

Performance Analysis of Probabilistic Timed Automata using Digital Clocks*

Marta Kwiatkowska¹, Gethin Norman¹, David Parker¹ and Jeremy Sproston²

¹ School of Computer Science, University of Birmingham, Edgbaston,
Birmingham B15 2TT, United Kingdom
{M.Z.Kwiatkowska,G.Norman,D.A.Parker}@cs.bham.ac.uk

² Dipartimento di Informatica, Università di Torino, 10149 Torino, Italy
sproston@di.unito.it

July 21, 2003

Abstract

Probabilistic timed automata, a variant of timed automata extended with discrete probability distributions, is a specification formalism suitable for describing both nondeterministic and probabilistic aspects of real-time systems, and is amenable to model checking against probabilistic timed temporal logic properties. In the case of classical (non-probabilistic) timed automata, it has been shown that for a large class of real-time verification problems correctness can be established using an integer-time model, inducing a notion of *digital clocks*, as opposed to the standard dense model of time. Based on these results, we address the question of under what conditions digital clocks are sufficient for the performance analysis of probabilistic timed automata. We extend previous results concerning the integer-time semantics of an important subclass of probabilistic timed automata to consider the computation of *expected costs* or *rewards*. We illustrate this approach through the analysis of the dynamic configuration protocol for IPv4 link-local addresses.

1 Introduction

Network protocols increasingly often rely on the use of randomness and timing delays, for example exponential back-off in Ethernet and IEEE 802.11, and IEEE 1394 FireWire root contention. Since these protocols execute in a distributed environment, it is important to also consider nondeterminism when modelling their behaviour. For example, we may wish to model a system for which the likelihood of a certain event occurring changes with respect to the

*Supported in part by the EPSRC grants GR/N22960 and GR/S11107.

amount of time elapsed. This notion is particularly important when considering timed fault-tolerant systems. A natural model for systems that exhibit nondeterminism, probability and real-time, called *probabilistic timed automata* – a probabilistic extension of timed automata [AD94] – has been proposed in [KNSS02]. In the probabilistic timed automata model real-valued clocks measure the passage of time and transitions can be probabilistic, that is, be expressed as a discrete probability distribution on the set of target states. In [KNSS02] model checking algorithms for verifying the likelihood of certain temporal properties being satisfied by such system models are introduced. However, these model checking algorithms are either based on *region equivalence* [AD94], and hence suffer from the state-space explosion problem, or on *forwards reachability*, which leads to approximate results [KNSS02, DKN02]. An alternative approach, based on *backwards reachability*, is given in [KNS01]; while this is more efficient than the region equivalence approach and leads to exact results, the approach has only been applied to calculating reachability probabilities.

When modelling real-time systems, there is often a trade-off between the expressiveness of the model and the complexity of the associated solution algorithms. A *dense-time* model is more expressive than an *integer-time* model; however, it is often the case that an integer-time model is easier to verify, since it can lead to a finite-state system and allows one to apply efficient symbolic methods developed for untimed systems. We refer to the clocks of an integer-time model as *digital clocks*. Henzinger et al. [HMP92] study the question of which real-time properties can be verified by considering system behaviours featuring only integer durations. These results are applied to timed automata in [Bos99, OW03], and it is shown that an approach using digital clocks is applicable to the verification of *closed, diagonal-free* timed automata; intuitively, these are automata whose constraints do not compare clocks or use strict comparison operators.

We have previously shown that probabilistic reachability properties, such as ‘with probability 0.05 or less, the system aborts’, of closed, diagonal-free probabilistic timed automata can be analysed faithfully using digital clocks [KNS03]. The main contribution of this paper is to extend this research by showing that digital clocks are also sufficient for verifying *expected reachability* properties such as ‘the expected time until a data packet is delivered is at most 0.05 seconds’, ‘the expected cost of a host choosing an IP address is at most 40’, or ‘the expected number of packets sent within the first 200 seconds is at least 400’.

In [dA97], de Alfaro presents a model-checking algorithm for verifying probabilistic and expected reachability properties of finite-state models. We implemented the algorithms of de Alfaro in the probabilistic model checking tool PRISM [KNP02, Pri], allowing us to automatically verify expected-cost properties of interest for integer-time models.

The paper proceeds by introducing preliminary concepts that we will use in the remainder of the paper. In Section 3 we introduce probabilistic timed automata and their semantics both for a dense model of time. Probabilistic and expected reachability properties for probabilistic timed automata are presented in Section 4, and the correctness of the digital clock interpretation of probabilis-

tic timed automata with respect to these properties is considered. In Section 5, we present a case study, in which PRISM is used to analyse the performance of the dynamic configuration protocol for IPv4 link-local addresses. Finally, in Section 6, we conclude the paper.

2 Preliminaries

2.1 Time, clocks and Zones

Let $\mathbb{T} \in \{\mathbb{R}, \mathbb{N}\}$ be the *time domain* of either the non-negative reals or naturals. Let \mathcal{X} be a finite set of variables called *clocks* which take values from the time domain \mathbb{T} . A point $v \in \mathbb{T}^{|\mathcal{X}|}$ is referred to as a *clock valuation*. Let $\mathbf{0} \in \mathbb{T}^{|\mathcal{X}|}$ be the clock valuation which assigns 0 to all clocks in \mathcal{X} . For any $v \in \mathbb{T}^{|\mathcal{X}|}$ and $t \in \mathbb{T}$, the clock valuation $v \oplus t$ denotes the *time increment* of values in v by t . We use $v[X := 0]$ to denote the clock valuation obtained from v by resetting all of the clocks in $X \subseteq \mathcal{X}$ to 0.

Let $Zones(\mathcal{X})$ be the set of *zones* over \mathcal{X} , which are conjunctions of atomic constraints of the form $x \sim c$ for $x \in \mathcal{X}$, $\sim \in \{\leq, =, \geq\}$, and $c \in \mathbb{N}$. The clock valuation v *satisfies* the zone ζ , written $v \models \zeta$, if and only if ζ resolves to true after substituting each clock $x \in \mathcal{X}$ with the corresponding clock value from v . Readers familiar with timed automata will note that we consider the syntax of *closed, diagonal-free zones*, which *do not* feature atomic constraints of the form $x > c$ or $x < c$ (closed) or $x - y \sim c$ (diagonal free) for $x, y \in \mathcal{X}$, $\sim \in \{\leq, =, \geq\}$ and $c \in \mathbb{N}$.

2.2 Probability and Measure Theory

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 possibly uncountable set Q' , let $\text{Dist}(Q')$ be the set of such distributions over countable subsets of Q' . For some element $q \in Q$, let the *point distribution* $\eta_q \in \text{Dist}(Q)$ be the distribution which assigns probability 1 to q .

We assume some familiarity with probability and measure theory, see e.g. [Bil79].

Definition 1 *Let Ω be an arbitrary non-empty set and \mathcal{F} a family of subsets of Ω . We say that \mathcal{F} is a field on Ω if:*

1. *the empty set \emptyset is in \mathcal{F} ;*
2. *whenever A is an element of \mathcal{F} , the complement of A is in \mathcal{F} ;*
3. *whenever A and B are elements of \mathcal{F} , then $A \cup B$ is in \mathcal{F} .*

A field of subsets \mathcal{F} is called a σ -field if it is field which is closed under countable union: whenever $A_i \in \mathcal{F}$ for $i \in \mathbb{N}$, then $\bigcup_{i \in \mathbb{N}} A_i$ is also in \mathcal{F} .

The elements of a σ -field are called *measurable sets*, and (Ω, \mathcal{F}) is called a *measurable space*.

Definition 2 Let (Ω, \mathcal{F}) be a measurable space. A function $\mu : \mathcal{F} \rightarrow [0, 1]$ is a probability measure on (Ω, \mathcal{F}) and $(\Omega, \mathcal{F}, \mu)$ a probability space, if μ satisfies the following properties:

1. $\mu(\Omega) = 1$
2. if A_1, A_2, \dots is a disjoint sequence of elements of \mathcal{F} , then $\mu(\cap_i A_i) = \sum_i \mu(A_i)$.

The measure μ is also referred to as a *probability distribution*. The set Ω is called the *sample space*, and the elements of \mathcal{F} *events*.

Definition 3 Let (Ω, \mathcal{F}) and (Ω', \mathcal{F}') be two measurable spaces. A function $f : \Omega \rightarrow \Omega'$ is said to be a measurable function from (Ω, \mathcal{F}) to (Ω', \mathcal{F}') if $f^{-1}(A') \in \mathcal{F}$ for all $A' \in \mathcal{F}'$.

Theorem 4 ([Bil79]) Let (Ω, \mathcal{F}) and (Ω', \mathcal{F}') be measurable spaces, and suppose that P is a measure on (Ω, \mathcal{F}) and the function $T : \Omega \rightarrow \Omega'$ is measurable. If f is a real non-negative measurable function on (Ω', \mathcal{F}') , then:

$$\int_{\omega \in \Omega} f(T\omega) dP = \int_{\omega' \in \Omega'} f(\omega') dPT^{-1}.$$

3 Probabilistic Timed Automata

In this section we review the definition of probabilistic timed automata [KNSS02] a modelling framework for real-time systems exhibiting both non-deterministic and stochastic behaviour. The formalism is derived from classical timed automata [ACD93, AD94] extended with discrete probability distributions over edges. An added feature is that of urgent events, which are a well-established concept for classical timed automata [HHWT95, DY95].

3.1 Syntax of probabilistic timed automata.

Definition 5 A probabilistic timed automaton is a tuple $(L, \bar{l}, \mathcal{X}, \Sigma, I, prob)$ where:

- L is a finite set of locations including the initial location \bar{l} ;
- \mathcal{X} is a set of clocks;
- Σ is a finite set of events, of which $\Sigma_u \subseteq \Sigma$ are declared as being urgent;
- the function $I : L \rightarrow Zones(\mathcal{X})$ is the invariant condition;
- the finite set $prob \subseteq L \times Zones(\mathcal{X}) \times \Sigma \times Dist(2^{\mathcal{X}} \times L)$ is the probabilistic edge relation.

A *state* of a probabilistic timed automaton is a pair (l, v) where $l \in L$ and $v \in \mathbb{T}^{|\mathcal{X}|}$ are such that $v \models I(l)$. Informally, the behaviour of a probabilistic timed automaton can be understood as follows. The model starts in the state $(\bar{l}, \mathbf{0})$; that is, in the initial location \bar{l} with all clocks set to 0. In any state (l, v) , there is a nondeterministic choice of either (1) making a *discrete transition* or (2) letting *time pass*. In case (1), a discrete transition can be made according to any $(l, g, \sigma, p) \in \text{prob}$ which is *enabled*; that is, the zone g is satisfied by the current clock valuation v . Then the probability of moving to the location l' and resetting all of the clocks in X to 0 is given by $p(X, l')$. In case (2), the option of letting time pass is available only if the invariant condition $I(l)$ is satisfied while time elapses and there does not exist an enabled probabilistic edge with an urgent event. Note that we often refer to the model presented above as *closed, diagonal-free probabilistic timed automata*, in order to distinguish the zones used with those in previous work [KNSS02].

