Strict Divergence for Probabilistic Timed Automata^{*}

Jeremy Sproston

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

Abstract. Probabilistic timed automata are an extension of timed automata with discrete probability distributions. In previous work, a probabilistic notion of time divergence for probabilistic timed automata has been considered, which requires the divergence of time with probability 1. We show that this notion can lead to cases in which the probabilistic timed automaton satisfies a correctness requirement by making an infinite number of probabilistic transitions in a finite amount of time. To avoid such cases, we consider strict time divergence which concerns the divergence of time over all paths, rather than time divergence of paths with probability 1. We present new model-checking algorithms for probabilistic timed automata both for probabilistic and strict divergence. The algorithms have the same complexity as the previous model-checking algorithms for probabilistic timed automata.

1 Introduction

Model checking is an automatic verification technique for establishing that a model of a system satisfies a formally-specified property [1]. Two particular classes of systems have been subject to extensions of the basic model-checking paradigm. Firstly, methods for *timed systems*, in which the durations of system behaviours is critical for the system's correctness, have been developed, with particular emphasis on techniques for the system-description formalism of timed automata [2]. Secondly, methods for *probabilistic systems*, in which system behaviours have associated probabilities of occurrence, have been introduced, in this case concentrating on techniques for Markov chains (in which the choice between transitions is probabilistic) or Markov decision processes (in which the choice between transitions is both nondeterministic and probabilistic). In this paper, we consider methods for *probabilistic timed systems*, in which both timed and probabilistic behaviour coexist. In the context of probabilistic timed systems, a correctness requirement typically combines probabilistic and timing thresholds, such as “a request is followed by a response within 5 time units with probability 0.99 or greater”. A number of model-checking methods for system-description formalisms for such systems, which generally can differ in terms of the way in

^{*} Supported in part by the MIUR-PRIN project PaCo - Performability-Aware Computing: Logics, Models and Languages.

which the interaction of probability and time is modelled both in the system and in the correctness requirements, have been presented [3–9]. Our focus is on the system-description formalism of *probabilistic timed automata* [10, 6], which can be regarded as an extension of timed automata with discrete probability distributions, or, equivalently, an extension of Markov decision processes with timed automata-like clocks, constraints and resets. Probabilistic timed automata have been used to model systems such as the IEEE 1394 root contention protocol, the backoff procedure in IEEE 802.11 Wireless LANs, and the IPv4 Zeroconf protocol [11].

When modelling timed systems, the issue of time divergence is of importance. Roughly speaking, behaviours of the model which correspond to the case in which the amount of time elapsed converges do not correspond to phenomena that a real system can exhibit, and therefore should be excluded from model-checking analyses. Methods for model checking timed automata therefore are defined in such a way as to consider divergent behaviours only [12–14]. Recall that, for probabilistic timed automata (as for Markov decision processes), a strategy is a function which resolves the nondeterminism of the system, by mapping finite system behaviours to nondeterministic alternative transitions available in the last state of the behaviour. For probabilistic timed automata, a probabilistic notion of time divergence has been presented [6], which requires that time diverges with probability 1 for all strategies of the model. We henceforth refer to this notion as *probabilistic divergence*. Furthermore, a model-checking algorithm for the probabilistic timed temporal logic PTCTL based on the standard region graph construction [2, 12] is presented in [6], and which computes the correct probability of property satisfaction in the case in which all strategies of the model are probabilistically divergent. This is guaranteed when all structural loops in the graph of the probabilistic automaton require that at least one time unit elapses, which holds for many case studies considered [11]. The algorithm runs in EXPTIME, which is optimal by the results of [15]. Furthermore, a symbolic probabilistic model-checking method for probabilistic timed automata has been presented which can be applied also if not all strategies of the system are probabilistically divergent [16], although this algorithm does not run in EXPTIME.

There remain two questions in this context. The first concerns whether there exists an EXPTIME algorithm for PTCTL model checking of probabilistic timed automata with probabilistically divergent strategies. The second question concerns whether the notion of probabilistic divergence is generally applicable. Consider the probabilistic timed automaton of Figure 1 (where edges without a probability label correspond to probability 1). From the location l_0 , there exists a probabilistically divergent strategy to reach l_2 with probability 1. An example of such a strategy is the following: the first time l_0 is visited, let $\frac{1}{2}$ time units elapse, then select the rightmost transition; if the probabilistic choice is resolved so that l_0 is then visited, let $\frac{1}{4}$ time units elapse, then take the rightmost transition again; if l_0 is visited for a third time, then let $\frac{1}{8}$ time units elapse, then take the rightmost transition again, and so on. Then the value of the clock x will never be 1, and the strategy will be able to take the rightmost transition an

Fig. 1. Maximal reachability probabilities under probabilistic and strict divergence.

infinite number of times, resulting in the probability of $\sum_{k \geq 1} \frac{1}{2^k} = 1$ of reaching l_2 . However, in the case in which we assume that the selection of the rightmost transition corresponds to the change of some physical state in the system, the behaviour exhibited by this strategy should be excluded from the analysis of the system. Therefore we argue that it is important to have an alternative to probabilistic divergence. We propose *strict divergence*, which requires that *all behaviours* (rather than probability 1 of the behaviours) of a strategy should be time divergent. Note that, with the requirement of strict divergence, the maximum probability of reaching l_2 from l_0 can be made to be arbitrarily close to 1 (by making arbitrarily large the upper bound on the number of times the rightmost transition is selected before the leftmost transition is taken). However, there is no strictly divergent strategy which reaches l_2 from l_0 actually with probability 1. Therefore, the correctness property which specifies that l_2 is reached with probability 1 for some strategy is satisfied under probabilistic divergence, but not under strict divergence.

In this paper, after first recalling the definition of probabilistic timed automata and PTCTL, we present an EXPTIME algorithm for PTCTL model checking with probabilistic divergence in Section 3. Then, in Section 4, we present an EXPTIME algorithm for PTCTL model checking with strict divergence.

Related work. The distinction between probabilistic and strict divergence is inspired by the distinction between fairness and strict fairness, as introduced by Baier and Kwiatkowska in the context of model checking Markov decision processes [17, 18]. We note that [19] features randomized strategies in the context of 2-player timed games which are required to be either time divergent, or blameless for the convergence of time, over all paths. In [20], the probability of behaviours satisfying a (Büchi) correctness requirement *and* are divergent can be computed. We do not follow this approach, in which the correctness property is adapted to encode also the divergence of time, because it does not exclude strategies in which time converges with positive probability.

2 Probabilistic Timed Automata

Preliminaries. We use $\mathbb{R}_{\geq 0}$ to denote the set of non-negative real numbers, $\mathbb{N}$ to denote the set of natural numbers, and AP to denote a set of atomic propositions. Given a set Q and a function $\mu : Q \rightarrow \mathbb{R}_{\geq 0}$, we define $\text{support}(\mu) = \{q \in Q \mid \mu(q) > 0\}$. A (discrete) probability *distribution* over a countable set Q is a function $\mu : Q \rightarrow [0, 1]$ such that $\sum_{q \in Q} \mu(q) = 1$. Let $\text{Dist}(Q)$ be the set

of distributions over Q . If Q is an uncountable set, we define $\text{Dist}(Q)$ to be the set of functions $\mu : Q \rightarrow [0, 1]$, such that $\text{support}(\mu)$ is a countable set and μ restricted to $\text{support}(\mu)$ is a (discrete) probability distribution.

A *timed Markov decision process* (TMDP) $\mathsf{T} = (S, \rightarrow, \text{lab})$ comprises the following components: a (possibly uncountable) set of *states* S ; a (possibly uncountable) *timed probabilistic, nondeterministic transition relation* $\rightarrow \subseteq S \times \mathbb{R}_{\geq 0} \times \text{Dist}(S)$; and a *labelling function* $\text{lab} : S \rightarrow 2^{AP}$. The transitions from state to state of a TMDP are performed in two steps: given that the current state is s , the first step concerns a nondeterministic selection of $(s, d, \mu) \in \rightarrow$, where d corresponds to the duration of the transition; the second step comprises a probabilistic choice, made according to the distribution μ , as to which state to make the transition to (that is, we make a transition to a state $s' \in S$ with probability $\mu(s')$). We often denote such a completed transition by $s \xrightarrow{d, \mu} s'$. A TMDP is *total* if, for each state $s \in S$, there exists at least one transition $(s, _, _) \in \rightarrow$.

An *infinite path* of T is an infinite sequence of transitions $r = s_0 \xrightarrow{d_0, \mu_0} s_1 \xrightarrow{d_1, \mu_1} \dots$ such that the target state of one transition is the source state of the next. Similarly, a *finite path* of T is a finite sequence of consecutive transitions $r = s_0 \xrightarrow{d_0, \mu_0} s_1 \xrightarrow{d_1, \mu_1} \dots \xrightarrow{d_{n-1}, \mu_{n-1}} s_n$. If r is finite, the length of r , denoted by $|r|$, is equal to the number of transitions along r . If r is infinite, we let $|r| = \infty$. We use $\text{Path}_{\text{ful}}^{\mathsf{T}}$ to denote the set of infinite paths of T , and $\text{Path}_{\text{fin}}^{\mathsf{T}}$ the set of finite paths of T . When clear from the context, we omit the superscript T . If r is a finite path, we denote by $\text{last}(r)$ the last state of r . For any path r and $i \leq |r|$, let $r(i) = s_i$ be the $(i+1)$ th state along r , and let $\text{step}(r, i) = \mu_i$ be the $(i+1)$ th distribution taken along r . Let $\text{Path}_{\text{ful}}^{\mathsf{T}}(s)$ and $\text{Path}_{\text{fin}}^{\mathsf{T}}(s)$ refer to the sets of infinite and finite paths of T , respectively, commencing in state $s \in S$.

A *strategy* of a TMDP T is a function σ mapping every finite path $r \in \text{Path}_{\text{fin}}^{\mathsf{T}}$ to a transition $(\text{last}(r), d, \mu) \in \rightarrow$. Let Σ_{T} be the set of strategies of T (when the context is clear, we write simply Σ). For any strategy $\sigma \in \Sigma$, let $\text{Path}_{\text{ful}}^{\sigma}$ and $\text{Path}_{\text{fin}}^{\sigma}$ denote the sets of infinite and finite paths, respectively, resulting from the choices of σ . For a state $s \in S$, let $\text{Path}_{\text{ful}}^{\sigma}(s) = \text{Path}_{\text{ful}}^{\sigma} \cap \text{Path}_{\text{ful}}(s)$ and $\text{Path}_{\text{fin}}^{\sigma}(s) = \text{Path}_{\text{fin}}^{\sigma} \cap \text{Path}_{\text{fin}}(s)$. Given a strategy $\sigma \in \Sigma$ and a state $s \in S$, we define the probability measure Prob_s^{σ} over $\text{Path}_{\text{ful}}^{\sigma}(s)$ in the standard way [21].

