Efficient Analysis of Probabilistic Programs with an Unbounded Counter

Tomáš Brázdil[1]⋆, Stefan Kiefer[2]⋆⋆, and Antonín Kučera[1]⋆

[1] Faculty of Informatics, Masaryk University, Czech Republic.
{brazdil,kucera}@fi.muni.cz
[2] Department of Computer Science, University of Oxford, United Kingdom.
stefan.kiefer@cs.ox.ac.uk

Abstract. We show that a subclass of infinite-state probabilistic programs that can be modeled by probabilistic one-counter automata (pOC) admits an efficient quantitative analysis. In particular, we show that the expected termination time can be approximated up to an arbitrarily small relative error with polynomially many arithmetic operations, and the same holds for the probability of all runs that satisfy a given ω-regular property. Further, our results establish a powerful link between pOC and martingale theory, which leads to fundamental observations about quantitative properties of runs in pOC. In particular, we provide a "divergence gap theorem", which bounds a positive non-termination probability in pOC away from zero.

1 Introduction

In this paper we aim at designing *efficient* algorithms for analyzing basic properties of probabilistic programs operating on unbounded data domains that can be abstracted into a non-negative integer counter. Consider, e.g., the following recursive program *TreeEval* which evaluates a given AND-OR tree, i.e., a tree whose root is an AND node, all descendants of AND nodes are either leaves or OR nodes, and all descendants of OR nodes are either leaves or AND nodes.

```
procedure AND(node)
if node is a leaf
   then return node.value
else
   for each successor s of node do
      if OR(s) = 0 then return 0
   end for
   return 1
end if
```

```
procedure OR(node)
if node is a leaf
   then return node.value
else
   for each successor s of node do
      if AND(s) = 1 then return 1
   end for
   return 0
end if
```

⋆ Tomáš Brázdil and Antonín Kučera are supported by the research center Institute for Theoretical Computer Science (ITI, project No. 1M0545) and by the Czech Science Foundation, grant No. P202/10/1469.

⋆⋆ Stefan Kiefer is supported by a postdoctoral fellowship of the German Academic Exchange Service (DAAD).

Note that the program *TreeEval* evaluates a subtree only when necessary. In general, we cannot say anything about its expected termination time. If the input tree is infinite, the program may not even terminate, i.e., it may fail to evaluate the root node. Now assume that we *do* have some knowledge about the actual input domain of the program, which might have been gathered empirically:

- an AND node has about a descendants on average;
- an OR node has about o descendants on average;
- the length of a branch is b on average;
- the probability that a leaf evaluates to 1 is z.

Further, let us assume that the actual number of descendants and the actual length of a branch are *geometrically* distributed (which is a reasonably good approximation in many cases). Hence, the probability that an AND node has *exactly* n descendants is $(1 - x_a)^{n-1} x_a$ with $x_a = \frac{1}{a}$. Under these assumption, the behaviour of *TreeEval* is well-defined in the probabilistic sense, and we may ask the following questions:

1) Does the program terminate with probability one? If not, what is the termination probability?
2) If we restrict ourselves to terminating runs, what is the expected termination time?

These questions are not trivial, and at first glance it is not clear how to approach them. Apart of the expected termination time, which is a fundamental characteristic of terminating runs, we are also interested in the properties on *non-terminating* runs, specified by linear-time logics or automata on infinite words. Here, we ask for the probability of all runs satisfying a given linear-time property. Using the results of this paper, answers to such questions can be computed *efficiently* for a large class of programs, including the program *TreeEval*. More precisely, the first question about the probability of termination can be answered using the existing results [14]; the original contributions of this paper are efficient algorithms for computing answers to the remaining questions.

The abstract class of probabilistic programs considered in this paper corresponds to *probabilistic one-counter automata (pOC)*. Informally, a pOC has finitely many control states $p, q, \ldots$ that can store global data, and a single non-negative counter that can be incremented, decremented, and tested for zero. The dynamics of a given pOC is described by finite sets of *positive* and *zero* rules of the form $p \xrightarrow{x,c}_{>0} q$ and $p \xrightarrow{x,c}_{=0} q$, respectively, where p, q are control states, x is the *probability* of the rule, and $c \in \{-1, 0, 1\}$ is the *counter change* which must be non-negative in zero rules. A *configuration* $p(i)$ is given by the current control state p and the current counter value i. If i is positive/zero, then positive/zero rules can be applied to $p(i)$ in the natural way. Thus, every pOC determines an infinite-state Markov chain where states are the configurations and transitions are determined by the rules. As an example, consider a pOC model of the program *TreeEval*. We use the counter to abstract the stack of activation records. Since the procedures AND and OR alternate regularly in the stack, we keep just the current stack height in the counter, and maintain the "type" of the current procedure in the finite control (when we increase or decrease the counter, the "type" is swapped). The return values of the two procedures are also stored in the finite control. Thus, we obtain the following pOC model with 6 control states and 12 positive rules (zero rules are irrelevant and hence not shown).