3.2 Semantics of probabilistic timed automata.

The semantics of probabilistic timed automata is defined in terms of *timed probabilistic systems*, which exhibit timed, nondeterministic and probabilistic behaviour. They are a variant of Markov decision processes [Der70] and Segala's probabilistic timed automata [Seg95].

Definition 6 A timed probabilistic system $\text{PS} = (S, \bar{s}, \text{Act}, \mathbb{T}, \text{Steps})$ consists of a set S of states, an initial state $\bar{s} \in S$, a set Act of actions, a time domain \mathbb{T} , and a probabilistic transition relation $\text{Steps} \subseteq S \times (\text{Act} \cup \mathbb{T}) \times \text{Dist}(S)$, such that, if $(s, t, \mu) \in \text{Steps}$ for any $t \in \mathbb{T}$, then μ is a point distribution.

A *probabilistic transition* $s \xrightarrow{a, \mu} s'$ is made from a state $s \in S$ by first nondeterministically selecting an action-distribution or duration-distribution pair (a, μ) such that $(s, a, \mu) \in \text{Steps}$, and second by making a probabilistic choice of target state s' according to the distribution μ , such that $\mu(s') > 0$.

We consider two ways in which a timed probabilistic system's computation may be represented. A *path* represents a particular resolution of both nondeterminism *and* probability. Formally, a path of a timed probabilistic system is a finite or infinite sequence of probabilistic transitions

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

A path ω is *initialised in* s if $s_0 = s$. We denote by $\omega(i)$ the $(i + 1)$ th state of ω , $\text{last}(\omega)$ the last state of ω if ω is finite, and $\text{Steps}(\omega, i)$ the action associated with the i -th step. If ω is infinite, the duration up to the $(n + 1)$ th state of ω is defined by

$$\mathcal{D}_\omega(n + 1) \stackrel{\text{def}}{=} \sum \{a_i \mid 0 \leq i \leq n \wedge a_i \in \mathbb{T}\}.$$

Let $\text{Path}_{\text{fut}}(s)$ be the set of infinite paths initialised in s .

The second notion of a timed probabilistic system's computations is that of an *adversary*, which represents a particular resolution of nondeterminism *only*. Formally, an adversary is a function A mapping every finite path ω to a pair (a, μ) such that $(\text{last}(\omega), a, \mu) \in \text{Steps}$ [Var85]. For any adversary A ,

let $Path_{ful}^A(s)$ denote the set of infinite paths initialised in s associated with A . Then, we define the probability measure $Prob_s^A$ over $Path_{ful}^A(s)$ by classical techniques [KSK76].

We restrict our attention to *time-divergent adversaries*; a common restriction imposed in real-time systems so that unrealisable behaviour (corresponding to time not advancing beyond a bound) is disregarded during analysis. We say that a path ω is *divergent* if for any $t \in \mathbb{R}$, there exists $j \in \mathbb{N}$ such that $\mathcal{D}_\omega(j) > t$. It is easy to see that the set of divergent paths is measurable under any of the adversaries identified in the previous paragraph.

Definition 7 *An adversary A of a timed probabilistic system PS is divergent if and only if for each state s the probability $Prob_s^A$ of the divergent paths of $Path_{ful}^A(s)$ is 1. Furthermore, let Adv_{PS} be the set of divergent adversaries of PS.*

We now define the semantics of probabilistic timed automata defined in terms of timed probabilistic systems. Observe that the definition is parameterized both by a time domain \mathbb{T} and time increment \oplus , and that the summation in the definition of discrete transitions is required for the cases in which multiple clock resets result in the same target location.

Definition 8 *Let $PTA = (L, \bar{l}, \mathcal{X}, \Sigma, I, prob)$ be a probabilistic timed automaton. The semantics of PTA with respect to the time domain \mathbb{T} and the time increment \oplus is the timed probabilistic system $\llbracket PTA \rrbracket_{\mathbb{T}}^{\oplus} = (S, \bar{s}, \Sigma, \mathbb{T}, Steps)$ where: $S \subseteq L \times \mathbb{T}^{|\mathcal{X}|}$ and $(l, v) \in S$ if and only if $v \models I(l)$; $\bar{s} = (\bar{l}, \mathbf{0})$; and $((l, v), a, \mu) \in Steps$ if and only if one of the following conditions holds:*

Time transitions. $a \in \mathbb{T}$, $\mu = \eta_{(l, v \oplus a)}$ and

1. $v \oplus t \models I(l)$ for all $0 \leq t \leq a$;
2. for all probabilistic edges of the form $(l, g, a, -) \in prob$, if $v \models g$, then $\sigma \notin \Sigma_u$;

Discrete transitions. $a \in \Sigma$ and there exists $(l, g, \sigma, p) \in prob$ such that $v \models g$ and for any $(l', v') \in S$, we have

$$\mu(l', v') = \sum_{X \subseteq \mathcal{X} \text{ \& } v' = v[X:=0]} p(X, l').$$

Traditionally, the semantics of probabilistic timed automata assumes that the reals form the underlying model of time, paired with a time increment which is standard addition. The continuous semantics of a probabilistic timed automaton is a timed probabilistic system with generally uncountably many states.

Definition 9 *The continuous semantics of a probabilistic timed automaton PTA is defined as $\llbracket PTA \rrbracket_{\mathbb{R}}^+$; that is, $\mathbb{T} = \mathbb{R}$ and $\oplus = +$.*

The continuous semantics of the parallel composition of two probabilistic timed automata corresponds to the the parallel composition of their individual continuous semantics [KNS03].

3.3 Higher-level modelling.

To aid modelling, probabilistic timed automata can feature *integer variables*, *urgent locations*, and *committed locations* (all of which are standard features of UPPAAL timed automata [BDL⁺01]). The techniques of [Tri98] can be adapted to represent, syntactically, integer variables and committed locations within our definition of probabilistic timed automata.

It is often useful to define complex systems as the *parallel composition* of a number of interacting sub-components. The definition of the parallel composition operator \parallel uses ideas from the theory of (untimed) probabilistic systems [SL94] and classical timed automata [AD94]. Let $\text{PTA}_i = (L_i, \bar{l}_i, \mathcal{X}_i, \Sigma_i, I_i, \text{prob}_i)$ for $i \in \{1, 2\}$.

Definition 10 *The parallel composition of two probabilistic timed automata PTA_1 and PTA_2 is the probabilistic timed automaton $\text{PTA}_1 \parallel \text{PTA}_2 = (L_1 \times L_2, (\bar{l}_1, \bar{l}_2), \mathcal{X}_1 \cup \mathcal{X}_2, \Sigma_1 \cup \Sigma_2, I, \text{prob})$ where $I(l, l') = I_1(l) \wedge I_2(l')$ for all $(l, l') \in L_1 \times L_2$ and $((l_1, l_2), g, \sigma, p) \in \text{prob}$ if and only if one of the following conditions holds:*

- $\sigma \in \Sigma_1 \setminus \Sigma_2$ and there exists $(l_1, g, \sigma, p_1) \in \text{prob}_1$ such that $p = p_1 \otimes \eta_{(\emptyset, l_2)}$;
- $\sigma \in \Sigma_2 \setminus \Sigma_1$ and there exists $(l_2, g, \sigma, p_2) \in \text{prob}_2$ such that $p = \eta_{(\emptyset, l_1)} \otimes p_2$;
- $\sigma \in \Sigma_1 \cap \Sigma_2$ and there exists $(l_1, g_1, \sigma, p_1) \in \text{prob}_1$ and $(l_2, g_2, \sigma, p_2) \in \text{prob}_2$ such that $g = g_1 \wedge g_2$ and $p = p_1 \otimes p_2$

where for any $l_1 \in L_1$, $l_2 \in L_2$, $X_1 \subseteq \mathcal{X}_1$ and $X_2 \subseteq \mathcal{X}_2$:

$$p_1 \otimes p_2(X_1 \cup X_2, (l_1, l_2)) = p_1(X_1, l_1) \cdot p_2(X_2, l_2).$$

4 Performance Measures

In this section, we consider two performance measures for probabilistic timed automata. The first is *probabilistic reachability*, namely the maximal and minimal probability of reaching, from the initial state, a certain set of goal or target states. For a timed probabilistic system $\text{PS} = (S, \bar{s}, \text{Act}, \mathbb{T}, \text{Steps})$, set $F \subseteq S$ of target states, and adversary $A \in \text{Adv}_{\text{PS}}$, let:

$$p_{\bar{s}}^A(F) \stackrel{\text{def}}{=} \text{Prob}_{\bar{s}}^A \{ \omega \in \text{Path}_{\text{ful}}^A(\bar{s}) \mid \exists i \in \mathbb{N}. \omega(i) \in F \}.$$

Definition 11 *The maximal and minimal reachability probabilities of reaching the set of states F of the timed probabilistic system PS are defined as follows:*

$$p_{\text{PS}}^{\max}(F) = \sup_{A \in \text{Adv}_{\text{PS}}} p_{\bar{s}}^A(F) \quad \text{and} \quad p_{\text{PS}}^{\min}(F) = \inf_{A \in \text{Adv}_{\text{PS}}} p_{\bar{s}}^A(F).$$

This performance measure has been studied in the context of probabilistic timed automata by Kwiatkowska et al. [KNSS02, KNS03].

The second measure we consider is *expected reachability*, which allows us to compute the expected cost (or reward) accumulated before reaching a certain

set of states. Expected reachability is defined with respect to a cost function mapping actions and durations to real values, as well as a set $F \subseteq S$ of target states, and corresponds to the expected cost (with respect to the given cost function) of reaching a state in F . More formally, for a timed probabilistic system $\text{PS} = (S, \bar{s}, \text{Act}, \mathbb{T}, \text{Steps})$, cost function $c : \text{Act} \cup \mathbb{T} \rightarrow \mathbb{R}$, set $F \subseteq S$ of target states, and adversary $A \in \text{Adv}_{\text{PS}}$, let $e_{\bar{s}}^A(\text{cost}(c, F))$ denote the usual expectation with respect to the measure $\text{Prob}_{\bar{s}}^A$ over $\text{Path}_{\text{ful}}^A(\bar{s})$, where for any $\omega \in \text{Path}_{\text{ful}}^A(\bar{s})$:

$$\text{cost}(c, F)(\omega) = \begin{cases} \min\{j \mid \omega(j) \in F\} \sum_{i=1}^j c(\text{Steps}(\omega, i-1)) & \text{if } \exists j \in \mathbb{N}. \omega(j) \in F \\ \infty & \text{otherwise.} \end{cases}$$

The value of $\text{cost}(c, F)(\omega)$ equals the total cost, with respect to the cost function c , accumulated until a state in F is reached along the path ω . Note that we define the cost of a path which does not reach F to be ∞ , even though the total cost of the path may not be infinite. Hence, the expected cost of reaching F from s is finite if and only if a state in F is reached from s with probability 1. *Expected time reachability* (the expected time with which a given set of states can be reached) is a special case of expected reachability, corresponding to the case when $c(a) = 0$ for all $a \in \text{Act}$ and $c(t) = t$ for all $t \in \mathbb{T}$.