An *untimed Markov decision process* (MDP) $\mathsf{M} = (S, \rightarrow, \text{lab})$ is defined as a TMDP, but for which $\rightarrow \subseteq S \times \text{Dist}(S)$ (that is, the transition relation $\rightarrow$ does not contain timing information). A sub-MDP $(S', \rightarrow', \text{lab}|_{S'})$ of M is an MDP such that $S' \subseteq S$, $\rightarrow' \subseteq \rightarrow$, and $\text{lab}|_{S'}$ is equal to lab restricted to S' . Let $T \subseteq S$. The *sub-MDP of M induced by T* is the sub-MDP $(T, \rightarrow|_T, \text{lab}|_T)$ of M , where $\rightarrow|_T = \{(s, \nu) \in \rightarrow \mid s \in T \wedge \text{support}(\nu) \subseteq T\}$. Occasionally we omit the labelling function lab for MDPs. The graph of an MDP $(S, \rightarrow)$ is the pair (S, E) where $(s, s') \in E$ if and only if there exists $(s, \mu) \in \rightarrow$ such that $s' \in \text{support}(\mu)$. An *end component* (EC) of an MDP M is a sub-MDP $(C, D) \in 2^S \times 2^{\rightarrow}$ such that (1) if $(s, \mu) \in D$, then $s \in C$ and $\text{support}(\mu) \subseteq C$, and (2) the graph of (C, D) is strongly connected [5]. An end component (C, D) of M is *maximal* if there

does not exist any EC (C', D') of M such that $(C, D) \neq (C', D')$, $C \subseteq C'$ and $D \subseteq D'$.

Probabilistic timed automata. Let $\mathcal{X}$ be a finite set of real-valued variables called *clocks*, the values of which increase at the same rate as real-time. The set $CC(\mathcal{X})$ of *clock constraints* over $\mathcal{X}$ is defined as the set of conjunctions over atomic formulae of the form $x \sim c$, where $x, y \in \mathcal{X}$, $\sim \in \{<, \leq, >, \geq, =\}$, and $c \in \mathbb{N}$. A *probabilistic timed automaton* (PTA) $P = (L, \mathcal{X}, inv, prob, \mathcal{L})$ consists of the following components: a finite set L of *locations*; a finite set $\mathcal{X}$ of clocks; a function $inv : L \rightarrow CC(\mathcal{X})$ associating an *invariant condition* with each location; a finite set $prob \subseteq L \times CC(\mathcal{X}) \times \text{Dist}(2^{\mathcal{X}} \times L)$ of *probabilistic edges* such that, for each $l \in L$, there exists at least one $(l, _, _) \in prob$; and a *labelling function* $\mathcal{L} : L \rightarrow 2^{AP}$. A probabilistic edge $(l, g, p) \in prob$ is a triple containing (1) a source location l , (2) a clock constraint g , called a *guard*, and (3) a probability distribution p which assigns probability to pairs of the form (X, l') , where $X \subseteq \mathcal{X}$ is a clock set X and $l' \in L$ is a location. The behaviour of a probabilistic timed automaton takes a similar form to that of a timed automaton [2]: in any location time can advance as long as the invariant holds, and a probabilistic edge can be taken if its guard is satisfied by the current values of the clocks. However, probabilistic timed automata generalize timed automata in the sense that, once a probabilistic edge is nondeterministically selected, the choice of which clocks to reset and which target location to make the transition to is *probabilistic*.

We refer to a mapping $v : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ as a *clock valuation*. Let $\mathbb{R}_{\geq 0}^{\mathcal{X}}$ denote the set of clock valuations. For a clock valuation $v \in \mathbb{R}_{\geq 0}^{\mathcal{X}}$ and a value $d \in \mathbb{R}_{\geq 0}$, we use $v + d$ to denote the clock valuation such that $(v + d)(x) = v(x) + d$ for all clocks $x \in \mathcal{X}$. For a clock set $X \subseteq \mathcal{X}$, we let $v[X := 0]$ be the clock valuation obtained from v by resetting all clocks within X to 0; formally, we let $v[X := 0](x) = 0$ for all $x \in X$, and let $v[X := 0](x) = v(x)$ for all $x \in \mathcal{X} \setminus X$. The clock valuation v *satisfies* the clock constraint $\psi \in CC(\mathcal{X})$, written $v \models \psi$, if and only if ψ resolves to true after substituting each clock $x \in \mathcal{X}$ with the corresponding clock value $v(x)$.

The semantics of the probabilistic timed automaton $P = (L, \mathcal{X}, inv, prob, \mathcal{L})$ is the TMDP $T[P] = (S, \rightarrow, lab)$ where:

- $S = \{(l, v) \mid l \in L \text{ and } v \in \mathbb{R}_{\geq 0}^{\mathcal{X}} \text{ s.t. } v \models inv(l)\}$;
- $\rightarrow$ is the smallest set such that $((l, v), d, \mu) \in \rightarrow$ if there exist $d \in \mathbb{R}_{\geq 0}$ and a probabilistic edge $(l, g, p) \in prob$ where:
 1. $v + d \models g$, and $v + d' \models inv(l)$ for all $0 \leq d' \leq d$;
 2. for any $(X, l') \in 2^{\mathcal{X}} \times L$, we have that $p(X, l') > 0$ implies $(v + d)[X := 0] \models inv(l')$;
 3. for any $(l', v') \in S$, we have that $\mu(l', v') = \sum_{X \in \text{Reset}(v, d, v')} p(X, l')$, where $\text{Reset}(v, d, v') = \{X \subseteq \mathcal{X} \mid (v + d)[X := 0] = v'\}$.
- lab is such that $lab(l, v) = \mathcal{L}(l)$ for each state $(l, v) \in S$.

We restrict our attention to PTA P with a semantic TMDP $T[P]$ which is total. This can be guaranteed by a syntactic condition on PTA which has been presented in [22], and which holds generally for PTA models in practice [11].

$$\begin{aligned}
s, w \models_{\Sigma} a & \quad \text{iff } a \in \text{lab}(s) \\
s, w \models_{\Sigma} z \sim c & \quad \text{iff } w(z) \sim c \\
s, w \models_{\Sigma} z \cdot \Phi & \quad \text{iff } s, w[z := 0] \models_{\Sigma} \Phi \\
s, w \models_{\Sigma} \Phi_1 \wedge \Phi_2 & \quad \text{iff } s, w \models_{\Sigma} \Phi_1 \text{ and } s, w \models_{\Sigma} \Phi_2 \\
s, w \models_{\Sigma} \neg \Phi & \quad \text{iff } s, w \not\models_{\Sigma} \Phi \\
s, w \models_{\Sigma} \mathbb{P}_{\bowtie \zeta}(\varphi) & \quad \text{iff } \text{Prob}_s^{\sigma} \{r \in \text{Path}_{\text{ful}}^{\sigma}(s) \mid r, w \models_{\Sigma} \varphi\} \bowtie \zeta \text{ for all } \sigma \in \Sigma \\
r, w \models_{\Sigma} \Phi_1 \mathcal{U} \Phi_2 & \quad \text{iff } \exists \text{ position } (i, \delta) \text{ of } r \text{ s.t. } r(i, \delta), w + \text{Dur}(r, i, \delta) \models_{\Sigma} \Phi_2, \\
& \quad \text{and } \forall \text{ positions } (j, \delta') \text{ of } r \text{ s.t. } (j, \delta') \prec_r (i, \delta) \text{ we have} \\
& \quad r(j, \delta'), w + \text{Dur}(r, j, \delta') \models_{\Sigma} \Phi_1 \vee \Phi_2.
\end{aligned}$$

Fig. 2. Semantics of PTCTL.

We say that σ is a strategy of P if σ is a strategy of $\Sigma_{T[P]}$. Given a path $r = (l_0, v_0) \xrightarrow{d_0, \mu_0} (l_1, v_1) \xrightarrow{d_1, \mu_1} \dots$ of $T[P]$, for every $i \in \mathbb{N}$, we use $r(i, d)$, with $0 \leq d \leq d_i$, to denote the state $(l_i, v_i + d)$ reached from (l_i, v_i) after delaying d time units. A pair (i, d) is called a *position* of r . We define a total order on positions of r : given positions $(i, d), (j, d')$ of r , the position (i, d) precedes (j, d') — denoted $(i, d) \prec_r (j, d')$ — if and only if either $i < j$, or $i = j$ and $d < d'$.

To reason about time divergence in the remainder of the paper, we construct a modified PTA in the following manner [23, 24]. First we add a new atomic proposition *tick* to AP . Given a PTA $P = (L, \mathcal{X}, \text{inv}, \text{prob}, \mathcal{L})$, its *enlarged PTA* $P' = (L', \mathcal{X}', \text{inv}', \text{prob}', \mathcal{L}')$ is constructed as follows. For each location $l \in L$, we introduce a new location $\bar{l}$. Let $L' = L \cup \{\bar{l} \mid l \in L\}$, and $\mathcal{X}' = \mathcal{X} \cup \{\mathfrak{z}\}$. For each $l \in L$, let $\text{inv}'(l) = \text{inv}'(\bar{l}) = \text{inv}(l) \wedge (\mathfrak{z} \leq 1)$. Let $\text{prob}' = \text{prob} \cup \{(l, (\mathfrak{z} = 1), p_{(\emptyset, \bar{l})}), (\bar{l}, (\mathfrak{z} = 1), p_{(\{\mathfrak{z}\}, l))} \mid l \in L\}$, where $p_{(\emptyset, \bar{l})}$ and $p_{(\{\mathfrak{z}\}, l)}$ are the distributions assigning probability 1 to the elements $(\emptyset, \bar{l})$ and $(\{\mathfrak{z}\}, l)$, respectively. Finally, let $\mathcal{L}'(l) = \mathcal{L}(l)$ and $\mathcal{L}'(\bar{l}) = \mathcal{L}(l) \cup \{\text{tick}\}$ for each $l \in L$. Note that *tick* becomes true at all natural numbered time points after the start of execution of the PTA. In the remainder of the paper, we assume that all considered PTA are enlarged.

Probabilistic timed temporal logic. We now describe a *probabilistic, timed* temporal logic which combines PCTL [4, 25] and TCTL [12, 13], and which can be used to specify properties of probabilistic timed automata [6]. Assume that the PTA $P = (L, \mathcal{X}, \text{inv}, \text{prob}, \mathcal{L})$ is fixed. Let $\mathcal{Z}$ be a finite set of clocks disjoint from $\mathcal{X}$. Clocks in the set $\mathcal{Z}$ are called *formula clocks*. Valuations of formula clocks are denoted by $w : \mathcal{Z} \rightarrow \mathbb{R}_{\geq 0}$. The formulae of PTCTL (Probabilistic Timed Computation Tree Logic) are given by the following grammar:

$$\Phi ::= a \mid z \sim c \mid z \cdot \Phi \mid \Phi \wedge \Phi \mid \neg \Phi \mid \mathbb{P}_{\bowtie \zeta}(\Phi \mathcal{U} \Phi)$$

where $a \in AP$ is an atomic proposition, $z \in \mathcal{Z}$ is a formula clock, $\sim \in \{<, \leq, =, \geq, >\}$, $c \in \mathbb{N}$ is a natural number, $\bowtie \in \{<, \leq, \geq, >\}$, and $\zeta \in [0, 1]$ is a probability.