/* if we have a leaf, return 0 or 1 */
$(and,init) \xrightarrow{yz,-1} (or,return,1)$,
$(and,init) \xrightarrow{y(1-z),-1} (or,return,0)$
/* otherwise, call OR */
$(and,init) \xrightarrow{(1-y),1} (or,\ init)$
/* if OR returns 1, call another OR? */
$(and,return,1) \xrightarrow{(1-x_a),1} (or,init)$
$(and,return,1) \xrightarrow{x_a,-1} (or,return,1)$
/* if OR returns 0, return 0 immediately */
$(and,return,0) \xrightarrow{1,-1} (or,return,0)$

/* if we have a leaf, return 0 or 1 */
$(or,init) \xrightarrow{yz,-1} (and,return,1)$,
$(or,init) \xrightarrow{y(1-z),-1} (and,return,0)$
/* otherwise, call AND */
$(or,init) \xrightarrow{(1-y),1} (and,init)$
/* if AND returns 0, call another AND? */
$(or,return,0) \xrightarrow{(1-x_o),1} (and,\ init)$
$(or,return,0) \xrightarrow{x_o,-1} (and,return,0)$
/* if AND returns 1, return 1 immediately */
$(or,return,1) \xrightarrow{1,-1} (and,return,1)$

The initial configuration is $(and,init)(1)$, and the pOC terminates either in $(or,return,0)(0)$ or $(or,return,1)(0)$, which corresponds to evaluating the input tree to 0 and 1, respectively. We set $x_a := 1/a$, $x_o := 1/o$ and $y := 1/b$ in order to obtain the average numbers a, o, b from the beginning.

As we already indicated, pOC can model recursive programs operating on unbounded data structures such as trees, queues, or lists, assuming that the structure can be faithfully abstracted into a counter. Let us note that modeling general recursive programs requires more powerful formalisms such as *probabilistic pushdown automata (pPDA)* [12] or *recursive Markov chains (RMC)* [17]. However, as it is mentioned below, pPDA and RMC do not admit *efficient* quantitative analysis for fundamental reasons. Hence, we must inevitably sacrifice a part of pPDA modeling power to gain efficiency in algorithmic analysis, and pOC seem to be a good candidate.

The relevance of pOC is not limited just to recursive programs. As observed in [14], pOC are equivalent, in a well-defined sense, to discrete-time *Quasi-Birth-Death processes (QBDs)*, a well-established stochastic model that has been deeply studied since late 60s. Thus, the applicability of pOC extends to queuing theory, performance evaluation, etc., where QBDs are considered as a fundamental formalism. Very recently, games over (probabilistic) one-counter automata, also called "energy games", were considered in several independent works [9, 10, 4, 3]. The study is motivated by optimizing the use of resources (such as energy) in modern computational devices.

Previous work. In [12, 17], it has been shown that the vector of termination probabilities in pPDA and RMC is the least solution of an effectively constructible system of quadratic equations. The termination probabilities may take irrational values, but can be effectively approximated up to an arbitrarily small absolute error $\varepsilon > 0$ in polynomial space by employing the decision procedure for the existential fragment of Tarski algebra (i.e., first order theory of the reals) [8]. Due to the results of [17], it is possible to approximate termination probabilities in pPDA and RMC "iteratively" by using the decomposed Newton's method. However, this approach may need exponentially many iterations of the method before it starts to produce one bit of precision per iteration [19]. Further, any non-trivial approximation of the non-termination probabilities is at least as hard as the SquareRootSum problem [17], whose exact complexity is a long-standing open question in exact numerical computations (the best known upper bound for SquareRootSum is PSPACE). Computing termination probabilities in pPDA and RMC up to a given *relative* error $\varepsilon > 0$, which is more relevant from the point of

view of this paper, is *provably* infeasible because the termination probabilities can be doubly-exponentially small in the size of a given pPDA or RMC [17].

The expected termination time and the expected reward per transition in pPDA and RMC has been studied in [13]. In particular, it has been shown that the tuple of expected termination times is the least solution of an effectively constructible system of linear equations, where the (products of) termination probabilities are used as coefficients. Hence, the equational system can be represented only symbolically, and the corresponding approximation algorithm again employs the decision procedure for Tarski algebra. There also other results for pPDA and RMC, which concern model-checking problems for linear-time [15, 16] and branching-time [7] logics, long-run average properties [5], discounted properties of runs [2], etc.