Definition 12 *The maximal and minimal expected costs of reaching a set of states F under the cost function c in the timed probabilistic system PS are defined as follows:*

$$e_{\text{PS}}^{\max}(c, F) = \sup_{A \in \text{Adv}_{\text{PS}}} e_{\bar{s}}^A(\text{cost}(c, F)) \quad \text{and} \quad e_{\text{PS}}^{\min}(c, F) = \inf_{A \in \text{Adv}_{\text{PS}}} e_{\bar{s}}^A(\text{cost}(c, F)).$$

We note that calculating expected reachability is equivalent to the *stochastic shortest path problem* for Markov decision processes; see for example [BT91].

At the level of probabilistic timed automata, one can define a cost function using a pair (r, c_{Σ}) , where $r \in \mathbb{R}$ gives the rate at which cost is accumulated as time passes, and $c_{\Sigma} : \Sigma \rightarrow \mathbb{R}$ is a function assigning the cost of executing each event in Σ . The associated cost function $c_{r, c_{\Sigma}}$ is defined by $c_{r, c_{\Sigma}}(t) = t \cdot r$ for all $t \in \mathbb{T}$, and $c_{r, c_{\Sigma}}(\sigma) = c_{\Sigma}(\sigma)$ for all $\sigma \in \Sigma$. A probabilistic timed automaton equipped with a pair (r, c_{Σ}) is a probabilistic generalisation of uniformly priced timed automata [BFH⁺01a].

For both probabilistic and expected reachability, we can consider reaching a state satisfying a formula which is a conjunction of propositions identifying locations and clock constraints of the form $x \sim c$ for $x \in \mathcal{X}$, $\sim \in \{\leq, =, \geq\}$ and $c \in \mathbb{N}$. Instead of considering these cases separately, we just note that such reachability problems can be reduced to those referring to locations only by modifying syntactically the probabilistic timed automaton of interest (see [KNSS02]). Note that, if the PTA and clock constraints present in the formula are diagonal-free and closed then the modified PTA is also diagonal-free and closed.

In the case of probabilistic reachability the types of properties which can be expressed can be classified as follows:

Reachability The system can reach a certain set of states with a given maximal or minimal probability. For example, ‘with probability at least 0.9999, a data packet is correctly delivered.’

Time bounded reachability The system can reach a certain set of states within a certain time deadline and probability threshold. For example, ‘with probability 0.01 or less, a data packet is lost within 5 time units.’

Cost bounded reachability The system can reach a certain set of states within a certain cost and probability bound. For example, ‘with probability 0.75 or greater, a data packet is correctly delivered with at most 4 retransmissions.’

Invariance The system does not leave a certain set of states with a given probability. For example, ‘with probability 0.875 or greater, the system never aborts.’

Bounded response The system inevitably reaches a certain set of states within a certain time deadline with a given probability. For example, ‘with probability 0.99 or greater, a data packet will always be delivered within 5 time units.’

On the other hand, expected time reachability allows us to express for example, ‘the expected time until a host can use an IP address is at most 0.05 seconds’ and ‘the expected time until a packet collision occurs is at least 100 seconds’. In general, expected reachability allows us to validate properties including: ‘the expected number of retransmissions before the message is correctly delivered is sent is less than 3’, ‘the expected number of packets sent before failure is at least 300’ and ‘the expected number of lost messages within the first 200 seconds is at most 10.’

To illustrate the expected reachability approach, in the case of the final property, we would first need to modify the probabilistic timed automaton under study by adding a distinct clock (to represent global time) and a location such that, from all locations, once the global clock has reached 200 seconds, the only transition is to this new location. The set of target states would then be the set containing only the new location and the cost function would equal 0 on all time transitions and events except those events corresponding to a message being lost; the costs for those actions would be set to 1.

4.1 Performance measures and digital clocks.

We now show, under the restriction that the probabilistic timed automaton under study is diagonal-free and closed, that it suffices just to consider the integer-time semantics when verifying expected reachability properties.

Definition 13 For any $x \in \mathcal{X}$, let \mathbf{k}_x denote the greatest constant that the clock x is compared to in the zones of PTA. Define $\oplus_{\mathbb{N}}$ such that, for any clock valuation $v \in \mathbb{N}^{|\mathcal{X}|}$ and time duration $t \in \mathbb{N}$, the clock valuation $v \oplus_{\mathbb{N}} t$ assigns the value $\min\{v_x + t, \mathbf{k}_x + 1\}$ to all clocks $x \in \mathcal{X}$. The integer-time semantics of PTA is then defined as $\llbracket \text{PTA} \rrbracket_{\mathbb{N}}^{\oplus_{\mathbb{N}}}$; that is, $\mathbb{T} = \mathbb{N}$ and $\oplus = \oplus_{\mathbb{N}}$.

Note that, as we always use the same type of time increment for a particular choice of time domain, we omit the $+$ and $\oplus_{\mathbb{N}}$ superscripts from the notation. The fact that the integral semantics of a probabilistic timed automaton is finite-state can be derived from the definitions. The integer semantics of the parallel composition of two probabilistic timed automata corresponds to the parallel composition of their integer semantics [KNS03].

Let $\text{PTA} = (L, \bar{l}, \mathcal{X}, \Sigma, I, \text{prob})$ be a (closed, diagonal-free) probabilistic timed automaton. For any set of locations $L' \subseteq L$, we denote by $F_{\mathbb{T}}^{L'}$ the set of all states of $\llbracket \text{PTA} \rrbracket_{\mathbb{T}}$ which correspond to these locations; that is $F_{\mathbb{T}}^{L'} = \{(l, v) \mid l \in L', v \in \mathbb{T}^{|\mathcal{X}|} \wedge v \models I(l)\}$.

Theorem 14 *For any (closed, diagonal-free) probabilistic timed automaton PTA, set of locations $L' \subseteq L$ and cost function $c : \Sigma \cup \mathbb{R} \rightarrow \mathbb{R}$ which satisfies $c(t + t') = c(t) + c(t')$ for all $t, t' \in \mathbb{R}$:*

$$e_{\llbracket \text{PTA} \rrbracket_{\mathbb{R}}}^{\max}(c, F_{\mathbb{R}}^{L'}) = e_{\llbracket \text{PTA} \rrbracket_{\mathbb{N}}}^{\max}(c, F_{\mathbb{N}}^{L'}) \quad \text{and} \quad e_{\llbracket \text{PTA} \rrbracket_{\mathbb{R}}}^{\min}(c, F_{\mathbb{R}}^{L'}) = e_{\llbracket \text{PTA} \rrbracket_{\mathbb{N}}}^{\min}(c, F_{\mathbb{N}}^{L'}).$$

Note that any cost functions defined by a pair (r, c_{Σ}) , where $r \in \mathbb{R}$ and $c_{\Sigma} : \Sigma \rightarrow \mathbb{R}$, will satisfy the condition $c(t + t') = c(t) + c(t')$ for all $t, t' \in \mathbb{R}$. The analogous result for probabilistic reachability (initially proved in [KNS03]) is as follows.

Theorem 15 *For any (closed, diagonal-free) probabilistic timed automaton PTA, set of locations $L' \subseteq L$:*

$$p_{\llbracket \text{PTA} \rrbracket_{\mathbb{R}}}^{\max}(F_{\mathbb{R}}^{L'}) = p_{\llbracket \text{PTA} \rrbracket_{\mathbb{N}}}^{\max}(F_{\mathbb{N}}^{L'}) \quad \text{and} \quad p_{\llbracket \text{PTA} \rrbracket_{\mathbb{R}}}^{\min}(F_{\mathbb{R}}^{L'}) = p_{\llbracket \text{PTA} \rrbracket_{\mathbb{N}}}^{\min}(F_{\mathbb{N}}^{L'}).$$

We now prove the above theorems based on the techniques developed in the classical timed automata case. Note that, an alternative proof of Theorem 15 appears in [KNS03]. First consider the following definition and lemma which are taken from [HMP92, Hen91].

Definition 16 *For any $t \in \mathbb{R}$ and $\varepsilon \in [0, 1]$ let:*

$$[t]_{\varepsilon} = \begin{cases} \lfloor t \rfloor & \text{if } t \leq \lfloor t \rfloor + \varepsilon \\ \lceil t \rceil & \text{otherwise.} \end{cases}$$

Note that, from Definition 16 it trivially follows that:

$$[t]_1 \leq t \leq [t]_0 \quad \text{for all } t \in \mathbb{R}. \quad (1)$$

Lemma 17 *For any $t, t' \in \mathbb{R}$, $c \in \mathbb{N}$ and $\sim \in \{\leq, =, \geq\}$, if $t - t' \sim c$ then $[t]_{\varepsilon} - [t']_{\varepsilon} \sim c$ for all $\varepsilon \in [0, 1]$.*

Proof. See for example [Hen91]. □

Next, we introduce the following property on paths of the PTA.

Lemma 18 For any path $\omega = (l_0, v_0) \xrightarrow{a_0, \mu_0} (l_1, v_1) \xrightarrow{a_1, \mu_1} \dots$, $x \in \mathcal{X}$ and $i \in \mathbb{N}$, there exists $j \leq i$ such that $v_i(x) = \mathcal{D}_\omega(i) - \mathcal{D}_\omega(j)$.

Proof. The proof follows from choosing $j \leq i$ such that s_j is the most recent state where the clock x was reset. \square

Using the above we now define the ε -digitization of a path [HMP92, Hen91].

Definition 19 (ε -digitization) For any path:

$$\omega = (\bar{l}, \mathbf{0}) \xrightarrow{a_0, \mu_0} (l_1, v_1) \xrightarrow{a_1, \mu_1} \dots$$

of $\text{PTA}_{\mathbb{R}}$, its ε -digitization is the path

$$[\omega]_\varepsilon = (\bar{l}, \mathbf{0}) \xrightarrow{a'_0, \mu_0} (l_1, [v_1]_\varepsilon) \xrightarrow{a'_1, \mu_1} \dots$$

of $\text{PTA}_{\mathbb{N}}$ where for any $i \in \mathbb{N}$ and $x \in \mathcal{X}$:

- $[v_i]_\varepsilon(x) = \min([\mathcal{D}_\omega(i)]_\varepsilon - [\mathcal{D}_\omega(j)]_\varepsilon, \mathbf{k}_x + 1)$ and $j \leq i$ such that $v_i(x) = \mathcal{D}_\omega(i) - \mathcal{D}_\omega(j)$, which exists by Lemma 18;
- if $a_i \in \Sigma$, then $a'_i = a_i$;
- if $a_i \in \mathbb{R}$, then $a'_i = [\mathcal{D}_\omega(i+1)]_\varepsilon - [\mathcal{D}_\omega(i)]_\varepsilon$.