We proceed to define the satisfaction relation of PTCTL for TMDPs. Given the infinite path $r = s_0 \xrightarrow{d_0, \mu_0} s_1 \xrightarrow{d_1, \mu_1} \dots$ of the TMDP T , let $\text{Dur}(r, i, d) = d + \sum_{0 \leq k < i} d_k$ be the accumulated duration along r until position (i, d) . Given a set of strategies $\Sigma \subseteq \Sigma_{\mathsf{T[P]}}$ of P , and a PTCTL formula Φ , we define the satisfaction relation $\models_\Sigma$ of PTCTL as in Figure 2.

In the following, for simplicity, we generally encode formula clock valuations within the state of a PTA: that is, a state $(l, v) \in S$ consists of a location l and a clock valuation $v \in \mathbb{R}_{\geq 0}^{(\mathcal{X} \cup \mathcal{Z})}$. This allows us to write $s \models_\Sigma \Phi$ rather than $s, w \models_\Sigma \Phi$.

The model-checking problem for a PTA P and a PTCTL formula Φ , given a set $\Sigma \subseteq \Sigma_{\mathsf{T[P]}}$ of strategies, consists of computing the set $\llbracket \Phi \rrbracket_\Sigma = \{s \in S \mid s \models_\Sigma \Phi\}$. When clear from the context, we write $\llbracket \Phi \rrbracket$ rather than $\llbracket \Phi \rrbracket_\Sigma$. From [6, 15], the PTCTL model-checking problem with respect to the full set $\Sigma_{\mathsf{T[P]}}$ of strategies is EXPTIME-complete.

3 Probabilistic Divergence

In this section, we give an EXPTIME algorithm for PTCTL model checking of PTA with the definition of probabilistically-divergent strategies of [6]. Throughout this section, we assume that the PTA $\mathsf{P} = (L, \mathcal{X}, \text{inv}, \text{prob}, \mathcal{L})$ and the PTCTL formula Φ , which has a set $\mathcal{Z}$ of formula clocks, are fixed. A path $r \in \text{Path}_{\text{ful}}$ is *divergent* if $\lim_{k \rightarrow \infty} \text{Dur}(r, k, 0) = \infty$. Let Timediv be the set of divergent paths. A strategy $\sigma \in \Sigma_{\mathsf{T[P]}}$ is *probabilistically divergent* if, for all states $s \in S$, we have $\text{Prob}_s^\sigma(\text{Timediv}) = 1$. The set of all probabilistically divergent strategies of P is denoted by $\Sigma_{\mathsf{P}}^{\text{pd}}$.

Region MDP. Our first task is to construct an MDP from P and Φ by using the standard region graph construction [2, 6]. For $t \in \mathbb{R}_{\geq 0}$, we let $\text{frac}(t) = t - \lfloor t \rfloor$. For each clock $x \in \mathcal{X} \cup \mathcal{Z}$, we let c_x be the maximal constant to which x is compared in any of the guards of probabilistic edges or invariants of P , or in a clock constraint in the formula Φ (if x is not involved in any clock constraint of P or Φ , we let $c_x = 1$). Two clock valuations $v, v' \in \mathbb{R}_{\geq 0}^{(\mathcal{X} \cup \mathcal{Z})}$ are *clock equivalent* if the following conditions are satisfied: (1) for all clocks $x \in \mathcal{X} \cup \mathcal{Z}$, we have $v(x) \leq c_x$ if and only if $v'(x) \leq c_x$; (2) for all clocks $x \in \mathcal{X} \cup \mathcal{Z}$ with $v(x) \leq c_x$, we have $\lfloor v(x) \rfloor = \lfloor v'(x) \rfloor$; (3) for all clocks $x, y \in \mathcal{X} \cup \mathcal{Z}$ with $v(x) \leq c_x$ and $v(y) \leq c_y$, we have $\text{frac}(v(x)) \leq \text{frac}(v(y))$ if and only if $\text{frac}(v'(x)) \leq \text{frac}(v'(y))$; and (4) for all clocks $x \in \mathcal{X} \cup \mathcal{Z}$ with $v(x) \leq c_x$, we have $\text{frac}(v(x)) = 0$ if and only if $\text{frac}(v'(x)) = 0$. We use α and β to refer to classes of clock equivalence.

Two states $(l, v), (l', v')$ are *region equivalent* if (1) $l = l'$, and (2) v and v' are clock equivalent. A *region* is an equivalence class of region equivalence. Let Regions be the set of regions of P and Φ . The number of regions corresponding to the PTA P and the PTCTL formula Φ is bounded by $|L| \cdot \prod_{x \in \mathcal{X} \cup \mathcal{Z}} (c_x + 1) \cdot |\mathcal{X} \cup \mathcal{Z}|! \cdot 2^{|\mathcal{X} \cup \mathcal{Z}|}$.

The set of regions of a PTA P and the PTCTL formula Φ can be used to construct an untimed, finite-state MDP $\text{Reg}[P, \Phi] = (\text{Regions}, \rightarrow_{\text{Reg}}, \text{lab}_{\text{Reg}})$ in the following way. The set of states of $\text{Reg}[P, \Phi]$ is the set Regions of regions. The transition relation $\rightarrow_{\text{Reg}} \subseteq \text{Regions} \times \text{Dist}(\text{Regions})$ is the smallest set such that $((l, \alpha), \nu) \in \rightarrow_{\text{Reg}}$ if there exists $((l, v), d, \mu) \in \rightarrow$ such that (1) $v \in \alpha$, and (2) for each $(l', \beta) \in \text{Regions}$ such that there exists $(l', v') \in \text{support}(\mu)$ and $v' \in \beta$ (by definition, this (l', v') will be unique), we have $\nu(l', \beta) = \mu(l', v')$, otherwise $\nu(l', \beta) = 0$. For each region $(l, \alpha) \in \text{Regions}$, we let $\text{lab}_{\text{Reg}}(l, \alpha) = \mathcal{L}(l)$.

Given a clock valuation v , the unique clock equivalence class to which v belongs is denoted by $[v]$. Given a state $(l, v) \in S$, the unique region to which (l, v) belongs is $(l, [v])$, and is denoted by $[(l, v)]$. An infinite path $r = s_0 \xrightarrow{d_0, \mu_0} s_1 \xrightarrow{d_1, \mu_1} \dots$ of $T[P]$ corresponds to a unique infinite path $[r] = [s_0] \xrightarrow{\nu_0} [s_1] \xrightarrow{\nu_1} \dots$. Similarly, a finite path $r = s_0 \xrightarrow{d_0, \mu_0} s_1 \xrightarrow{d_1, \mu_1} \dots \xrightarrow{d_{n-1}, \mu_{n-1}} s_n$ of $T[P]$ corresponds to a unique finite path $[r] = [s_0] \xrightarrow{\nu_0} [s_1] \xrightarrow{\nu_1} \dots \xrightarrow{\nu_{n-1}} [s_n]$.

Probabilistically divergent strategies on $\text{Reg}[P, \Phi]$. In the following, we use LTL notation (see, for example, [1]), which is interpreted on paths of $\text{Reg}[P, \Phi]$ in the standard way. An infinite path r of $\text{Reg}[P, \Phi]$ is *region divergent* if it satisfies the condition $\Box \Diamond \text{tick}$. Note that an infinite path r of $T[P]$ is divergent if and only if $[r]$ is region divergent. Hence $[\text{Timediv}] = \bigcup_{r \in \text{Timediv}} [r] = \{r \in \text{Path}_{\text{ful}}^{\text{Reg}[P, \Phi]} \mid r \models \Box \Diamond \text{tick}\}$ is the set of all region divergent runs (where $\models$ is the standard satisfaction for LTL properties on finite-state systems [1]). A strategy $\sigma \in \Sigma_{\text{Reg}[P, \Phi]}$ is *probabilistically region divergent* if, for all regions $R \in \text{Regions}$, we have $\text{Prob}_R^\sigma(\Box \Diamond \text{tick}) = 1$. The set of all probabilistically region divergent strategies of $\text{Reg}[P, \Phi]$ is denoted by $\Sigma_{\text{Reg}[P, \Phi]}^{\text{Pd}}$.

We can check whether there exists a probabilistically region divergent strategy of $\text{Reg}[P, \Phi]$ by computing the set of regions from which it is possible to satisfy $\Box \Diamond \text{tick}$ with probability 1, then comparing this set to Regions . Formally, we compute the set of regions of $\text{Reg}[P, \Phi]$, denoted by $\llbracket \neg \mathbb{P}_{<1}(\Box \Diamond \text{tick}) \rrbracket$, for which $R \in \llbracket \neg \mathbb{P}_{<1}(\Box \Diamond \text{tick}) \rrbracket$ if and only if there exists a strategy $\sigma \in \Sigma_{\text{Reg}[P, \Phi]}$ such that $\text{Prob}_R^\sigma(\Box \Diamond \text{tick}) = 1$. If $\llbracket \neg \mathbb{P}_{<1}(\Box \Diamond \text{tick}) \rrbracket \neq \text{Regions}$, then $\text{Reg}[P, \Phi]$ does not have a probabilistically region divergent strategy. We note that, for the region $R = (l, \alpha)$, we have $(l, \alpha) \in \llbracket \neg \mathbb{P}_{<1}(\Box \Diamond \text{tick}) \rrbracket$ if and only if there exists $\sigma \in \Sigma_{T[P]}$ such that $\text{Prob}_{(l, v)}^\sigma(\text{Timediv}) = 1$ for all $v \in \alpha$. The set $\llbracket \neg \mathbb{P}_{<1}(\Box \Diamond \text{tick}) \rrbracket$ can be computed on $\text{Reg}[P, \Phi]$ using polynomial-time algorithms for Büchi objectives of Markov decision processes [26]. In the remainder of this section, we assume that $\text{Reg}[P, \Phi]$ has at least one probabilistically region divergent strategy; that is, $\llbracket \neg \mathbb{P}_{<1}(\Box \Diamond \text{tick}) \rrbracket = \text{Regions}$. If this is not the case, we compute the sub-MDP of $\text{Reg}[P, \Phi]$ induced by $\llbracket \neg \mathbb{P}_{<1}(\Box \Diamond \text{tick}) \rrbracket$ and use it in the place of $\text{Reg}[P, \Phi]$.

