'95)*. IEEE Press, New York, pp 66–75
31. Eisentraut C, Hermanns H, Zhang L (2010) On probabilistic automata in continuous time. In: *Proc of 25th annual IEEE symposium on logic in computer science (LICS'10)*. IEEE Comput. Soc., Los Alamitos, pp 342–351
32. Forejt V, Kwiatkowska M, Norman G, Parker D (2011) Automated verification techniques for probabilistic systems. In: Bernardo M, Issarny V (eds) *Formal methods for eternal networked software systems (SFM'11)*. LNCS, vol 6659. Springer, Berlin, pp 53–113
33. Fruth M (2006) Probabilistic model checking of contention resolution in the IEEE 802.15.4 low-rate wireless personal area network protocol. In: *Proc of 2nd int symp on leveraging applications of formal methods, verification and validation (ISOLA'06)*
34. Glässer U, Rastkar S, Vajihollahi M (2008) Modeling and validation of aviation security. In: *Intelligence and security informatics*, vol 135. Springer, Berlin, pp 337–355
35. Gregersen H, Jensen HE (1995) Formal design of reliable real time systems. Master's thesis, Department of Mathematics and Computer Science, Aalborg University
36. Hahn EM, Norman G, Parker D, Wachter B, Zhang L (2011) Game-based abstraction and controller synthesis for probabilistic hybrid systems. In: *Proc of 8th int conf on quantitative evaluation of systems (QEST'11)*. IEEE Press, New York, pp 69–78
37. Hansson H, Jonsson B (1994) A logic for reasoning about time and reliability. *Form Asp Comput* 6(5):512–535
38. Hartmanns A, Hermanns H (2009) A modest approach to checking probabilistic timed automata. In: *Proc of 6th int conf on quantitative evaluation of systems (QEST'09)*. IEEE Press, New York, pp 187–196
39. Henzinger T, Manna Z, Pnueli A (1992) What good are digital clocks? In: Kuich W (ed) *Proc of 19th int colloq on automata, languages and programming (ICALP'92)*. LNCS, vol 623. Springer, Berlin, pp 545–558
40. Henzinger T, Nicollin X, Sifakis J, Yovine S (1994) Symbolic model checking for real-time systems. *Inf Comput* 111(2):193–244
41. Hermanns H (2002) Interactive Markov chains and the quest for quantified quality. LNCS, vol 2428. Springer, Berlin
42. Jensen H (1996) Model checking probabilistic real time systems. In: Bjerner B, Larsson M, Nordström B (eds) *Proc of 7th Nordic workshop on programming theory*. Report, vol 86. Chalmers University of Technology, Chalmers, pp 247–261
43. Jurdziński M, Kwiatkowska M, Norman G, Trivedi A (2009) Concavely-priced probabilistic timed automata. In: Bravetti M, Zavattaro G (eds) *Proc of 20th int conf on concurrency theory (CONCUR'09)*. LNCS, vol 5710. Springer, Berlin, pp 415–430
44. Jurdzinski M, Sproston J, Laroussinie F (2008) Model checking probabilistic timed automata with one or two clocks. *Logical Methods in Computer Science* 4(3)
45. Katoen JP, Hahn EM, Hermanns H, Jansen D, Zapreev I (2009) The ins and outs of the probabilistic model checker MRMC. In: *Proc of 6th int conf on quantitative evaluation of systems (QEST'09)*. IEEE Press, New York, pp 167–176
46. Kattenbelt M, Kwiatkowska M, Norman G, Parker D (2010) A game-based abstraction-refinement framework for Markov decision processes. *Form Methods Syst Des* 36(3):246–280
47. Kemeny J, Snell J, Knapp A (1976) *Denumerable Markov chains*, 2nd edn. Springer, Berlin
48. Krause C, Giese H (2011) Model checking probabilistic real-time properties for service-oriented systems with service level agreements. In: *Proc of 13th int workshop on verification of infinite-state systems (INFINITY'11)*, EPTCS, vol 73. Elsevier, Amsterdam, pp 64–78
49. Kwiatkowska M, Norman G, Parker D (2006) Game-based abstraction for Markov decision processes. In: *Proc of 3rd int conf on quantitative evaluation of systems (QEST'06)*. IEEE Press, New York, pp 157–166