The well-definedness of this construction, that is, the fact that $[\omega]_\varepsilon$ is a path of $\text{PTA}_{\mathbb{N}}$, follows from Lemma 17, Lemma 18, and the fact that PTA is diagonal free and closed. For example, for any $x \in \mathcal{X}$ and $i \geq 0$ where $a_i \in \mathbb{R}$, by Definition 19 there exists $j \leq i$ such that:

$$\begin{aligned} \min([v_i]_\varepsilon(x) + a'_i, \mathbf{k}_x) &= \min([\mathcal{D}_\omega(i)]_\varepsilon - [\mathcal{D}_\omega(j)]_\varepsilon + a'_i, \mathbf{k}_x) \\ &= \min([\mathcal{D}_\omega(i)]_\varepsilon - [\mathcal{D}_\omega(j)]_\varepsilon + [\mathcal{D}_\omega(i+1)]_\varepsilon - [\mathcal{D}_\omega(i)]_\varepsilon, \mathbf{k}_x) && \text{by construction} \\ &= \min([\mathcal{D}_\omega(i+1)]_\varepsilon - [\mathcal{D}_\omega(j)]_\varepsilon, \mathbf{k}_x) && \text{rearranging} \\ &= [v_{i+1}]_\varepsilon(x) && \text{by Definition 19.} \end{aligned}$$

We now introduce the following lemmas which relate the time and cost of a path with that of its digitization.

Lemma 20 For any path $\omega \in \text{Path}_{ful}(\bar{s})$, $\varepsilon \in [0, 1]$ and $i \in \mathbb{N}$: $\mathcal{D}_{[\omega]_\varepsilon}(i) = [\mathcal{D}_\omega(i)]_\varepsilon$.

Proof. Consider any path $\omega = (\bar{l}, \mathbf{0}) \xrightarrow{a_0, \mu_0} (l_1, v_1) \xrightarrow{a_1, \mu_1} \dots \in \text{Path}_{ful}(\bar{s})$. We prove the lemma by induction on $i \in \mathbb{N}$. If $i = 0$, then by definition $\mathcal{D}_{[\omega]_\varepsilon}(i) = [\mathcal{D}_\omega(i)]_\varepsilon = 0$ as required.

Now suppose that the lemma holds for some $i \in \mathbb{N}$. We have two cases to consider.

- If $a_i \in \Sigma$, then $[\mathcal{D}_\omega(i+1)]_\varepsilon = [\mathcal{D}_\omega(i)]_\varepsilon$ and from Definition 19 we have $\mathcal{D}_{[\omega]_\varepsilon}(i+1) = \mathcal{D}_{[\omega]_\varepsilon}(i)$, and hence by induction the lemma holds in this case.

- If $a_i \in \mathbb{R}$, then by Definition 19 we have

$$\begin{aligned}
\mathcal{D}_{[\omega]_\varepsilon}(i+1) &= \mathcal{D}_{[\omega]_\varepsilon}(i) + ([\mathcal{D}_\omega(i+1)]_\varepsilon - [\mathcal{D}_\omega(i)]_\varepsilon) \\
&= [\mathcal{D}_\omega(i)]_\varepsilon + ([\mathcal{D}_\omega(i+1)]_\varepsilon - [\mathcal{D}_\omega(i)]_\varepsilon) && \text{by induction} \\
&= [\mathcal{D}_\omega(i+1)]_\varepsilon && \text{as required.}
\end{aligned}$$

Since, these are the only cases to consider, the lemma holds by induction on $i \in \mathbb{N}$. \square

Lemma 21 *For any path $\omega \in Path_{ful}(\bar{s})$, set of locations $L' \subseteq L$ and non-negative cost function $c : Act \cup \mathbb{R} \rightarrow \mathbb{R}$ such that $c(t+t') = c(t) + c(t')$ for all $t, t' \in \mathbb{R}$, we have:*

$$cost(c, F_{\mathbb{N}}^{L'})([\omega]_1) \leq cost(c, F_{\mathbb{R}}^{L'}) (\omega) \leq cost(c, F_{\mathbb{N}}^{L'})([\omega]_0).$$

Proof. Consider any path $\omega \in Path_{ful}(\bar{s})$, set of locations $L' \subseteq L$ and cost function $c : Act \cup \mathbb{T} \rightarrow \mathbb{R}$ such that $c(t+t') = c(t) + c(t')$ for all $t, t' \in \mathbb{R}$. If $\omega(i) \notin F_{\mathbb{R}}^{L'}$ for all $i \geq 0$, then by Definition 19 we have $[\omega]_\varepsilon(i) \notin F_{\mathbb{N}}^{L'}$ for all $i \geq 0$ and $\varepsilon \in [0, 1]$, and hence

$$cost(c, F_{\mathbb{N}}^{L'})([\omega]_1) = cost(c, F_{\mathbb{R}}^{L'}) (\omega) = cost(c, F_{\mathbb{N}}^{L'})([\omega]_0) = \infty$$

as required.

Therefore we are let to consider the case when $\omega(i) \in F_{\mathbb{R}}^{L'}$ for some $i \geq 0$. Now, let $c_{\uparrow Act} : Act \cup \mathbb{R} \rightarrow \mathbb{R}$ be the cost function defined as follows: $c_{\uparrow Act}(a) = c(a)$ for all $a \in Act$ and $c_{\uparrow Act}(t) = 0$ for all $t \in \mathbb{R}$. First, by definition of $cost(c, F)$ and of \mathcal{D} , and since $c(t+t') = c(t) + c(t')$ for all $t, t' \in \mathbb{R}$, it follows that:

$$cost(c, F_{\mathbb{R}}^{L'}) (\omega) = cost(c_{\uparrow Act}, F_{\mathbb{R}}^{L'}) (\omega) + c(\mathcal{D}_\omega(i_{\min})) \quad (2)$$

where $i_{\min} = \min\{i \mid \omega(i) \in F_{\mathbb{R}}^{L'}\}$. Next, from Definition 19 we have for any $\varepsilon \in [0, 1]$, $i \in \mathbb{N}$ and $a \in Act$:

- $\omega(i) \in F_{\mathbb{R}}^{L'}$ if and only if $[\omega]_\varepsilon(i) \in F_{\mathbb{N}}^{L'}$;
- $Steps(\omega, i) = a$ if and only if $Steps([\omega]_\varepsilon, i) = a$.

It then follows that:

$$cost(c, F_{\mathbb{N}}^{L'})([\omega]_\varepsilon) = cost(c_{\uparrow Act}, F_{\mathbb{N}}^{L'})([\omega]_\varepsilon) + c(\mathcal{D}_{[\omega]_\varepsilon}(i_{\min})) \quad \text{for all } \varepsilon \in [0, 1] \quad (3)$$

and

$$cost(c_{\uparrow Act}, F_{\mathbb{N}}^{L'})([\omega]_\varepsilon) = cost(c_{\uparrow Act}, F_{\mathbb{R}}^{L'}) (\omega) \quad \text{for all } \varepsilon \in [0, 1]. \quad (4)$$

Furthermore, using Lemma 20 and (1) we have:

$$\mathcal{D}_{[\omega]_1}(i_{\min}) \leq \mathcal{D}_\omega(i_{\min}) \leq \mathcal{D}_{[\omega]_0}(i_{\min}) \quad . \quad (5)$$

Finally, since c is non-negative and $c(t + t') = c(t) + c(t')$ for all $t, t' \in \mathbb{R}$, it follows that $c(t) \leq c(t')$ for all $t \leq t' \in \mathbb{R}$, and combining this fact with (2), (3), (4) and (5) we have

$$\text{cost}(c, F_{\mathbb{N}}^{L'})([\omega]_1) \leq \text{cost}(c, F_{\mathbb{R}}^{L'})(\omega) \leq \text{cost}(c, F_{\mathbb{N}}^{L'})([\omega]_0)$$

as required. \square

Unlike in the non-probabilistic setting, we must now extend the notion of digitization from paths to adversaries. In the following proofs, we use *randomized adversaries*, where a randomized adversary B is a function B mapping every finite path ω to a distribution over $\{(a, \mu) \mid (last(\pi), a, \mu) \in Steps\}$. Similarly to the case of ‘non-randomized’ adversaries presented in Section 3, we can associate with any randomized adversary a probability measure over the set of paths of the adversary.

To simplify the presentation, when considering a fixed adversary $A \in Adv_{\llbracket PTA \rrbracket_{\mathbb{R}}}$, we suppose that the domain of the mapping $[\cdot]_{\varepsilon}$ is restricted to the set of paths $Path_{ful}^A(\bar{s})$. Using this interpretation we now extend the notion of digitization to adversaries through the following proposition.

Proposition 22 *For any adversary $A \in Adv_{\llbracket PTA \rrbracket_{\mathbb{R}}}$ and $\varepsilon \in [0, 1]$, there exists a (randomized) adversary $B^{\varepsilon} \in Adv_{\llbracket PTA \rrbracket_{\mathbb{N}}}$ such that: $Prob^{B^{\varepsilon}}(\Pi) = Prob^A([\Pi]_{\varepsilon}^{-1})$ for all $\Pi \in \mathcal{F}_{Path_{ful}^{B^{\varepsilon}}(\bar{s})}$.*

Proof. Consider any adversary A of $PTA_{\mathbb{R}}$. First, to ease notation, for any set of finite paths $\Omega \in Path_{fin}^A(\bar{s})$ let:

$$Prob^A(\Omega) = Prob^A\{\omega' \in Path_{ful}^A(\bar{s}) \mid \omega \leq \omega' \text{ for some } \omega \in \Omega\}$$

and for any path π and $\mu \in Steps(last(\pi))$:

$$[\pi \xrightarrow{\mu}]_{\varepsilon}^{-1} = \{\omega \in Path_{fin}^A(\bar{s}) \mid \exists s \in S. [\omega]_{\varepsilon} = \pi \xrightarrow{\mu} s\}.$$

We now define the adversary B^{ε} as follows. The set of paths of B^{ε} is given by $\{[\omega]_{\varepsilon} \mid \omega \in Path_{ful}^A(\bar{s})\}$ and for any $\pi \in Path_{fin}^{B^{\varepsilon}}(\bar{s})$ and $\mu \in Steps(last(\pi))$, the probability of choosing μ after π has been performed is given by:

$$B^{\varepsilon}(\pi)(\mu) \stackrel{\text{def}}{=} \frac{Prob^A([\pi \xrightarrow{\mu}]_{\varepsilon}^{-1})}{Prob^A([\pi]_{\varepsilon}^{-1})}.$$

We now prove that $Prob^{B^{\varepsilon}}(\Pi) = Prob^A([\Pi]_{\varepsilon}^{-1})$ for all $\Pi \in \mathcal{F}_{Path_{ful}^{B^{\varepsilon}}(\bar{s})}$. From the cone construction [KSK76] and the construction of the adversary B^{ε} above, it follows that it is sufficient to show that:

$$Prob^{B^{\varepsilon}}(\pi) = Prob^A([\pi]_{\varepsilon}^{-1}) \text{ for all } \pi \in Path_{fin}^{B^{\varepsilon}}(\bar{s}) \quad (6)$$

which we now prove by induction on the length of π . Therefore, consider any path $\pi \in Path_{fin}^{B^{\varepsilon}}(\bar{s})$, if $|\pi| = 0$ then $Prob^{B^{\varepsilon}}(\pi) = 1 = Prob^A([\pi]_{\varepsilon}^{-1})$ as required.