PTCTL model checking with probabilistic divergence. The only formula which depends on the notion of strategy is $\mathbb{P}_{\bowtie \zeta}(\Phi_1 \mathcal{U} \Phi_2)$, and hence we consider sub-formulae of Φ of this form. We assume that, for each state $s \in S$, we have

$s \models_{\Sigma^{\text{Pd}}} \Phi_i$ if and only if $[s] \models_{\Sigma^{\text{Pd}}_{\text{Reg}[\mathbf{P}, \Phi]}} \Phi_i$ for $i \in \{1, 2\}$. From standard reasoning, for any path $r \in \text{Path}_{\text{ful}}$ of $\mathbf{T}[\mathbf{P}]$, we have $r \models_{\Sigma} \Phi_1 \mathcal{U} \Phi_2$ if and only if $[r] \models \Phi_1 \mathcal{U} \Phi_2$, where Σ is an arbitrary set of strategies.

Proposition 1. (1) Let $\sigma \in \Sigma^{\text{Pd}}_{\mathbf{P}}$ be a probabilistically divergent strategy. Then there exists a probabilistically region divergent strategy $\sigma' \in \Sigma^{\text{Pd}}_{\text{Reg}[\mathbf{P}, \Phi]}$ such that $\text{Prob}_{\sigma}^{\sigma}(\Phi_1 \mathcal{U} \Phi_2) = \text{Prob}_{\sigma'}^{\sigma'}(\Phi_1 \mathcal{U} \Phi_2)$ for all states $s \in S$. (2) Let $\sigma \in \Sigma^{\text{Pd}}_{\text{Reg}[\mathbf{P}, \Phi]}$ be a probabilistically region divergent strategy. Then there exists a probabilistically divergent strategy $\sigma' \in \Sigma^{\text{Pd}}_{\mathbf{P}}$ such that $\text{Prob}_{\sigma}^{\sigma}(\Phi_1 \mathcal{U} \Phi_2) = \text{Prob}_{\sigma'}^{\sigma'}(\Phi_1 \mathcal{U} \Phi_2)$ for all states $s \in S$.

Corollary 1. For any $s \in S$, we have $s \models_{\Sigma^{\text{Pd}}_{\mathbf{P}}} \mathbb{P}_{\bowtie \zeta}(\Phi_1 \mathcal{U} \Phi_2)$ if and only if $[s] \models_{\Sigma^{\text{Pd}}_{\text{Reg}[\mathbf{P}, \Phi]}} \mathbb{P}_{\bowtie \zeta}(\Phi_1 \mathcal{U} \Phi_2)$.

Corollary 1 follows from Proposition 1 and the semantics of PTCTL. Therefore it suffices to consider resolving properties of the form $\mathbb{P}_{\bowtie \zeta}(\Phi_1 \mathcal{U} \Phi_2)$ on $\text{Reg}[\mathbf{P}, \Phi]$. We now make a case distinction based on whether $\mathbb{P}_{\bowtie \zeta}(\Phi_1 \mathcal{U} \Phi_2)$ is of the form (A) $\mathbb{P}_{\leq \zeta}(\Phi_1 \mathcal{U} \Phi_2)$ or $\mathbb{P}_{< \zeta}(\Phi_1 \mathcal{U} \Phi_2)$, or (B) $\mathbb{P}_{\geq \zeta}(\Phi_1 \mathcal{U} \Phi_2)$ or $\mathbb{P}_{> \zeta}(\Phi_1 \mathcal{U} \Phi_2)$. In the remainder of this section, we generally omit the subscript from the sets of strategies of $\text{Reg}[\mathbf{P}, \Phi]$, and write Σ for $\Sigma_{\text{Reg}[\mathbf{P}, \Phi]}$, and Σ^{Pd} for $\Sigma^{\text{Pd}}_{\text{Reg}[\mathbf{P}, \Phi]}$.

Case (A): properties of the form $\mathbb{P}_{\leq \zeta}(\Phi_1 \mathcal{U} \Phi_2)$ or $\mathbb{P}_{< \zeta}(\Phi_1 \mathcal{U} \Phi_2)$. The following proposition states that any strategy of $\text{Reg}[\mathbf{P}, \Phi]$ can be transformed into a probabilistically region divergent strategy which assigns the same or greater probability to the satisfaction of $\Phi_1 \mathcal{U} \Phi_2$.

Proposition 2. Let $\sigma \in \Sigma$ be a strategy of $\text{Reg}[\mathbf{P}, \Phi]$ and $R \in \text{Regions}$ be a region. There exists a probabilistically region divergent strategy $\sigma' \in \Sigma^{\text{Pd}}$ such that $\text{Prob}_{\sigma}^{\sigma}(\Phi_1 \mathcal{U} \Phi_2) \leq \text{Prob}_{\sigma'}^{\sigma'}(\Phi_1 \mathcal{U} \Phi_2)$.

Proposition 2, together with the fact that $\Sigma^{\text{Pd}} \subseteq \Sigma$, establishes that there exists a probabilistically region divergent strategy which assigns the same probability to satisfying $\Phi_1 \mathcal{U} \Phi_2$ as a maximal – but not necessarily divergent – strategy of $\text{Reg}[\mathbf{P}, \Phi]$. Combining this fact with standard methods for finite-state MDPs [25, 17], we conclude that $\text{Reg}[\mathbf{P}, \Phi]$ can be used directly to compute maximal probabilities of until formulae.

Theorem 1. Let $R \in \text{Regions}$ be a region. Then $R \models_{\Sigma^{\text{Pd}}} \mathbb{P}_{\leq \zeta}(\Phi_1 \mathcal{U} \Phi_2)$ (respectively, $\mathbb{P}_{< \zeta}(\Phi_1 \mathcal{U} \Phi_2)$) if and only if $R \models_{\Sigma} \mathbb{P}_{\leq \zeta}(\Phi_1 \mathcal{U} \Phi_2)$ (respectively, $\mathbb{P}_{< \zeta}(\Phi_1 \mathcal{U} \Phi_2)$).

Case (B): properties of the form $\mathbb{P}_{\geq \zeta}(\Phi_1 \mathcal{U} \Phi_2)$ or $\mathbb{P}_{> \zeta}(\Phi_1 \mathcal{U} \Phi_2)$. To avoid overloading the subsequent notation, we consider properties in the form $\mathbb{P}_{\geq \zeta}(\neg \Phi_1 \mathcal{U} \neg \Phi_2)$ or $\mathbb{P}_{> \zeta}(\neg \Phi_1 \mathcal{U} \neg \Phi_2)$. Observe that $\neg(\neg \Phi_1 \mathcal{U} \neg \Phi_2) \equiv \Phi_2 \mathcal{U}(\Phi_1 \wedge \Phi_2) \vee \square(\neg \Phi_1 \wedge \Phi_2)$ (from classical reasoning about temporal logic). Therefore, the maximal probability over probabilistically region divergent strategies of satisfying $\Phi_2 \mathcal{U}(\Phi_1 \wedge \Phi_2) \vee \square(\neg \Phi_1 \wedge \Phi_2)$ equals 1 minus the minimal probability over probabilistically

region divergent strategies of satisfying $\neg\Phi_1\mathcal{U}\neg\Phi_2$. Hence, our aim is to compute the maximal probability over probabilistically region divergent strategies of satisfying $\Phi_2\mathcal{U}(\Phi_1 \wedge \Phi_2) \vee \Box(\neg\Phi_1 \wedge \Phi_2)$.

We introduce a notion of *time-divergent EC*. A time-divergent EC (C, D) is an EC of $\text{Reg}[\mathbf{P}, \Phi]$ such that $\text{tick} \in \text{lab}_{\text{Reg}}(R)$ for some region $R \in C$ (a similar definition is featured in [5]). For an infinite path $r \in \text{Path}_{\text{ful}}^{\text{Reg}[\mathbf{P}, \Phi]}$, let $C_r = \{R \mid \exists i \geq 0. r(i) = R\}$ and $D_r = \{(R, \nu) \mid \exists i \geq 0. R \in C_r \wedge \text{step}(r, i) = \nu\}$. Let $\text{Inf}(r) = (C_r, D_r)$. Note that a path $r \in \text{Path}_{\text{ful}}^{\text{Reg}[\mathbf{P}, \Phi]}$ of $\text{Reg}[\mathbf{P}, \Phi]$ is region divergent if and only if $\text{Inf}(r)$ is a time-divergent EC. For $C \subseteq \text{Regions}$ and $D \subseteq \rightarrow_{\text{Reg}}$, let $\text{Path}_{\text{ful}}^{(C, D)}(R) = \{r \in \text{Path}_{\text{ful}}^{\text{Reg}[\mathbf{P}, \Phi]}(R) \mid \text{Inf}(r) = (C, D)\}$. The next lemma adapts to probabilistic divergence a fundamental result for ECs [5], and states that a probabilistically region divergent strategy will be confined eventually to time-divergent ECs with probability 1.

Lemma 1. *Let $\mathcal{E}$ be the set of time-divergent ECs of $\text{Reg}[\mathbf{P}, \Phi]$, let $R \in \text{Regions}$ and let $\sigma \in \Sigma^{\text{Pd}}$. Then $\text{Prob}_R^\sigma(\bigcup_{(C, D) \in \mathcal{E}} \text{Path}_{\text{ful}}^{(C, D)}(R)) = 1$.*

Let $\mathbf{U} \subseteq \text{Regions}$ be a set of regions. The set $\mathcal{M}_{\mathbf{U}}$ of time-divergent maximal ECs within $\mathbf{U}$ can be computed as follows: first compute the set of maximal ECs of the sub-MDP of $\text{Reg}[\mathbf{P}, \Phi]$ induced by $\mathbf{U}$ by the standard maximal EC computation algorithm of [5], then include in $\mathcal{M}_{\mathbf{U}}$ only those maximal ECs with at least one region labelled by *tick*.

We compute the set $\mathcal{M}_{[\neg\Phi_1 \wedge \Phi_2]}$ of time-divergent maximal ECs within the set of states satisfying $\neg\Phi_1 \wedge \Phi_2$. Let $\mathbf{U}_{\neg\Phi_1 \wedge \Phi_2} = \bigcup_{(C, D) \in \mathcal{M}_{[\neg\Phi_1 \wedge \Phi_2]}} C$ be the set of regions corresponding to $\mathcal{M}_{[\neg\Phi_1 \wedge \Phi_2]}$. By abuse of notation, we use $\mathbf{U}_{\neg\Phi_1 \wedge \Phi_2}$ as an atomic proposition such that $R \models_{\Sigma} \mathbf{U}_{\neg\Phi_1 \wedge \Phi_2}$ if and only if $R \in \mathbf{U}_{\neg\Phi_1 \wedge \Phi_2}$. Note that, by Proposition 2, for any strategy $\sigma \in \Sigma$ and $R \in \text{Regions}$, there exists a probabilistically region divergent strategy $\sigma' \in \Sigma^{\text{Pd}}$ such that:

$$\text{Prob}_R^\sigma(\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee \mathbf{U}_{\neg\Phi_1 \wedge \Phi_2})) \leq \text{Prob}_R^{\sigma'}(\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee \mathbf{U}_{\neg\Phi_1 \wedge \Phi_2})).$$

Proposition 3. *Let $R \in \text{Regions}$ be a region and $\sigma \in \Sigma^{\text{Pd}}$ be a probabilistically region divergent strategy of $\text{Reg}[\mathbf{P}, \Phi]$. There exists a probabilistically region divergent strategy $\sigma' \in \Sigma^{\text{Pd}}$ such that:*

$$\text{Prob}_R^\sigma(\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee \mathbf{U}_{\neg\Phi_1 \wedge \Phi_2})) \leq \text{Prob}_R^{\sigma'}(\Phi_2\mathcal{U}(\Phi_1 \wedge \Phi_2) \vee \Box(\neg\Phi_1 \wedge \Phi_2)).$$

Proposition 4. *Let $R \in \text{Regions}$ be a region and $\sigma \in \Sigma^{\text{Pd}}$ be a probabilistically region divergent strategy of $\text{Reg}[\mathbf{P}, \Phi]$. Then:*

$$\text{Prob}_R^\sigma(\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee \mathbf{U}_{\neg\Phi_1 \wedge \Phi_2})) \geq \text{Prob}_R^\sigma(\Phi_2\mathcal{U}(\Phi_1 \wedge \Phi_2) \vee \Box(\neg\Phi_1 \wedge \Phi_2)).$$

The subsequent theorem then follows from Proposition 2, Proposition 3, Proposition 4, and the fact that $\Sigma^{\text{Pd}} \subseteq \Sigma$.