50. Kwiatkowska M, Norman G, Parker D (2009) Stochastic games for verification of probabilistic timed automata. In: Ouaknine J, Vaandrager F (eds) Proc of 7th int conf on formal modelling and analysis of timed systems (FORMATS'09). LNCS, vol 5813. Springer, Berlin, pp 212–227
51. Kwiatkowska M, Norman G, Parker D (2011) PRISM 4.0: verification of probabilistic real-time systems. In: Gopalakrishnan G, Qadeer S (eds) Proc of 23rd int conf on computer aided verification (CAV'11). LNCS, vol 6806. Springer, Berlin, pp 585–591
52. Kwiatkowska M, Norman G, Parker D, Sproston J (2006) Performance analysis of probabilistic timed automata using digital clocks. *Form Methods Syst Des* 29:33–78
53. Kwiatkowska M, Norman G, Segala R, Sproston J (2000) Verifying quantitative properties of continuous probabilistic timed automata. In: Palamidessi C (ed) Proc of 11th international conference on concurrency theory (CONCUR'00). LNCS, vol 1877. Springer, Berlin, pp 123–137
54. Kwiatkowska M, Norman G, Segala R, Sproston J (2002) Automatic verification of real-time systems with discrete probability distributions. *Theor Comput Sci* 282:101–150
55. Kwiatkowska M, Norman G, Sproston J (2003) Probabilistic model checking of deadline properties in the IEEE 1394 FireWire root contention protocol. *Form Asp Comput* 14(3):295–318
56. Kwiatkowska M, Norman G, Sproston J, Wang F (2007) Symbolic model checking for probabilistic timed automata. *Inf Comput* 205(7):1027–1077
57. Lanotte R, Maggiolo-Schettini A, Troina A (2005) Automatic analysis of a non-repudiation protocol. In: Proc of second workshop on quantitative aspects of programming languages (QAPL 2004). ENTCS, vol 112, pp 113–129
58. Larsen K, Skou A (1991) Bisimulation through probabilistic testing. *Inf Comput* 94:1–28
59. Larsen K, Yi W (1997) Time-abstracted bisimulation: implicit specifications and decidability. *Inf Comput* 134(2):75–101
60. Maler O, Larsen K, Krogh B (2010) On zone-based analysis of duration probabilistic automata. In: Proc of 12th int workshop on verification of infinite-state systems (INFINITY'10), EPTCS, vol 39, pp 33–46
61. Markowitch O, Roggeman Y (1999) Probabilistic non-repudiation without trusted third party. In: Proc of 2nd workshop on security in communication networks
62. Segala R, Lynch N (1995) Probabilistic simulations for probabilistic processes. *Nord J Comput* 2(2):250–273
63. Sproston J (2009) Strict divergence for probabilistic timed automata. In: Bravetti M, Zavattaro G (eds) Proc of 20th int conf on concurrency theory (CONCUR 2009). LNCS, vol 5710. Springer, Berlin, pp 620–636
64. Sproston J (2011) Discrete-time verification and control for probabilistic rectangular hybrid automata. In: Proc of 8th int conf on quantitative evaluation of systems (QEST'11). IEEE Press, New York, pp 79–88
65. Tripakis S (1998) The analysis of timed systems in practice. PhD thesis, Université Joseph Fourier, Grenoble
66. Tripakis S (1999) Verifying progress in timed systems. In: Katoen JP (ed) Proc of 5th int AMAST workshop on real-time and probabilistic systems (ARTS'99). LNCS, vol 1601. Springer, Berlin, pp 299–314
67. Tripakis S, Yovine S, Bouajjani A (2005) Checking timed Büchi automata emptiness efficiently. *Form Methods Syst Des* 26(3):267–292
68. UPPAAL PRO web site. <http://people.cs.aau.dk/~arild/uppaal-probabilistic/>
69. Zhang M, Hung DV (2006) Formal analysis of streaming downloading protocol for system upgrading. In: Proc of 4th workshop on quantitative aspects of programming languages (QAPL 06), ENTCS, vol 164, pp 205–224