Next, suppose by induction the lemma holds for all paths of length n and π is of length $n + 1$, then π is of the form $\pi' \xrightarrow{\mu} s'$ for some path π' of length n , $\mu \in \text{Steps}(\text{last}(\pi'))$ and $s' \in S$. Therefore, in this case we have:

$$\begin{aligned}
\text{Prob}^{B^\varepsilon}(\pi) &= \text{Prob}^{B^\varepsilon}(\pi') \cdot B^\varepsilon(\pi')(\mu) \cdot \mu(s') \\
&= \text{Prob}^A([\pi']_\varepsilon^{-1}) \cdot B^\varepsilon(\pi')(\mu) \cdot \mu(s') && \text{by induction} \\
&= \text{Prob}^A([\pi']_\varepsilon^{-1}) \cdot \frac{\text{Prob}^A([\pi' \xrightarrow{\mu}]_\varepsilon^{-1})}{\text{Prob}^A([\pi']_\varepsilon^{-1})} \cdot \mu(s') && \text{by definition of } B^\varepsilon \\
&= \text{Prob}^A([\pi' \xrightarrow{\mu}]_\varepsilon^{-1}) \cdot \mu(s') && \text{rearranging} \\
&= \text{Prob}^A\{\omega \in \text{Path}_{fn}^A(\bar{s}) \mid [\omega]_\varepsilon = \pi' \xrightarrow{\mu} s'\} && \text{by definition of } \text{Prob}^A \\
&= \text{Prob}^A([\pi]_\varepsilon^{-1}) && \text{by construction of } \pi
\end{aligned}$$

and hence (6) holds by induction.

Finally, to show that $B^\varepsilon \in \text{Adv}_{\llbracket \text{PTA} \rrbracket_{\mathbb{N}}}$, we must show that B^ε is divergent. This result follows from (6), the fact that A is divergent, and since from Lemma 20 any (infinite) path is divergent if and only if its ε -digitization is divergent. \square

We are now in a position to prove Theorem 15, namely that it is sufficient to consider the integral semantics when considering probabilistic reachability properties.

Proof of Theorem 15. Consider any set of locations $L' \subseteq L$, since $B^\varepsilon \in \text{Adv}_{\llbracket \text{PTA} \rrbracket_{\mathbb{N}}}$ for any $A \in \text{Adv}_{\llbracket \text{PTA} \rrbracket_{\mathbb{R}}}$ and $\omega(i) \in F_{\mathbb{R}}^{L'}$ if and only if $[\omega]_\varepsilon(i) \in F_{\mathbb{N}}^{L'}$ for any $\omega \in \text{Path}_{ful}^A(\bar{s})$, from Proposition 22 it follows that:

$$\inf_{A \in \text{Adv}_{\llbracket \text{PTA} \rrbracket_{\mathbb{R}}}} p^A(F_{\mathbb{R}}^{L'}) \geq \inf_{B \in \text{Adv}_{\llbracket \text{PTA} \rrbracket_{\mathbb{N}}}} p^B(F_{\mathbb{N}}^{L'})$$

and

$$\sup_{A \in \text{Adv}_{\llbracket \text{PTA} \rrbracket_{\mathbb{R}}}} p^A(F_{\mathbb{R}}^{L'}) \leq \sup_{B \in \text{Adv}_{\llbracket \text{PTA} \rrbracket_{\mathbb{N}}}} p^B(F_{\mathbb{N}}^{L'}).$$

On the other hand by definition of the continuous and integral semantics for any adversary $B \in \text{Adv}_{\llbracket \text{PTA} \rrbracket_{\mathbb{N}}}$, there exists an adversary $A \in \text{Adv}_{\llbracket \text{PTA} \rrbracket_{\mathbb{R}}}$ such that:

$$p^A(F_{\mathbb{R}}^{L'}) = p^B(F_{\mathbb{N}}^{L'}).$$

Intuitively, the adversary A behaves like the integral semantics adversary B ; that is, chooses to make timed transitions of only integer duration. The result then follows since this was for an arbitrary integral semantics adversary. \square

Before we consider the proof of Theorem 14 we require the following lemma and proposition.

Lemma 23 *For any adversary $A \in \text{Adv}_{\llbracket \text{PTA} \rrbracket_{\mathbb{R}}}$, the mapping $[\cdot]_\varepsilon$ is a measurable function from $(\text{Path}_{ful}^A(\bar{s}), \mathcal{F}_{\text{Path}_{ful}^A(\bar{s})})$ to the measurable space induced from the set of paths $\{[\omega]_\varepsilon \mid \omega \in \text{Path}_{ful}^A(\bar{s})\}$.*

Proof. The proof follows from the fact that for any finite path π in the set $\{[\omega]_\varepsilon \mid \omega \in \text{Path}_{fin}^A(\bar{s})\}$:

$$[\text{Cone}(\pi)]_\varepsilon^{-1} = \bigcup_{\omega \in \text{Path}_{fin}^A(\bar{s}) \wedge [\omega]_\varepsilon = \pi} \text{Cone}(\omega),$$

that is, the set $[\text{Cone}(\pi)]_\varepsilon^{-1}$ is the finite union of measurable sets of $\text{Path}_{ful}^A(\bar{s})$, and hence is a measurable set. \square

For the following proposition, we use the construction of the adversary B^ε , for some $\varepsilon \in [0, 1]$, as defined in the proof of Proposition 22.

Proposition 24 *For any set of locations $L' \subseteq L$, adversary $A \in \text{Adv}_{[\text{PTA}]_{\mathbb{R}}}$ and non-negative cost function $c : \text{Act} \cup \mathbb{R} \rightarrow \mathbb{R}$ such that $c(t + t') = c(t) + c(t')$ for all $t, t' \in \mathbb{R}$:*

$$e_s^{B^1}(\text{cost}(c, F_{\mathbb{N}}^{L'})) \leq e_s^A(\text{cost}(c, F_{\mathbb{R}}^{L'})) \leq e_s^{B^0}(\text{cost}(c, F_{\mathbb{N}}^{L'})).$$

Proof. Consider any adversary $A \in \text{Adv}_{[\text{PTA}]_{\mathbb{R}}}$, and non-negative cost function $c : \text{Act} \cup \mathbb{R} \rightarrow \mathbb{R}$ such that $c(t + t') = c(t) + c(t')$ for all $t, t' \in \mathbb{R}$. First, for any $\varepsilon \in [0, 1]$, since c is non-negative, $\text{cost}(c, F_{\mathbb{N}}^{L'})$ is a real non-negative function on $(\text{Path}_{ful}^{B^\varepsilon}(\bar{s}), \mathcal{F}_{\text{Path}_{ful}^{B^\varepsilon}(\bar{s})})$, and hence using Proposition 22 and Lemma 23 we can apply Theorem 4 to give:

$$\int_{\omega \in \text{Path}_{ful}^A(\bar{s})} \text{cost}(c, F_{\mathbb{R}}^{L'})([\omega]_\varepsilon) d\text{Prob}^A = \int_{\pi \in \text{Path}_{ful}^{B^\varepsilon}(\bar{s})} \text{cost}(c, F_{\mathbb{N}}^{L'}) (\pi) d\text{Prob}^{B^\varepsilon}. \quad (7)$$

Now by definition of e_s^A :

$$\begin{aligned} e_s^A(\text{cost}(c, F_{\mathbb{R}}^{L'})) &= \int_{\omega \in \text{Path}_{ful}^A(\bar{s})} \text{cost}(c, F_{\mathbb{R}}^{L'}) (\omega) d\text{Prob}^A \\ &\leq \int_{\omega \in \text{Path}_{ful}^A(\bar{s})} \text{cost}(c, F_{\mathbb{N}}^{L'})([\omega]_0) d\text{Prob}^A \quad \text{by Lemma 21} \\ &= \int_{\pi \in \text{Path}_{ful}^{B^0}(\bar{s})} \text{cost}(c, F_{\mathbb{N}}^{L'}) (\pi) d\text{Prob}^{B^0} \quad \text{by (7)} \\ &= e_s^{B^0}(\text{cost}(c, F_{\mathbb{N}}^{L'})) \quad \text{by definition.} \end{aligned}$$

Similarly, we can show $e_s^{B^1}(\text{cost}(c, F_{\mathbb{N}}^{L'})) \leq e_s^A(\text{cost}(c, F_{\mathbb{R}}^{L'}))$ as required. \square

We are now in a position to prove Theorem 14.

Proof of Theorem 14. Consider any set of locations $L' \subseteq L$ and non-negative cost function $c : \text{Act} \cup \mathbb{R} \rightarrow \mathbb{R}$ such that $c(t + t') = c(t) + c(t')$ for all $t, t' \in \mathbb{R}$. First, since $B^\varepsilon \in \text{Adv}_{[\text{PTA}]_{\mathbb{N}}}$ for any $A \in \text{Adv}_{[\text{PTA}]_{\mathbb{R}}}$, from Proposition 24 it follows that:

$$\inf_{A \in \text{Adv}_{[\text{PTA}]_{\mathbb{R}}}} e_s^A(\text{cost}(c, F_{\mathbb{R}}^{L'})) \geq \inf_{B \in \text{Adv}_{[\text{PTA}]_{\mathbb{N}}}} e_s^B(\text{cost}(c, F_{\mathbb{N}}^{L'}))$$

and

$$\sup_{A \in Adv_{[PTA]_{\mathbb{R}}}} e_s^A(cost(c, F_{\mathbb{R}}^{L'})) \leq \sup_{B \in Adv_{[PTA]_{\mathbb{N}}}} e_s^B(cost(c, F_{\mathbb{N}}^{L'})) .$$

On the other hand, as in the proof of Theorem 15, for any adversary $B \in Adv_{[PTA]_{\mathbb{N}}}$, there exists an adversary $A \in Adv_{[PTA]_{\mathbb{R}}}$ such that $e_s^A(cost(c, F_{\mathbb{R}}^{L'})) = e_s^B(cost(c, F_{\mathbb{N}}^{L'}))$. \square

5 Case study: Dynamic configuration of link-local addresses in IPv4

In this section, we illustrate the utility of the integer-time semantics of probabilistic timed automata with an analysis of the dynamic configuration protocol for IPv4 link-local addresses [CAG].

The dynamic configuration protocol for IPv4 addresses offers a distributed ‘plug-and-play’ solution in which IP address configuration is managed by individual devices connected to a local network. Upon connecting to the network, a device, henceforth called a *host*, first randomly chooses an IP address from a pool of 65024 available (the Internet Assigned Number Authority has allocated the addresses from 169.254.1.0 to 169.254.254.255 for the purpose of such link-local networks). The host waits a random time of between 0 and 2 seconds before sending four *Address Resolution Protocol* (ARP) packets, called *probes*, to all of the other hosts of the network. Probes contain the IP address selected by the host, operate as requests to use the address, and are sent at 2 second intervals. A host which is already using the address will respond with an ARP reply packet, asserting its claim to the address, and the original host will restart the protocol by reconfiguring its chosen address and sending new probes. If the host sends four probes without receiving an ARP reply packet, then it commences to use the chosen IP address. The host then sends confirmations of this fact to the other hosts of the network by means of two *gratuitous* ARPs, also at 2 second intervals. The protocol has an inherent degree of redundancy, for example with regard to the number of repeated ARP packets sent, in order to cope with message loss. Indeed, message loss makes possible the undesirable situation in which two or more hosts use the same IP address simultaneously.