Theorem 2. *Let $R \in \text{Regions}$ be a region. Then $R \models_{\Sigma^{\text{Pd}}} \mathbb{P}_{\geq \zeta}(\neg\Phi_1\mathcal{U}\neg\Phi_2)$ (respectively, $\mathbb{P}_{> \zeta}(\neg\Phi_1\mathcal{U}\neg\Phi_2)$) if and only if $R \models_{\Sigma} \mathbb{P}_{\leq 1-\zeta}(\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee \mathbf{U}_{\neg\Phi_1 \wedge \Phi_2}))$ (respectively, $\mathbb{P}_{< 1-\zeta}(\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee \mathbf{U}_{\neg\Phi_1 \wedge \Phi_2}))$).*

Fig. 3. Minimal reachability probabilities under probabilistic and strict divergence.

The following theorem is a consequence of Corollary 1, Theorem 1, Theorem 2 and the following facts: the size of $\text{Reg}[\mathbf{P}, \Phi]$ is exponential in the size of $\mathbf{P}$ and Φ ; computing the maximal probability of a formula of the form $\Phi_1 \mathcal{U} \Phi_2$ on MDPs is in polynomial-time [25, 17]; model checking PTA against properties of the form $\neg \mathbb{P}_{<1}(\Diamond a)$ is EXPTIME-hard [15].

Theorem 3. *Let $\mathbf{P}$ be a PTA and Φ be a formula of PTCTL. Then the problem of computing the set $\llbracket \Phi \rrbracket$ for $\mathbf{P}$ under probabilistically divergent strategies is EXPTIME-complete.*

4 Strict Divergence

We now extend the model-checking algorithm for probabilistically divergent strategies to provide a model-checking algorithm for strictly divergent adversaries. Given a PTA $\mathbf{P}$, a strategy $\sigma \in \Sigma_{\mathbf{T}[\mathbf{P}]}$ is *strictly divergent* if, for all states $s \in S$, we have $\text{Path}_{\text{ful}}^\sigma(s) \subseteq \text{Timediv}$. The set of all probabilistically divergent strategies of $\mathbf{P}$ is denoted by $\Sigma_{\mathbf{P}}^{\text{Sd}}$. An example of the difference between probabilistic and strict divergence for maximal reachability probabilities (or maximal probabilities of satisfying formulae of the form $\Phi_1 \mathcal{U} \Phi_2$) has been presented in the introduction. For minimal reachability (or $\Phi_1 \mathcal{U} \Phi_2$) probabilities, note that, in the PTA of Figure 3, the minimum probability of reaching l_2 from l_1 is 0 under probabilistically divergent strategies, but is 1 under strictly divergent strategies.

Strictly divergent strategies on $\text{Reg}[\mathbf{P}, \Phi]$. A strategy $\sigma \in \Sigma_{\text{Reg}[\mathbf{P}, \Phi]}$ is *strictly region divergent* if, for all regions $R \in \text{Regions}$, all paths $r \in \text{Path}_{\text{ful}}^\sigma(R)$ are region divergent. The set of strictly region divergent strategies of $\text{Reg}[\mathbf{P}, \Phi]$ is denoted by $\Sigma_{\text{Reg}[\mathbf{P}, \Phi]}^{\text{Sd}}$. We generally write Σ instead of $\Sigma_{\text{Reg}[\mathbf{P}, \Phi]}$, and Σ^{Sd} instead of $\Sigma_{\text{Reg}[\mathbf{P}, \Phi]}^{\text{Sd}}$.

Similarly to the case of Section 3, we can check whether there exists a strictly region divergent strategy of $\text{Reg}[\mathbf{P}, \Phi]$ by computing the set of regions from which it is possible to satisfy $\Box \Diamond \text{tick}$ on all paths. For this purpose, we compute the set of regions satisfying the ATL [27] formula $\langle\langle N \rangle\rangle(\Box \Diamond \text{tick})$, where $\text{Reg}[\mathbf{P}, \Phi]$ is interpreted as a turn-based game with 2 players: player N corresponds to nondeterministic choice between transitions from a region, whereas player P refers to choice between probabilistic alternatives corresponding to a transition. Then the formula $\langle\langle N \rangle\rangle(\Box \Diamond \text{tick})$ expresses the property that player N has the aim of ensuring region time divergence, regardless of the choices of

player P . Formally, $\llbracket \langle N \rangle (\Box \Diamond tick) \rrbracket = \{R \in \text{Regions} \mid \exists \sigma \in \Sigma_{\text{Reg}[P, \Phi]} \cdot \forall r \in \text{Path}_{\text{ful}}^\sigma(R) \cdot r \models_\Sigma \Box \Diamond tick\}$. We note that, for the region $R = (l, \alpha)$, we have $(l, \alpha) \in \llbracket \langle N \rangle (\Box \Diamond tick) \rrbracket$ if and only if there exists $\sigma \in \Sigma_{\text{TP}[P]}$ such that r is divergent for all paths $r \in \text{Path}_{\text{ful}}^\sigma(l, v)$, for all $v \in \alpha$. In order to compute $\llbracket \langle N \rangle (\Box \Diamond tick) \rrbracket$, we rely on standard methods for obtaining the winning states in 2-player turn-based games with Büchi objectives [26]. In the remainder of this section, we assume that $\text{Reg}[P, \Phi]$ has at least one strictly region divergent strategy; that is, $\llbracket \langle N \rangle (\Box \Diamond tick) \rrbracket = \text{Regions}$. If this is not the case, we compute the sub-MDP of $\text{Reg}[P, \Phi]$ induced by $\llbracket \langle N \rangle (\Box \Diamond tick) \rrbracket$ and use the new sub-MDP in the place of $\text{Reg}[P, \Phi]$.

PTCTL model checking with strict divergence. We now describe a PTCTL model-checking algorithm for the semantics under strictly divergent strategies. The mechanism that we add to the PTCTL model-checking algorithm in order to cater for strict divergence is inspired by similar results of [17, 18], and takes the form of the following: a set T^{max} of regions of $\text{Reg}[P, \Phi]$ is computed from which it is guaranteed that there exists an optimal (maximal or minimal probability), strictly divergent strategy. From regions not in T^{max} , there does not exist such a strategy. However, from regions not in T^{max} , we show that we can approximate arbitrarily closely an optimal, probabilistically divergent strategy.

We first present an analogue of Proposition 1 adapted to strict divergence.

Proposition 5. (1) Let $\sigma \in \Sigma_P^{\text{Sd}}$ be a strictly divergent strategy. Then there exists a strictly region divergent strategy $\sigma' \in \Sigma_{\text{Reg}[P, \Phi]}^{\text{Sd}}$ such that $\text{Prob}_s^\sigma(\Phi_1 \mathcal{U} \Phi_2) = \text{Prob}_{[s]}^{\sigma'}(\Phi_1 \mathcal{U} \Phi_2)$ for all states $s \in S$. (2) Let $\sigma \in \Sigma_{\text{Reg}[P, \Phi]}^{\text{Sd}}$ be a strictly region divergent strategy. Then there exists a strictly divergent strategy $\sigma' \in \Sigma_P^{\text{Sd}}$ such that $\text{Prob}_{[s]}^{\sigma'}(\Phi_1 \mathcal{U} \Phi_2) = \text{Prob}_s^{\sigma'}(\Phi_1 \mathcal{U} \Phi_2)$ for all states $s \in S$.

Corollary 2. For any $s \in S$, we have $s \models_{\Sigma_P^{\text{Sd}}} \mathbb{P}_{\bowtie \zeta}(\Phi_1 \mathcal{U} \Phi_2)$ if and only if $[s] \models_{\Sigma_{\text{Reg}[P, \Phi]}^{\text{Sd}}} \mathbb{P}_{\bowtie \zeta}(\Phi_1 \mathcal{U} \Phi_2)$.

Therefore, as in Section 3, it suffices to resolve properties of the form $\mathbb{P}_{\bowtie \zeta}(\Phi_1 \mathcal{U} \Phi_2)$ on $\text{Reg}[P, \Phi]$. We make a case distinction based on whether $\mathbb{P}_{\bowtie \zeta}(\Phi_1 \mathcal{U} \Phi_2)$ is of the form (A) $\mathbb{P}_{\leq \zeta}(\Phi_1 \mathcal{U} \Phi_2)$ or $\mathbb{P}_{< \zeta}(\Phi_1 \mathcal{U} \Phi_2)$, or (B) $\mathbb{P}_{\geq \zeta}(\Phi_1 \mathcal{U} \Phi_2)$ or $\mathbb{P}_{> \zeta}(\Phi_1 \mathcal{U} \Phi_2)$.

Case (A): properties of the form $\mathbb{P}_{\leq \zeta}(\Phi_1 \mathcal{U} \Phi_2)$ or $\mathbb{P}_{< \zeta}(\Phi_1 \mathcal{U} \Phi_2)$. Recall that the example in the introduction shows that a PTA may exhibit a strategy with a certain probability of reaching a location under probabilistic divergence, but, under strict divergence, there may not exist a strategy with the same probability. However, we show that strictly region divergent strategies can approximate from below the probability of satisfying $\Phi_1 \mathcal{U} \Phi_2$ of an arbitrary strategy on $\text{Reg}[P, \Phi]$.

Lemma 2. Let $R \in \text{Regions}$ and let $\sigma \in \Sigma$ be a strategy of $\text{Reg}[P, \Phi]$. Then, for every $n \in \mathbb{N}$, there exists a strictly region divergent strategy $\sigma_n \in \Sigma^{\text{Sd}}$ such that $\text{Prob}_R^{\sigma_n}(\Phi_1 \mathcal{U} \Phi_2) \geq \text{Prob}_R^\sigma(\Phi_1 \mathcal{U} \Phi_2) - \frac{1}{n}$.