A host which has commenced using an IP address must reply to ARP packets containing the same IP addresses that it receives from other hosts. It continues using the address unless it receives any ARP packet other than a probe (for example, a gratuitous ARP) containing the IP address that it is using currently. In such a case, the host can either *defend* its IP address, or *defer* to the host which sent the conflicting ARP packet. The host may only defend its address if it has not received a previous conflicting ARP packet within the previous ten seconds; otherwise it is forced to defer. A defending host replies by the sending an ARP packet, thereby indicating that it is using the IP address. A deferring host does not send a reply; instead, it ceases using its current IP address, and reconfigures its IP address by restarting the protocol.

As in [ZV03], we assume a ‘broadcast’-based communication medium with no routers (for example, a single wire), in which messages arrive in the order

in which they are sent. In contrast to the analytic analysis of the protocol of Bohnenkamp et al. [BvdSHV03], we model the possibility that a device could surrender an IP address that it is using to another host; and in contrast to timed-automata-based analysis of Zhang and Vaandrager [ZV03], we model some important probabilistic characteristics of the protocol, and consider parameters more faithful to the standard (such as the maximum number of times a device can witness an ARP packet with the same IP address as that which it wishes to use before ‘backing off’ and remaining idle for at least one minute).

In the standard [CAG], there is no mention of what a host should do with messages corresponding to its current IP address (i.e. the probes and gratuitous ARP packets specified in the standard) which are in its output buffer (i.e. those that have yet to be sent), when it reconfigures (chooses a new IP address). However, when the host does reconfigure, unless it picks the same IP address, which happens with the very small probability $1/65024$, these messages are not relevant. In fact, such messages will slow down the network and may even make hosts reconfigure when they do not need to. We therefore considered two different versions of the protocol: one where the host does not do anything about these messages (`no_reset`) and another where the host clears its buffer (removes the messages) when it is about to choose a new IP address (`reset`).

5.1 Modelling the dynamic configuration protocol

We consider in detail one *concrete host*, which is attempting to configure an IP address for a network in which there are N *abstract hosts* (they are called abstract because we do not study their behaviour in depth) which have already configured IP addresses. Therefore, when the concrete host picks an address, the probability of this address being *fresh* (not in use by an abstract host) is $(65024 - N)/65024$. We also assume that the concrete host never picks the same IP address twice, as this happens only with a very small probability. Also, the (continuous) uniform choice over $[0,2]$ (made before the concrete host sends its first probe) is abstracted to a discrete uniform choice over $\{0, 1, 2\}$.

To enable the analysis of the protocol under different network scenarios we consider two different values of N corresponding to a networks with both a large number and small number of hosts. More precisely, we consider the case when $N = 1000$ (the value taken in [BvdSHV03]) and in the case when $N = 20$.

Following the above assumptions, we require only three abstract IP addresses:

- 0 – an address of an abstract host which the concrete host previously chose;
- 1 – an address of an abstract host which is the concrete host’s current choice;
- 2 – a fresh address which is the concrete host’s current choice.

As in the standard [CAG], we suppose that it takes between 0 and 1 seconds to send a packet between hosts (where the choice of the exact time delay is nondeterministic). Since the abstract hosts have already picked their IP address, by supposing that they always defend their addresses, the concrete host will never receive probes. It then follows that we do not need to record the

variable	description	range
$coll$	the number of address collisions detected by the concrete host	$0 \dots 10$
iph	the current address of the concrete host	$1 \dots 2$
$defend$	equals 1 when the host is defending its address (0 otherwise)	$0 \dots 1$
$probes$	the number of probes/ARPs sent by the concrete host	$0 \dots K$
ip	the address of the ARP packet currently being sent	$0 \dots 2$
n	the number of packets in the concrete host's output buffer	$0 \dots 8$
$b[i]$	the address of packet i in the concrete host's output buffer	$0 \dots 1$
m_0	the number of packets containing an IP address of type 0 in all of the buffers of the abstract hosts	$0 \dots 20$
m_1	the number of packets containing an IP address of type 1 in all of the buffers of the abstract hosts	$0 \dots 8$

Table 1: Integer variables used in the probabilistic timed automata

event	description
rec	concrete host receives an ARP packet from the environment
$send$	concrete host sends an ARP packet to the environment
$reset$	concrete host (re)configures a new IP address
$urgent$	the environment starts sending a packet

Table 2: Events used in the probabilistic timed automata

type of message being sent, but instead only the IP address in the message, and whether it is sent from the concrete host to the abstract hosts or vice versa.

As in [ZV03], we consider the case in which hosts use output buffers to store the packets they want to send. We have chosen the size of the buffers such that the probability of any buffer becoming full is negligible. We suppose that the concrete host can send a packet to all the abstract hosts at the same time and only one of the abstract hosts can send a packet to the concrete host at a time.

Variables. The set of variables of our probabilistic timed automata includes both clocks (x , y and z) and *integer variables* which are described in Table 1. Note that the range of the integer variable $probes$ is changed for different verification instances, and since the abstract IP address 2 corresponds to a fresh address chosen by the concrete host we need only two buffers for the abstract hosts (corresponding to addresses of type 0 and 1).

Events We model the protocol using two probabilistic timed automata, one for the concrete host and one for the environment (which comprises both abstract hosts and the output buffer of the concrete host); these models communicate through the 3 events rec , $send$ and $reset$ as described in Table 2. In Table 2 we have also included the event $urgent$, which although it is not a communicating event (it will appear only in the automaton representing the environment), is an urgent event (no time can pass if it enabled), since a packet should be sent as soon as it is possible. Note that, as we are considering the model $reset$, we use the event $reset$ to model the environment resetting the packets in the buffers of both the concrete and abstract hosts when the concrete host reconfigures.

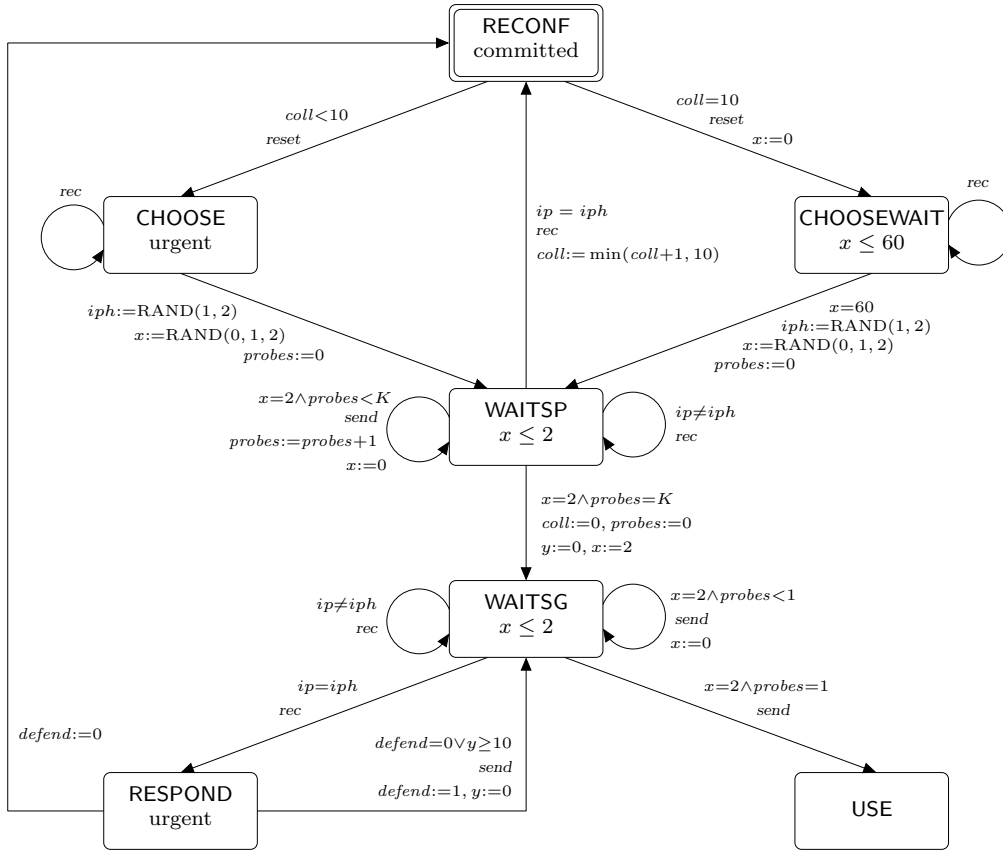


Figure 1: Probabilistic timed automaton for the concrete host

5.2 Probabilistic timed automata for the protocol.

In the following, we describe the modelling of the *reset* version of the protocol only. Recall that we use two probabilistic timed automata, one to model the concrete host and one to model the environment (the abstract hosts and the output buffers of *all* hosts). Note that in the description below, we have omitted the labelling of the non-urgent events on which the two automata do not synchronize, that is, the non-urgent events which do not appear in the set of events of each automata.

The concrete host. The model for the concrete host is shown in Figure 1. The host commences in the location **RECONF** (the double border indicates it is the initial location); this is a committed location, and therefore must be left immediately. In **RECONF**, the host chooses a new IP address by moving to the location **CHOOSE** if it has experienced less than ten address collisions, and to **CHOOSEWAIT** otherwise. These transitions are labelled with the event *reset* to inform the environment that the host’s buffer is to be reset (all messages in its buffer are to be removed).

In both **CHOOSE** and **CHOOSEWAIT**, the address selection is represented

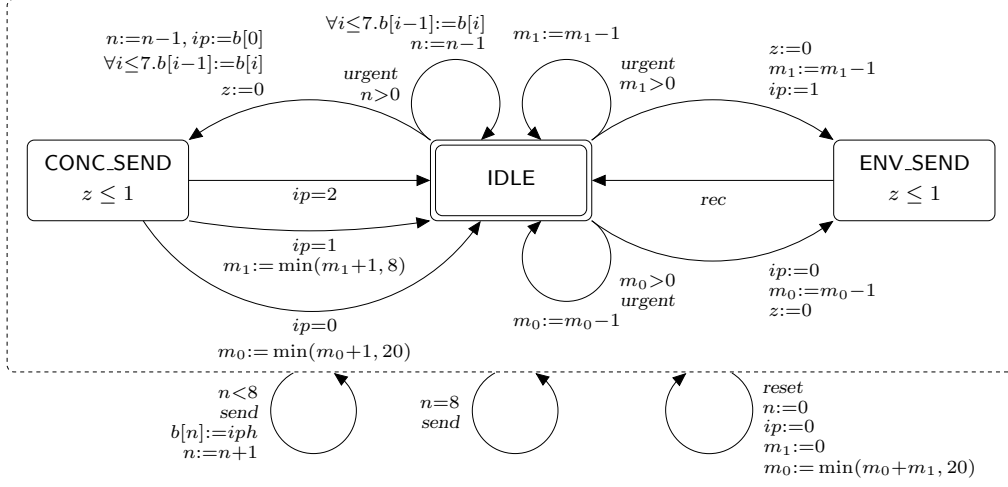


Figure 2: Probabilistic timed automaton for the environment

by the assignment $iph := \text{RAND}(1, 2)$, which corresponds to the host randomly selecting an IP address (using the probabilities given at the start of this section). The assignment to the clock x (a uniform choice between $\{0, 1, 2\}$) approximates the random delay of between 0 and 2 made by the host before sending the first probe. Note that, in **CHOOSEWAIT**, since the host has already experienced at least ten address collisions, it waits 60 seconds before choosing a new address.

In the location **WAITSP** the host sends K probes at 2 second intervals (the self-loop labelled with *send*). The host may also receive packets by means of the event *rec*. If it receives a packet which has a different IP address ($ip \neq iph$), then the host ignores the packet (and remains in **WAITSP**); however, if it has the same address, the host immediately reconfigures (moves to **RECONF**). When sending the K th probe, the host proceeds to location **WAITSG**, waits 2 seconds and then sends two gratuitous ARPs (re-using the variable *probes* to count these ARPs). After these ARPs have been sent, the host moves to **USE**. However, if while in **WAITSG** the host receives a packet with the same IP address, it moves to **RESPOND**. In this location, the host can decide to reconfigure (return to **RECONF**), or defend its IP address (by sending an ARP packet) if it has either not yet defended the address ($defend = 0$) or 10 seconds have passed since it previously defended the address ($y \geq 10$). This defence takes the form of sending of a defending packet, as denoted by the *send* labelled transition from **RESPOND** to **WAITSG**.

The environment. The model for the environment is shown in Figure 2. The dotted box labelled with three transitions which surrounds the model denotes that these transitions are available in *all* of the locations of the model. More precisely, in all locations, the environment may: receive a *send* event from the concrete host and, if the host's buffer is not full ($n < 8$), the corresponding packet is added to the buffer (otherwise it is lost); receive a *reset* event and clear the buffer of the concrete host ($n := 0$) and, since we assume that the

concrete host will never choose the same IP address twice, sets the IP address in any packet being sent or to be sent to type 0 (i.e. $ip := 0$, $m_1 := 0$ and $m_0 := \min(m_0 + m_1, 20)$).

The behaviour of the environment commences in the location IDLE. The transition which probabilistically moves to either IDLE or CONC_SEND corresponds to the environment sending a packet from the concrete host’s buffer. The *urgent* labelling denotes that the transition should be taken as soon as it is enabled, i.e. it should be taken as soon as there is something to send. Similarly, the transitions which move probabilistically to either IDLE or ENV_SEND correspond to an abstract host sending a packet, and are again urgent. There are two such transitions, since the address in the packet can either be of type 0 ($m_0 > 0$) or 1 ($m_1 > 0$). For each of these transitions, the loop (remaining in IDLE) corresponds to the packet being lost by the medium, while the other edge corresponds to the packet being sent correctly (therefore the required buffers are updated when one of these transitions is taken). Note that, since each of these transitions corresponds to a message from a different host, when more than one of these transitions is enabled, there is a nondeterministic choice as to which one is taken. We vary the probability of message loss depending on the verification instance. Once in either CONC_SEND or ENV_SEND, after a delay of between 0 and 1 seconds, the model returns to IDLE; this corresponds to the message taking between 0 and 1 seconds to send.

5.3 Verification using PRISM

In this section, we outline our results of using PRISM [KNP02] to verify the integer-time model of the probabilistic timed automata of the dynamic configuration protocol given in Section 5.2. In the experiments, as explained above we consider the cases when the number of hosts (N) equals 1000 and 20 and vary both the number of probes a host sends (K), and the probability of message loss. Further details can be found at the PRISM web page [Pri]. The algorithms used by PRISM for both probabilistic and expected reachability are taken from the literature; for probabilistic reachability see [BdA95], and for expected reachability see [dA97, dA99]. In both cases, verification reduces to solving a linear optimization problem on which one can apply iterative methods.

To apply model-checking methods we must ensure that the model under study has only finitely-many states and is finitely branching. From the construction given in Section 4, the integer-time model will have finitely many states. To ensure finite branching, we restrict the delays from \mathbb{N} to some finite set. More precisely, we allow delays of duration 1 only. Then, since any transition of duration $t \in \mathbb{N}$ can be modelled by a sequence of transitions of duration 1 and we restrict our attention to divergent adversaries, nothing is lost by omitting delays greater than 1 or equal to 0.

Note that, because we have abstracted certain aspects of the network (for example, the time taken to send a message), the presented results will give upper and lower bounds on the performance of the protocol, for example the actual reachability probability will lie in between the minimum and maximum reachability probabilities computed for the model under study.

number of abstract hosts equals 10000								
number of probes sent	probability of message loss							
	0.1		0.01		0.001		0	
	no_reset	reset	no_reset	reset	no_reset	reset	no_reset	reset
1	5.6e-4	5.6e-4	6.1e-6	6.2e-6	6.1e-8	6.2e-8	0	0
2	1.1e-4	1.1e-4	1.2e-7	1.2e-7	1.2e-10	1.2e-10	0	0
3	2.0e-5	2.0e-5	2.4e-9	2.4e-9	2.4e-13	2.5e-13	0	0
4	3.8e-6	3.9e-6	4.8e-11	4.9e-11	4.9e-16	5.0e-16	0	0

number of abstract hosts equals 20								
number of probes sent	probability of message loss							
	0.1		0.01		0.001		0	
	no_reset	reset	no_reset	reset	no_reset	reset	no_reset	reset
1	1.1e-5	1.1e-5	1.2e-7	1.2e-7	1.2e-9	1.2e-9	0	0
2	2.1e-6	2.1e-6	2.4e-9	2.4e-9	2.4e-12	2.4e-12	0	0
3	4.0e-7	4.0e-7	4.8e-11	4.8e-11	4.9e-15	4.9e-15	0	0
4	7.6e-8	7.6e-8	9.6e-13	9.6e-13	9.8e-18	9.8e-18	0	0

Table 3: Minimum probabilistic reachability results

number of abstract hosts equals 10000								
number of probes sent	probability of message loss							
	0.1		0.01		0.001		0	
	no_reset	reset	no_reset	reset	no_reset	reset	no_reset	reset
1	0.0154	0.0154	0.0154	0.0154	0.0154	0.0154	0.0154	0.0154
2	0.00298	0.00296	3.8e-4	3.1e-4	1.1e-4	3.1e-5	8.0e-5	0
3	5.6e-4	5.6e-4	7.2e-6	6.2e-6	1.3e-6	6.2e-8	1.2e-6	0
4	1.1e-4	1.1e-4	5.0e-7	1.2e-7	4.1e-7	1.2e-10	4.2e-7	0

number of abstract hosts equals 20								
number of probes sent	probability of message loss							
	0.1		0.01		0.001		0	
	no_reset	reset	no_reset	reset	no_reset	reset	no_reset	reset
1	3.1e-4	3.1e-4	3.1e-4	3.1e-4	3.1e-4	3.1e-4	3.1e-4	3.1e-4
2	5.8e-5	5.8e-5	6.2e-6	5.8e-5	6.5e-7	6.2e-7	3.2e-8	0
3	1.1e-5	1.1e-5	1.2e-7	1.1e-5	1.2e-9	1.2e-9	9.7e-12	0
4	2.1e-6	2.1e-6	2.4e-9	2.1e-6	5.7e-12	2.5e-12	3.2e-12	0

Table 4: Maximum probabilistic reachability results

5.3.1 Probabilistic reachability

The probabilistic reachability property we consider is the (minimum and maximum) probability of the host using an IP address which is already in use by another host. In Table 3 and Table 4 we present the minimum and maximum probabilistic reachability results obtained in the cases when the number of abstract hosts (N) is fixed at 1000 and at 20. Further details including the model checking times are available from the PRISM web page [Pri].

The results obtained show the expected result: increasing the number of probes sent decreases the probability of the host using an IP address which is already in use (recall that the number specified by the standard is four). When the probability of message loss is 0, Table 4 shows that the maximum probability is 0 for the the model reset (the model where the host clears its buffer) provided the host sends more than one probe. On the other hand, for




Figure 3: Minimum time-bounded reachability results (model `reset` and $N = 1000$)




Figure 4: Minimum time-bounded reachability results (model `reset` and $N = 20$)

the model `no_reset` (when the host does not clear its buffer), even if the host sends more than one probe, this maximum reachability probability is greater than 0. To understand this result, consider the fact that, if a host does not clear its buffer, then there is a chance that the probes corresponding to its new IP address will get delayed, and hence the host will not receive a reply to these probes until after it starts using the address (as the probability is 0, the host will eventually get a reply).

In the cases when message loss is greater than 0, the results presented in Table 3 and Table 4 again demonstrate that, by allowing the host to clear its buffer, the performance of the protocol improves; that is, the minimum and maximum reachability probability increases and decreases respectively.

For both models the results demonstrate that the probabilities decrease as the number of probes increases, which is to be expected since if one sends more probes there is a greater chance of receiving a reply to a probe when an IP address already in use is chosen (i.e. the chance that not all the probes and responses get lost). Furthermore, the probabilities in the case when $N = 20$ are smaller than when $N = 1000$, again this is to be expected, since when N is smaller there are fewer abstract hosts, and hence fewer IP addresses in use and therefore there is a smaller chance of the host choosing an IP address which is already in use.




Figure 5: Minimum time-bounded reachability results (model reset and $N = 1000$)




Figure 6: Maximum time-bounded reachability results (model reset and $N = 20$)

5.3.2 Time-bounded probabilistic reachability

The time-bounded probabilistic reachability property we consider is the (minimum and maximum) probability of the host using a fresh IP address within time T . For the model reset, the minimal time-bounded reachability results are presented in Figures 5 and 6, while the maximal time-bounded reachability results are given in Figures ?? and ?. Note that the graphs use a log scale and actually plot 1 minus the actual probabilities under study.

The results demonstrate that, for small time bounds, the probability is higher when K is smaller. This is to be expected since sending more probes takes more time. However, for larger time bounds the probability is larger when more probes are sent. This is due to the fact that, when more probes are sent, there is less chance of using an IP address already in use, and hence not reaching a state where the host uses a fresh IP address.

Also note that the probabilities are higher when $N = 20$ as opposed to 1000, the reasoning is the same as in the probabilistic reachability case, namely when $N = 20$ there is a greater chance that the host will choose an IP address.

The results for the model `no_reset` are similar to those presented for the `reset`, although the minimum probabilities are smaller and the maximum probabilities are larger for the model `no_reset`, for reasons similar to those given for the probabilistic reachability results.




Figure 7: Expected reachability results (model reset and $E = 10^6$)




Figure 8: Expected reachability results (model reset and $E = 10^{12}$)

5.3.3 Expected reachability

We consider the expected cost of a host choosing an IP address and using it. As in [BvdSHV03], the cost is defined as the time to start using an IP address plus an additional cost (E) associated with the host using an address which is already in use. We consider two different values for E , namely 10^6 and 10^{12} . Note that the choice of the value of this additional cost will depend on how damaging it is for two hosts to use the same IP address, which in turn depends on the network and the nature of its devices.