The intuition underlying the proof of the lemma is that we construct σ_n to behave the same as σ on all paths whose length does not exceed some constant c_n which depends on n . From the last regions of paths of length c_n , the strategy σ_n then behaves as a strictly region divergent strategy. From Lemma 1 of [28], given σ and n , such a constant can be found such that the lemma holds.

Given that it is possible to approximate arbitrarily closely using strictly divergent strategies the probability of satisfying a property $\Phi_1 \mathcal{U} \Phi_2$ of an arbitrary strategy, the case $\mathbb{P}_{\leq \zeta}(\Phi_1 \mathcal{U} \Phi_2)$ is not of interest: a maximal arbitrary strategy satisfies $\Phi_1 \mathcal{U} \Phi_2$ with probability greater than ζ if and only if there exists a strictly divergent strategy satisfying $\Phi_1 \mathcal{U} \Phi_2$ with probability greater than ζ . Hence, we concentrate on the case of $\mathbb{P}_{< \zeta}(\Phi_1 \mathcal{U} \Phi_2)$.

Let $R \in \text{Regions}$, and $p_R^{\max}(\Phi_1 \mathcal{U} \Phi_2) = \max_{\sigma \in \Sigma} \text{Prob}_R^\sigma(\Phi_1 \mathcal{U} \Phi_2)$. Following [17, 18], for each region $R \in \text{Regions}$ of $\text{Reg}[\mathbf{P}, \Phi]$, we define the set $\text{Max}(R, \Phi_1 \mathcal{U} \Phi_2)$. If $R \in \text{Regions} \setminus \llbracket \Phi_1 \rrbracket$, then we let $\text{Max}(R, \Phi_1 \mathcal{U} \Phi_2) = \{(R', \nu) \in \rightarrow_{\text{Reg}} \mid R = R'\}$. If $R \in \llbracket \Phi_1 \rrbracket$, then we let $\text{Max}(R, \Phi_1 \mathcal{U} \Phi_2)$ equal:

$$\left\{ (R, \nu) \in \rightarrow_{\text{Reg}} \mid p_R^{\max}(\Phi_1 \mathcal{U} \Phi_2) = \sum_{R'' \in \text{Regions}} \nu(R'') \cdot p_{R''}^{\max}(\Phi_1 \mathcal{U} \Phi_2) \right\}.$$

We use $\mathbf{M}^{\max}$ to denote the MDP obtained by removing from $\text{Reg}[\mathbf{P}, \Phi]$ all transitions which are not in $\text{Max}(\cdot, \Phi_1 \mathcal{U} \Phi_2)$. Formally, let $\mathbf{M}^{\max} = (\text{Regions}, \rightarrow_{\text{Reg}}^{\max}, \text{lab}_{\text{Reg}})$, where $\rightarrow_{\text{Reg}}^{\max} = \bigcup_{R \in \text{Regions}} \text{Max}(R, \Phi_1 \mathcal{U} \Phi_2)$.

Let $\mathbf{T} = \llbracket \Phi_2 \rrbracket \cup \text{Regions} \setminus \text{Regions}^+(\Phi_1, \Phi_2)$, and let $\llbracket \exists \Diamond \mathbf{T} \rrbracket = \{R \in \text{Regions} \mid \exists r \in \text{Path}_{\text{fin}}(R) \text{ s.t. } \text{last}(r) \in \mathbf{T}\}$. We apply the following algorithm to $\mathbf{M}^{\max}$.

1. Let $\mathbf{U}$ equal Regions and $\mathbf{M}$ equal $\mathbf{M}^{\max}$.
 2. Repeat the following:
 - (a) Let $\mathbf{U}$ equal either $\llbracket \langle \langle N \rangle \rangle (\Diamond \text{tick}) \rrbracket$ or $\llbracket \exists \Diamond \mathbf{T} \rrbracket$, computed on $\mathbf{M}$.
 - (b) Compute the sub-MDP of $\mathbf{M}$ induced by $\mathbf{U}$, and call this sub-MDP $\mathbf{M}$.
- Until $\mathbf{M}$ cannot be changed by the above.

Let $\mathbf{T}^{\max}$ be the set of regions of the MDP obtained on termination of the algorithm. The strategies of this MDP do not select non-optimal transitions (from the definition of $\mathbf{M}^{\max}$), and can be both strictly region divergent *and* reach $\mathbf{T}$ with probability 1. We state this formally in relation to strategies of $\text{Reg}[\mathbf{P}, \Phi]$ in the following lemma.

Lemma 3. *Let $R \in \text{Regions}$. Then $R \in \mathbf{T}^{\max}$ if and only if there exists a strategy $\sigma \in \Sigma$ such that (1) for each $r \in \text{Path}_{\text{fin}}^\sigma(R)$, we have $\sigma(r) \in \text{Max}(\text{last}(r), \Phi_1 \mathcal{U} \Phi_2)$, (2) $\sigma \in \Sigma^{\text{Sd}}$ (σ is strictly region divergent) and (3) $\text{Prob}_R^\sigma(\Diamond \mathbf{T}) = 1$.*

We note the importance of reaching $\mathbf{T}$ with probability 1: strategies satisfying this requirement do not idle in a cycle of transitions. Such cycles, in which a strategy does not attempt to reach $\llbracket \Phi_2 \rrbracket$, may be locally optimal (in the sense that transitions of $\mathbf{M}^{\max}$ are taken), but will be globally sub-optimal: another strategy which does not cycle, but attempts to reach $\llbracket \Phi_2 \rrbracket$, will correspond to a higher probability of satisfying $\Phi_1 \mathcal{U} \Phi_2$.

Our next task is to show that, from regions in $\mathbf{T}^{\max}$, there exists a strictly region divergent strategy with the same probability of satisfying $\Phi_1 \mathcal{U} \Phi_2$ as for an optimal strategy which is not necessarily strictly region divergent.

Proposition 6. *Let $R \in \mathbf{T}^{\max}$ and $\sigma \in \Sigma^{\text{Sd}}$ be such that $\text{Prob}_R^\sigma(\Diamond \top) = 1$ and, for each $r \in \text{Path}_{\text{fin}}^\sigma(R)$, we have $\sigma(r) \in \text{Max}(\text{last}(r), \Phi_1 \mathcal{U} \Phi_2)$. Then $p_R^{\max}(\Phi_1 \mathcal{U} \Phi_2) = \text{Prob}_R^\sigma(\Phi_1 \mathcal{U} \Phi_2)$.*

Next, we show that, from regions not in $\mathbf{T}^{\max}$, it is not possible to find strictly region divergent strategies which obtain the probability of satisfying $\Phi_1 \mathcal{U} \Phi_2$ computed over arbitrary strategies on $\text{Reg}[\mathbf{P}, \Phi]$.

Lemma 4. *Let $R \in \text{Regions} \setminus \mathbf{T}^{\max}$. Then, for $\sigma \in \Sigma^{\text{Sd}}$, we have $p_R^{\max}(\Phi_1 \mathcal{U} \Phi_2) > \text{Prob}_R^\sigma(\Phi_1 \mathcal{U} \Phi_2)$.*

The combination of Lemma 2, Proposition 6 and Lemma 4 allows us to obtain the following result.

Theorem 4. *Let $R \in \text{Regions}$ be a region. Then:*

$$\begin{aligned} R \models_{\Sigma^{\text{Sd}}} \mathbb{P}_{\leq \zeta}(\Phi_1 \mathcal{U} \Phi_2) &\Leftrightarrow p_R^{\max}(\Phi_1 \mathcal{U} \Phi_2) \leq \zeta \\ R \models_{\Sigma^{\text{Sd}}} \mathbb{P}_{< \zeta}(\Phi_1 \mathcal{U} \Phi_2) &\Leftrightarrow \begin{cases} p_R^{\max}(\Phi_1 \mathcal{U} \Phi_2) < \zeta & \text{if } R \in \mathbf{T}^{\max} \\ p_R^{\max}(\Phi_1 \mathcal{U} \Phi_2) \leq \zeta & \text{otherwise.} \end{cases} \end{aligned}$$

Case (B): properties of the form $\mathbb{P}_{\geq \zeta}(\Phi_1 \mathcal{U} \Phi_2)$ or $\mathbb{P}_{> \zeta}(\Phi_1 \mathcal{U} \Phi_2)$. Case (B) has similarities with previous results in the paper in the following ways. Firstly, analogously to Section 3, we use the equivalence $\neg(\neg\Phi_1 \mathcal{U} \neg\Phi_2) \equiv \Phi_2 \mathcal{U}(\Phi_1 \wedge \Phi_2) \vee \Box(\neg\Phi_1 \wedge \Phi_2)$, and then consider the maximal probability over strictly region divergent strategies of satisfying the formula $\Phi_2 \mathcal{U}(\Phi_1 \wedge \Phi_2) \vee \Box(\neg\Phi_1 \wedge \Phi_2)$. This resulting probability corresponds to the 1 minus the minimal probability of strictly region divergent strategies satisfying $\neg\Phi_1 \mathcal{U} \neg\Phi_2$. Secondly, again as in Section 3, we compute the probability $p_R^{\max}(\Phi_2 \mathcal{U}((\Phi_1 \wedge \Phi_2) \vee \mathbf{U}_{\neg\Phi_1 \wedge \Phi_2}))$: we show that this probability is equal to the supremum of the probabilities of satisfying $\Phi_2 \mathcal{U}(\Phi_1 \wedge \Phi_2) \vee \Box(\neg\Phi_1 \wedge \Phi_2)$ over strictly region divergent strategies. Thirdly, analogously to Case (A) of this section, it is possible that no strictly region divergent strategy attains this supremum.

We first observe that, by Proposition 4 and the fact that any strictly region divergent strategy is also a probabilistically region divergent strategy, for any $R \in \text{Regions}$ and $\sigma \in \Sigma^{\text{Sd}}$, we have $\text{Prob}_R^\sigma(\Phi_2 \mathcal{U}(\Phi_1 \wedge \Phi_2) \vee \Box(\neg\Phi_1 \wedge \Phi_2)) \leq p_R^{\max}(\Phi_2 \mathcal{U}((\Phi_1 \wedge \Phi_2) \vee \mathbf{U}_{\neg\Phi_1 \wedge \Phi_2}))$. We now show that the probabilities assigned by strictly region divergent strategies to $\Phi_2 \mathcal{U}(\Phi_1 \wedge \Phi_2) \vee \Box(\neg\Phi_1 \wedge \Phi_2)$ approximate from below the maximal probability of satisfying $\Phi_2 \mathcal{U}((\Phi_1 \wedge \Phi_2) \vee \mathbf{U}_{\neg\Phi_1 \wedge \Phi_2})$ over arbitrary strategies.