The results for the model reset are presented in Figure 7 and Figure 8 in the cases when E equals 10^6 and 10^{12} respectively. Note that, in each graph a log scale has been use to improve readability.

Note that the results for the model no_reset are similar, although the minimum costs are smaller and the maximum costs are larger for the model no_reset (see [Pri] for further details). This is to be expected, since the results for probabilistic reachability show that, when the host does not clear its buffer, there is a greater chance of it using an IP address which is already in use, and hence of incurring a greater cost.

These results are similar to those of [BvdSHV03]: as the message loss probability increases, one must increase the number of probes sent in order to reduce the expected cost; however, by sending too many probes the expected cost may then start to increase. The rationale for this is that, although increasing the number of probes sent decreases the probability of the host using an IP address which is already in use (that is, decreases the chance of incurring the additional

cost), it increases the expected time to choose an IP address (sending more probes takes more time).

The results also show that, as the probability of message loss increases, to minimize the expected costs one must send more probes (increase the value of K). Similarly, the results presented demonstrate the fact that when the cost of using an IP address which is already in use by another host increases one must send more probes to minimize the expected cost.

6 Conclusions

We have presented results demonstrating that digital clocks are sufficient for analysing a large class of probabilistic timed automata and performance properties. Since many of today's protocols include both timing and probabilistic behaviour, this approach is widely applicable, a fact which we illustrate by analysing the performance of the IPv4 dynamic configuration protocol.

By using probabilistic timed automata the model of the IPv4 protocol we consider differs from that present in the literature, in that, we are able to correctly model both the important probabilistic characteristics of the protocol and (timing) parameters faithful to the standard. For example, in [AHK03, BvdSHV03] the probabilistic characteristics of the protocol are correctly modelled but the timing parameters of the protocol are abstracted, whereas in [ZV03] the probabilistic attributes of the protocol are not considered but the timing parameters are correctly modelled. Furthermore, since our model is an abstraction of the actual protocol, by computing both minimum and maximum values for the properties that we consider, our results can be seen as giving bounds on the actual behaviour of the protocol.

Future work could consider extending the cost functions in order to vary the rate of cost accumulation in different locations, as in priced or weighted timed automata [BFH⁺01b, ATP01]. There are still limitations as to the size of the models that can be considered using digital clocks. In the case of probabilistic reachability, a generally more efficient approach is to consider *zones*, and in particular the backwards reachability approach introduced in [KNS01]. The application of zones to the verification of priced timed automata [LBB⁺01] may be instructive to this line of research.

References

- [ACD93] R. Alur, C. Courcoubetis, and D. Dill. Model checking in dense real time. *Information and Computation*, 104(1):2–34, 1993.
- [AD94] R. Alur and D. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.
- [AHK03] S. Andova, H. Hermanns, and J.-P. Katoen. Discrete-time rewards model-checked. In K. Larsen and P. Niebert, editors, *Proc. 1st Int. Workshop Formal Modeling and Analysis of Timed Systems*

- (*FORMATS 2003*), Lecture Notes in Computer Science. Springer-Verlag, 2003. to appear.
- [ATP01] R. Alur, S. Torre, and G. Pappas. Optimal paths in weighted timed automata. In M. Benedetto and A. Sangiovanni-Vincentelli, editors, *Proc. 4th Int. Workshop Hybrid Systems: Computation and Control (HSCC'01)*, volume 2034 of *Lecture Notes in Computer Science*, pages 49–62. Springer-Verlag, 2001.
- [BdA95] A. Bianco and L. de Alfaro. Model checking of probabilistic and nondeterministic systems. In P. Thiagarajan, editor, *Proc. 15th Conference on Foundations of Software Technology and Theoretical Computer Science*, volume 1026 of *Lecture Notes in Computer Science*, pages 499–513. Springer-Verlag, 1995.
- [BDL⁺01] G. Behrmann, A. David, K. Larsen, O. Möller, P. Pettersson, and W. Yi. UPPAAL - present and future. In *Proc. 40th IEEE Conf. Decision and Control (CDC'2001)*. IEEE Computer Society Press, 2001.
- [BFH⁺01a] G. Behrmann, A. Fehnker, T. Hune, K. Larsen, P. Pettersson, and J. Romijn. Efficient guiding towards cost-optimality in UPPAAL. In T. Margaria and W. Yi, editors, *Proc. 7th Int. Conf. Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2001)*, volume 2031 of *Lecture Notes in Computer Science*, pages 174–188. Springer-Verlag, 2001.
- [BFH⁺01b] G. Behrmann, A. Fehnker, T. Hune, K. Larsen, P. Pettersson, J. Romijn, and F. Vaandrager. Minimum-cost reachability for priced timed automata. In M. Benedetto and A. Sangiovanni-Vincentelli, editors, *Proc. 4th Int. Workshop Hybrid Systems: Computation and Control (HSCC'01)*, volume 2034 of *Lecture Notes in Computer Science*, pages 147–161. Springer-Verlag, 2001.
- [Bil79] P. Billingsley. *Probability and Measure*. John Wiley and Sons, 1979.
- [Bos99] D. Bosnacki. Digitization of timed automata. In S. Gnesi and D. Latella, editors, *Proc. 4th International Workshop on Formal Methods for Industrial Critical Systems (FMICS'99)*, pages 283–302. S.T.A.R.- CNR, 1999.
- [BT91] D. Bertsekas and J. Tsitsiklis. An analysis of stochastic shortest path problems. *Mathematics of Operations Research*, 16(3):580–595, 1991.
- [BvdSHV03] H. Bohnenkamp, P. v. d. Stok, H. Hermanns, and F. Vaandrager. Cost-optimisation of the IPv4 zeroconf protocol. In *Proc. Int. Performance and Dependability Symposium (IPDS 2003)*, pages 531–540. IEEE Computer Society Press, 2003.

- [CAG] S. Cheshire, B. Adoba, and E. Guttman. Dynamic configuration of IPv4 link-local addresses. Draft, August 2002. Available from www.ietf.org/internet-drafts/draft-ietf-zeroconf-ipv4-linklocal-07.txt.
- [dA97] L. de Alfaro. *Formal Verification of Probabilistic Systems*. PhD thesis, Stanford University, 1997.
- [dA99] L. de Alfaro. Computing minimum and maximum reachability times in probabilistic systems. In J. Baeten and S. Mauw, editors, *Proc. 10th International Conference on Concurrency Theory (CONCUR'99)*, volume 1664 of *Lecture Notes in Computer Science*, pages 66–81. Springer-Verlag, 1999.
- [Der70] C. Derman. *Finite-State Markovian Decision Processes*. Academic Press, 1970.
- [DKN02] C. Daws, M. Kwiatkowska, and G. Norman. Automatic verification of the IEEE 1394 root contention protocol with KRONOS and PRISM. In R. Cleaveland and H. Garavel, editors, *Proc. 7th International Workshop on Formal Methods for Industrial Critical Systems (FMICS'02)*, volume 66(2) of *Electronic Notes in Theoretical Computer Science*. Elsevier Science, 2002.
- [DY95] C. Daws and S. Yovine. Two examples of verification of multirate timed automata with KRONOS. In *Proc. 16th IEEE Real-Time Systems Symposium (RTSS'95)*, pages 66–75. IEEE Computer Society Press, 1995.
- [Hen91] T.A. Henzinger. *The Temporal Specification and Verification of Real-time Systems*. PhD thesis, Stanford University, 1991.
- [HHWT95] T. Henzinger, P.-H. Ho, and H. Wong-Toi. A user guide to HYTECH. In *Proc. TACAS'95*, volume 1019 of *LNCS*, pages 41–71. Springer, 1995.
- [HMP92] T. Henzinger, Z. Manna, and A. Puneli. What good are digital clocks? In W. Kuich, editor, *Proc. Automata, Languages and Programming (ICALP'92)*, volume 623 of *Lecture Notes in Computer Science*, pages 545–558. Springer-Verlag, 1992.
- [KNP02] M. Kwiatkowska, G. Norman, and D. Parker. PRISM: Probabilistic symbolic model checker. In T. Field, P. Harrison, J. Bradley, and U. Harder, editors, *Proc. 12th International Conference on Modelling Techniques and Tools for Computer Performance Evaluation (TOOLS'02)*, volume 2324 of *Lecture Notes in Computer Science*, pages 200–204. Springer-Verlag, 2002.
- [KNS01] M. Kwiatkowska, G. Norman, and J. Sproston. Symbolic computation of maximal probabilistic reachability. In K. Larsen and

- M. Nielsen, editors, *Proc. 13th International Conference on Concurrency Theory (CONCUR'01)*, volume 2154 of *LNCS*, pages 169–183. Springer, 2001.
- [KNS03] M. Kwiatkowska, G. Norman, and J. Sproston. Probabilistic model checking of deadline properties in the IEEE 1394 FireWire root contention protocol. *Formal Aspects of Computing*, 14(3):295–318, 2003.
- [KNSS02] M. Kwiatkowska, G. Norman, R. Segala, and J. Sproston. Automatic verification of real-time systems with discrete probability distributions. *Theoretical Computer Science*, 282:101–150, 2002.
- [KSK76] J. Kemeny, J. Snell, and A. Knapp. *Denumerable Markov Chains*. Graduate Texts in Mathematics. Springer-Verlag, 2nd edition, 1976.
- [LBB⁺01] K. Larsen, G. Behrmann, E. Brinksma, A. Fehnker, T. Hune, P. Pettersson, and J. Romijn. As cheap as possible: Efficient cost-optimal reachability for priced timed automata. In G. Berry, H. Comon, and A. Finkel, editors, *Proc. 13th International Conference on Computer Aided Verification (CAV'01)*, volume 2102 of *Lecture Notes in Computer Science*, pages 493–505. Springer-Verlag, 2001.
- [OW03] J. Ouaknine and J. Worrell. Universality and language inclusion for open and closed timed automata. In O. Maler and A. Pnueli, editors, *Proc. 6th Int. Workshop Hybrid Systems: Computation and Control (HSCC'03)*, volume 2623 of *Lecture Notes in Computer Science*, pages 375–388. Springer-Verlag, 2003.
- [Pri] PRISM web page. <http://www.cs.bham.ac.uk/~dxp/prism/>.
- [Seg95] R. Segala. *Modelling and Verification of Randomized Distributed Real Time Systems*. PhD thesis, Massachusetts Institute of Technology, 1995.
- [SL94] R. Segala and N. Lynch. Probabilistic simulations for probabilistic processes. In B. Jonsson and J. Parrow, editors, *Proc. 5th International Conference on Concurrency Theory (CONCUR'94)*, volume 836 of *LNCS*, pages 481–496. Springer, 1994.
- [Tri98] S. Tripakis. *The formal analysis of timed systems in practice*. PhD thesis, Université Joseph Fourier, 1998.
- [Var85] M. Vardi. Automatic verification of probabilistic concurrent finite state programs. In *Proc. 26th Annual Symposium on Foundations of Computer Science (FOCS'85)*, pages 327–338. IEEE Computer Society Press, 1985.

- [ZV03] M. Zhang and F. Vaandrager. Analysis of a protocol for dynamic configuration of IPv4 link local addresses using UPPAAL. Technical Report, NIII, University of Nijmegen, 2003.