Lemma 5. *Let $R \in \text{Regions}$. Then, for every $n \in \mathbb{N}$, there exists a strictly region divergent strategy $\sigma_n \in \Sigma^{\text{Sd}}$ such that:*

$$\text{Prob}_R^{\sigma_n}(\Phi_2 \mathcal{U}(\Phi_1 \wedge \Phi_2) \vee \Box(\neg\Phi_1 \wedge \Phi_2)) \geq p_R^{\max}(\Phi_2 \mathcal{U}((\Phi_1 \wedge \Phi_2) \vee \mathbf{U}_{\neg\Phi_1 \wedge \Phi_2})) - \frac{1}{n}$$

The intuition underlying the proof of the lemma is as follows. First note that, from a region in a time-divergent MEC of $\text{Reg}[\mathbf{P}, \Phi]$, a strictly region divergent strategy can confine itself to that MEC with probability arbitrarily close to 1. However, a strategy which confines itself with probability 1 to a time-divergent MEC will not necessarily be strictly region divergent: the strategy may exhibit paths (with probability 0) which are not time divergent. Using these facts, we define a strategy $\tilde{\sigma}$ in the following way. Let σ^{Pd} be a probabilistically divergent strategy such that $\text{Prob}_R^{\sigma^{\text{Pd}}}(\Phi_2 \mathcal{U}((\Phi_1 \wedge \Phi_2) \vee \mathbf{U}_{\neg\Phi_1 \wedge \Phi_2})) = p_R^{\max}(\Phi_2 \mathcal{U}((\Phi_1 \wedge \Phi_2) \vee \mathbf{U}_{\neg\Phi_1 \wedge \Phi_2}))$. For all finite paths on which a region satisfying $(\neg\Phi_1 \wedge \neg\Phi_2)$, $(\Phi_1 \wedge \Phi_2)$ or $\mathbf{U}_{\neg\Phi_1 \wedge \Phi_2}$ is not reached, the strategy $\tilde{\sigma}$ behaves in the same way as σ^{Pd} . Upon reaching a region satisfying $(\neg\Phi_1 \wedge \neg\Phi_2)$ or $(\Phi_1 \wedge \Phi_2)$, the strategy $\tilde{\sigma}$ then subsequently behaves as an arbitrary strictly region divergent strategy. Upon reaching a region labelled by $\mathbf{U}_{\neg\Phi_1 \wedge \Phi_2}$, the strategy $\tilde{\sigma}$ confines itself, with probability greater than $1 - \frac{1}{n}$, to the time-divergent MEC reached (recall from Section 3 that $\mathbf{U}_{\neg\Phi_1 \wedge \Phi_2}$ labels regions in time-divergent MECs); on leaving the time-divergent MEC, the strategy $\tilde{\sigma}$ then behaves as an arbitrary strictly region divergent strategy. This construction permits us to define a strategy $\tilde{\sigma}$ for which the probability $\text{Prob}_R^{\tilde{\sigma}}(\Phi_2 \mathcal{U}(\Phi_1 \wedge \Phi_2) \vee \square(\neg\Phi_1 \wedge \Phi_2))$ is greater than $p_R^{\max}(\Phi_2 \mathcal{U}((\Phi_1 \wedge \Phi_2) \vee \mathbf{U}_{\neg\Phi_1 \wedge \Phi_2})) - \frac{1}{n}$. Note that $\tilde{\sigma}$ may not be strictly region divergent, because, before reaching regions satisfying $(\neg\Phi_1 \wedge \neg\Phi_2)$, $(\Phi_1 \wedge \Phi_2)$ or $\mathbf{U}_{\neg\Phi_1 \wedge \Phi_2}$, the probabilistically region divergent strategy σ^{Pd} may make choices generating paths which are not region divergent (although such paths have probability 0, because σ^{Pd} is probabilistically region divergent). However, taking into account the probabilities assigned by $\tilde{\sigma}$ to being confined in the time-divergent MECs of $\text{Reg}[\mathbf{P}, \Phi]$, this case can be dealt with in a similar manner to that of Lemma 2 to result in a strictly region divergent strategy.

We now proceed to define an algorithm for deciding, given $R \in \text{Regions}$, whether there exists a strictly region divergent strategy $\sigma \in \Sigma^{\text{Sd}}$ such that $\text{Prob}_R^{\sigma}(\Phi_2 \mathcal{U}(\Phi_1 \wedge \Phi_2) \vee \square(\neg\Phi_1 \wedge \Phi_2)) = p_R^{\max}(\Phi_2 \mathcal{U}((\Phi_1 \wedge \Phi_2) \vee \mathbf{U}_{\neg\Phi_1 \wedge \Phi_2}))$. In the following, as in Section 3, it will be useful to consider strategies which, after a certain path prefix, confine themselves to certain ECs. Recall that the notion of time-divergent EC is not sufficient, because a strategy which confines itself to a time-divergent EC with probability 1 may exhibit paths (with probability 0) which are not time divergent. Hence, to reason about ECs in which a strategy can confine itself with probability 1 *and* let time diverge on all paths, we define strictly-divergent ECs. Formally, a *strictly-divergent EC* (C, D) is an EC of $\text{Reg}[\mathbf{P}, \Phi]$ such that, for all regions $R \in C$, we have $R \in \llbracket \langle N \rangle (\Diamond tick) \rrbracket$ in the sub-MDP (C, D) (hence the strategy witnessing $\Diamond tick$ for all paths from R is a strategy of (C, D)). Intuitively, a strategy can guarantee strict region divergence and remain within a strictly-divergent EC (C, D) by choosing transitions according to a strategy of the sub-MDP (C, D) which witnesses $\langle N \rangle (\Diamond tick)$; then, after a *tick*-region is visited, the strategy again chooses transitions according to a strategy which, starting from the current region, witnesses $\langle N \rangle (\Diamond tick)$, and so on. For all paths of this strategy, a *tick*-region is visited infinitely often, and hence the strategy is strictly region divergent.

Let $U \subseteq \text{Regions}$. The set of strictly-divergent maximal ECs within the set U can be computed by the following algorithm.

1. Compute the set $\mathcal{M}_U$ of maximal ECs of $\text{Reg}[P, \Phi]$ within U .¹ Let $\mathcal{M} = \mathcal{M}_U$.
 2. Repeat the following:
 - (a) Remove some (C, D) from $\mathcal{M}$.
 - (b) Compute $\llbracket \langle N \rangle (\Diamond tick) \rrbracket$ obtained from the sub-MDP (C, D) .
 - (c) Compute the maximal ECs $(C_1, D_1), \dots, (C_n, D_n)$ of the sub-MDP (C, D) induced by $\llbracket \langle N \rangle (\Diamond tick) \rrbracket$, and add them to $\mathcal{M}$.
- Until $\mathcal{M}$ cannot be changed by the above iteration.

At the termination of the algorithm, the set $\mathcal{M}$ will be the set of strictly-divergent maximal ECs of $\text{Reg}[P, \Phi]$ induced by U . We use $\mathcal{S}_U$ to denote this set.

We then follow the approach of Section 3 by computing the set $\mathcal{S}_{\llbracket \neg \Phi_1 \wedge \Phi_2 \rrbracket}$ of strictly-divergent maximal ECs within the set of states satisfying $\neg \Phi_1 \wedge \Phi_2$. Let $V_{\neg \Phi_1 \wedge \Phi_2} = \bigcup_{(C, D) \in \mathcal{S}_{\llbracket \neg \Phi_1 \wedge \Phi_2 \rrbracket}} C$ be the set of regions corresponding to $\mathcal{S}_{\llbracket \neg \Phi_1 \wedge \Phi_2 \rrbracket}$. We use $V_{\neg \Phi_1 \wedge \Phi_2}$ as an atomic proposition such that $R \models_{\Sigma} V_{\neg \Phi_1 \wedge \Phi_2}$ if and only if $R \in V_{\neg \Phi_1 \wedge \Phi_2}$.

We also define the following set T_V of regions:

$$T_V = \llbracket (\Phi_1 \wedge \Phi_2) \vee V_{\neg \Phi_1 \wedge \Phi_2} \rrbracket \cup \text{Regions} \setminus \text{Regions}^+(\Phi_2, (\Phi_1 \wedge \Phi_2) \vee V_{\neg \Phi_1 \wedge \Phi_2}).$$

Recall the definition of $T^{\max}$ in Case (A), which was defined with regard to the set T of regions. Here we define $T_V^{\max}$ in a similar way, apart from the fact that $T_V^{\max}$ is defined with regard to T_V . Then let $\widetilde{T_V^{\max}}$ be the set of regions R satisfying the following two conditions:

1. $R \in T_V^{\max}$;
2. $p_R^{\max}(\Phi_2 \mathcal{U}((\Phi_1 \wedge \Phi_2) \vee V_{\neg \Phi_1 \wedge \Phi_2})) = p_R^{\max}(\Phi_2 \mathcal{U}((\Phi_1 \wedge \Phi_2) \vee U_{\neg \Phi_1 \wedge \Phi_2}))$.

We now present the key result of this subsection, which gives a condition for the existence of a strictly region divergent strategy attaining the probability $p_R^{\max}(\Phi_2 \mathcal{U}((\Phi_1 \wedge \Phi_2) \vee U_{\neg \Phi_1 \wedge \Phi_2}))$.

Proposition 7. *Let $R \in \text{Regions}$ be a region. Then $R \in \widetilde{T_V^{\max}}$ if and only if there exists $\sigma \in \Sigma^{\text{Sd}}$ such that:*

$$\text{Prob}_R^{\sigma}(\Phi_2 \mathcal{U}(\Phi_1 \wedge \Phi_2) \vee \Box(\neg \Phi_1 \wedge \Phi_2)) = p_R^{\max}(\Phi_2 \mathcal{U}((\Phi_1 \wedge \Phi_2) \vee U_{\neg \Phi_1 \wedge \Phi_2})).$$

Proposition 7 can be understood as follows. If $R \in \widetilde{T_V^{\max}}$, then a strictly region divergent strategy σ can be defined on paths from R in the following way. For finite paths along which regions satisfying either $(\neg \Phi_1 \wedge \neg \Phi_2)$, $(\Phi_1 \wedge \Phi_2)$ or $V_{\neg \Phi_1 \wedge \Phi_2}$ have not been reached, we let σ behave like a strategy satisfying $\Phi_2 \mathcal{U}((\Phi_1 \wedge \Phi_2) \vee V_{\neg \Phi_1 \wedge \Phi_2})$ with the maximal probability (by $R \in T_V^{\max}$ and Proposition 6 it is possible to construct σ so that no choice is taken by σ along

¹ Recall that an algorithm for this purpose can be found in [5].

these finite paths which prevents σ from being strictly region divergent). For finite paths along which regions satisfying $V_{\neg\Phi_1 \wedge \Phi_2}$ have been reached, the strategy σ confines itself with probability 1 to the strictly-divergent MEC which it has reached (recall that $V_{\neg\Phi_1 \wedge \Phi_2}$ labels regions in strictly-divergent MECs), thereby guaranteeing the satisfaction of $\Box(\neg\Phi_1 \wedge \Phi_2)$ and the divergence of time on all subsequent paths. Instead, for finite paths along which regions satisfying $(\neg\Phi_1 \wedge \neg\Phi_2)$ or $(\Phi_1 \wedge \Phi_2)$ have been reached, the strategy σ subsequently behaves as an arbitrary strictly region divergent strategy. The construction of σ ensures that $\text{Prob}_R^\sigma(\Phi_2\mathcal{U}(\Phi_1 \wedge \Phi_2) \vee \Box(\neg\Phi_1 \wedge \Phi_2)) = p_R^{\max}(\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee V_{\neg\Phi_1 \wedge \Phi_2}))$. Then condition 2 of the definition of $\widetilde{T}_V^{\max}$ gives us the required equality.

Next we consider the case in which $R \notin \widetilde{T}_V^{\max}$. The case in which $R \notin T_V^{\max}$ is similar to that of Lemma 4, and hence we concentrate on the case in which $p_R^{\max}(\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee V_{\neg\Phi_1 \wedge \Phi_2})) < p_R^{\max}(\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee U_{\neg\Phi_1 \wedge \Phi_2}))$.² We consider the class of strictly region divergent strategies of the following form. A strategy σ of the class behaves like a strategy satisfying $\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee U_{\neg\Phi_1 \wedge \Phi_2})$ with the maximal probability until a region satisfying either $(\neg\Phi_1 \wedge \neg\Phi_2)$, $(\Phi_1 \wedge \Phi_2)$ or $U_{\neg\Phi_1 \wedge \Phi_2}$ is reached. For finite paths along which regions satisfying $(\neg\Phi_1 \wedge \neg\Phi_2)$ or $(\Phi_1 \wedge \Phi_2)$ have been reached, the strategy σ subsequently behaves as an arbitrary strictly region divergent strategy. For finite paths along which regions labelled by $U_{\neg\Phi_1 \wedge \Phi_2}$ have been reached, the strategy σ subsequently confines itself to the reached time-divergent MEC with probability arbitrarily close to 1. From $p_R^{\max}(\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee V_{\neg\Phi_1 \wedge \Phi_2})) < p_R^{\max}(\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee U_{\neg\Phi_1 \wedge \Phi_2}))$, it must be the case that, on leaving the time-divergent MEC, the property $\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee V_{\neg\Phi_1 \wedge \Phi_2})$ cannot be satisfied with probability equal to 1. From this fact we conclude that the probability of satisfying $\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee \Box(\neg\Phi_1 \wedge \Phi_2))$ associated with σ is strictly less than $p_R^{\max}(\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee U_{\neg\Phi_1 \wedge \Phi_2}))$. Observe that this class of strictly region divergent strategies results in “maximal probabilities” of satisfying $\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee \Box(\neg\Phi_1 \wedge \Phi_2))$: strictly region divergent strategies not belonging to this class will also have probabilities of satisfying $\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee \Box(\neg\Phi_1 \wedge \Phi_2))$ strictly less than $p_R^{\max}(\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee U_{\neg\Phi_1 \wedge \Phi_2}))$.

The combination of Lemma 5, Proposition 7 and Theorem 4 then gives us the following.

Theorem 5. *Let $R \in \text{Regions}$ be a region. Then:*

$$\begin{aligned} R \models_{\Sigma^{\text{sd}}} \mathbb{P}_{\geq \zeta}(\neg\Phi_1\mathcal{U}\neg\Phi_2) &\Leftrightarrow 1 - p_R^{\max}(\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee U_{\neg\Phi_1 \wedge \Phi_2})) \geq \zeta \\ R \models_{\Sigma^{\text{sd}}} \mathbb{P}_{> \zeta}(\neg\Phi_1\mathcal{U}\neg\Phi_2) &\Leftrightarrow \begin{cases} 1 - p_R^{\max}(\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee U_{\neg\Phi_1 \wedge \Phi_2})) > \zeta & \text{if } R \in \widetilde{T}_V^{\max} \\ 1 - p_R^{\max}(\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee U_{\neg\Phi_1 \wedge \Phi_2})) \geq \zeta & \text{otherwise.} \end{cases} \end{aligned}$$

From Corollary 2, Theorem 4 and Theorem 5, and from the fact that the procedures for computing $\widetilde{T}_V^{\max}$ and the set of strictly-divergent maximal ECs presented above run in time polynomial in the size of $\text{Reg}[\mathbf{P}, \Phi]$, we then obtain the following result.

² The case $p_R^{\max}(\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee V_{\neg\Phi_1 \wedge \Phi_2})) > p_R^{\max}(\Phi_2\mathcal{U}((\Phi_1 \wedge \Phi_2) \vee U_{\neg\Phi_1 \wedge \Phi_2}))$ is not possible because $V_{\neg\Phi_1 \wedge \Phi_2} \subseteq U_{\neg\Phi_1 \wedge \Phi_2}$ by definition.

Theorem 6. *Let P be a PTA and Φ be a formula of PTCTL. Then the problem of computing the set $\llbracket \Phi \rrbracket$ for P under strictly divergent strategies is EXPTIME-complete.*

5 Conclusions

We have presented optimal model-checking algorithms for PTA for two notions of time divergence. As in previous methods [6], the algorithms rely on a computation of probabilities of satisfying an until property on a finite-state MDP resulting from the classical region graph construction. For probabilistic divergence and properties of the form $\mathbb{P}_{\geq \zeta}(\Phi_1 \mathcal{U} \Phi_2)$ and $\mathbb{P}_{> \zeta}(\Phi_1 \mathcal{U} \Phi_2)$, the algorithms rely also on the computation of maximal ECs in which a strategy can ensure time divergence with probability 1. For strict divergence and for properties of the form $\mathbb{P}_{< \zeta}(\Phi_1 \mathcal{U} \Phi_2)$, we compute the set of states in which the maximal until probability can be obtained by a strictly-divergent strategy; for all other states, strictly-divergent strategies can approximate arbitrarily closely the maximal probability. A similar technique can be used for properties of the form $\mathbb{P}_{> \zeta}(\Phi_1 \mathcal{U} \Phi_2)$ in combination with the computation of maximal ECs in which a strategy can ensure time divergence on all paths. The techniques of this paper are useful when considering models in which there are no lower time bounds on structural loops: these include abstract models of embedded controllers in which lower bounds on certain reaction times are left unspecified. In future work we intend to extend our notions of divergence to controller synthesis, and to consider symbolic, zone-based algorithms for strict divergence.

Acknowledgements. I thank François Laroussinie and Arnaud Sangnier for discussions concerning Case (B) of Section 4.

References

1. Clarke, E.M., Grumberg, O., Peled, D.: Model checking. MIT Press (1999)
2. Alur, R., Dill, D.L.: A theory of timed automata. TCS **126**(2) (1994) 183–235
3. Alur, R., Courcoubetis, C., Dill, D.L.: Model-checking for probabilistic real-time systems. In: Proc. ICALP’91. Volume 510 of LNCS., Springer (1991) 115–136
4. Hansson, H.A., Jonsson, B.: A logic for reasoning about time and reliability. Formal Aspects of Computing **6**(5) (1994) 512–535
5. de Alfaro, L.: Formal verification of probabilistic systems. PhD thesis, Stanford University, Department of Computer Science (1997)
6. Kwiatkowska, M., Norman, G., Segala, R., Sproston, J.: Automatic verification of real-time systems with discrete probability distributions. TCS **286** (2002) 101–150
7. Baier, C., Haverkort, B., Hermanns, H., Katoen, J.P.: Model-checking algorithms for continuous-time Markov chains. TSE **29**(6) (2003) 524–541
8. Donatelli, S., Haddad, S., Sproston, J.: Model checking stochastic and timed properties with CSL^{TA}. TSE **35**(2) (2009) 224–240
9. Chen, T., Han, T., Katoen, J.P., Mereacre, A.: Quantitative model checking of continuous-time Markov chains against timed automata specifications. In: Proc. LICS’09, IEEE (2009)

10. Jensen, H.E.: Model checking probabilistic real time systems. In: Proc. of the 7th Nordic Work. on Progr. Theory, Chalmers Institute of Technology (1996) 247–261
11. Kwiatkowska, M., Norman, G., Parker, D., Sproston, J.: Performance analysis of probabilistic timed automata using digital clocks. *FMSD* **29** (2006) 33–78
12. Alur, R., Courcoubetis, C., Dill, D.L.: Model-checking in dense real-time. *I & C* **104**(1) (1993) 2–34
13. Henzinger, T., Nicollin, X., Sifakis, J., Yovine, S.: Symbolic model checking for real-time systems. *I & C* **111**(2) (1994) 193–244
14. Tripakis, S., Yovine, S., Bouajjani, A.: Checking timed Büchi automata emptiness efficiently. *FMSD* **26**(3) (2005) 267–292
15. Laroussinie, F., Sproston, J.: State explosion in almost-sure probabilistic reachability. *IPL* **102**(6) (2007) 236–241
16. Kwiatkowska, M., Norman, G., Sproston, J., Wang, F.: Symbolic model checking for probabilistic timed automata. *I & C* **205**(7) (2007) 1027–1077
17. Baier, C., Kwiatkowska, M.: Model checking for a probabilistic branching time logic with fairness. *Dist. Comp.* **11**(3) (1998) 125–155
18. Baier, C.: On the algorithmic verification of probabilistic systems (1998) Habilitation thesis, Universität Mannheim.
19. Chatterjee, K., Henzinger, T., Prabhu, V.: Trading infinite memory for uniform randomness in timed games. In: Proc. HSCC’08. Volume 4981 of LNCS., Springer (2008) 87–100
20. Beauquier, D.: On probabilistic timed automata. *TCS* **292**(1) (2003) 65–84
21. Kemeny, J.G., Snell, J.L., Knapp, A.W.: Denumerable Markov Chains. 2nd edn. Graduate Texts in Mathematics. Springer (1976)
22. Jurdziński, M., Laroussinie, F., Sproston, J.: Model checking probabilistic timed automata with one or two clocks. *LMCS* **4**(3) (2008) 1–28
23. Alur, R., Henzinger, T.: Real-time system = discrete system + clock variables. *STTT* **1** (1997) 86–109
24. de Alfaro, L., Faella, M., Henzinger, T., Majumdar, R., Stoelinga, M.: The element of surprise in timed games. In: Proc. CONCUR’03. Volume 2761 of LNCS., Springer (2003) 144–158
25. Bianco, A., de Alfaro, L.: Model checking of probabilistic and nondeterministic systems. In: Proc. FSTTCS’95. Volume 1026 of LNCS., Springer (1995) 499–513
26. Chatterjee, K., Jurdziński, M., Henzinger, T.: Simple stochastic parity games. In: Proc. CSL’03. Volume 2803 of LNCS., Springer (2003) 100–113
27. Alur, R., Henzinger, T., Kupferman, O.: Alternating-time temporal logic. *JACM* **49** (2002) 672–713
28. Fecher, H., Huth, M., Piterman, N., Wagner, D.: Hintikka games for PCTL on labeled Markov chains. In: Proc. QEST’08, IEEE (2008) 169–178