Our contribution. In this paper, we build on the previously established results for pPDA and RMC, and on the recent results of [14] where is shown that the decomposed Newton method of [19] can be used to compute termination probabilities in pOC up to a given *relative* error $\varepsilon > 0$ in time which is *polynomial* in the size of pOC and $\log(1/\varepsilon)$, assuming the unit-cost rational arithmetic RAM (i.e., Blum-Shub-Smale) model of computation. Adopting the same model, we show the following:

1. The expected termination time in a pOC $\mathscr{A}$ is computable up to an arbitrarily small relative error $\varepsilon > 0$ in time polynomial in $|\mathscr{A}|$ and $\log(1/\varepsilon)$. Actually, we can even compute the expected termination time up to an arbitrarily small *absolute* error, which is a better estimate because the expected termination time is always at least 1.
2. The probability of all runs in a pOC $\mathscr{A}$ satisfying an ω-regular property encoded by a deterministic Rabin automaton $\mathcal{R}$ is computable up to an arbitrarily small relative error $\varepsilon > 0$ in time polynomial in $|\mathscr{A}|$, $|\mathcal{R}|$, and $\log(1/\varepsilon)$.

The crucial step towards obtaining these results is the construction of a suitable *martingale* for a given pOC, which allows to apply powerful results of martingale theory (such as the optional stopping theorem or Azuma's inequality, see, e.g., [20, 21]) to the quantitative analysis of pOC. In particular, we use this martingale to establish the crucial *divergence gap theorem* in Section 4, which bounds a positive divergence probability in pOC away from 0. The divergence gap theorem is indispensable in analysing properties of non-terminating runs, and together with the constructed martingale provide generic tools for designing efficient approximation algorithms for other interesting quantitative properties of pOC.

Although our algorithms have polynomial worst-case complexity, the obtained bounds look complicated and it is not immediately clear whether the algorithms are practically usable. Therefore, we created a simple experimental implementation which computes the expected termination time for pOC, and used this tool to analyse the pOC model of the program *TreeEval*. The details are given in Section 5.

Due to space limits, we could not include most of the proofs into the main body of the paper. These can be found in a full version of this paper [6].

2 Definitions

We use $\mathbb{Z}$, $\mathbb{N}$, $\mathbb{N}_0$, $\mathbb{Q}$, and $\mathbb{R}$ to denote the set of all integers, positive integers, non-negative integers, rational numbers, and real numbers, respectively. Let $\delta > 0$, $x \in \mathbb{Q}$,

and $y \in \mathbb{R}$. We say that x approximates y up to a relative error δ, if either $y \neq 0$ and $|x-y|/|y| \leq \delta$, or $x = y = 0$. Further, we say that x approximates y up to an absolute error δ if $|x-y| \leq \delta$. We use standard notation for intervals, e.g., $(0,1]$ denotes $\{x \in \mathbb{R} \mid 0 < x \leq 1\}$.

Given a finite set Q, we regard elements of $\mathbb{R}^Q$ as vectors over Q. We use boldface symbols like $\boldsymbol{u}, \boldsymbol{v}$ for vectors. In particular we write $\mathbf{1}$ for the vector whose entries are all 1. Similarly, elements of $\mathbb{R}^{Q \times Q}$ are regarded as square matrices.

Let $\mathcal{V} = (V, \to)$, where V is a non-empty set of vertices and $\to \subseteq V \times V$ a *total* relation (i.e., for every $v \in V$ there is some $u \in V$ such that $v \to u$). The reflexive and transitive closure of $\to$ is denoted by $\to^*$. A *finite path* in $\mathcal{V}$ of *length* $k \geq 0$ is a finite sequence of vertices $v_0, \ldots, v_k$, where $v_i \to v_{i+1}$ for all $0 \leq i < k$. The length of a finite path w is denoted by $\mathit{length}(w)$. A *run* in $\mathcal{V}$ is an infinite sequence w of vertices such that every finite prefix of w is a finite path in $\mathcal{V}$. The individual vertices of w are denoted by $w(0), w(1), \ldots$ The sets of all finite paths and all runs in $\mathcal{V}$ are denoted by $\mathit{FPath}_{\mathcal{V}}$ and $\mathit{Run}_{\mathcal{V}}$, respectively. The sets of all finite paths and all runs in $\mathcal{V}$ that start with a given finite path w are denoted by $\mathit{FPath}_{\mathcal{V}}(w)$ and $\mathit{Run}_{\mathcal{V}}(w)$, respectively. Let $U \subseteq V$. We say that U is *strongly connected* if $v \to^+ u$ for all $v, u \in U$ (here $v \to^+ u$ if there is a path of length greater than 1 from v to u). Further, we say that U is a *strongly connected component (SCC)* if $U \neq \emptyset$ is a maximal strongly connected subset of V, and U is a *bottom SCC (BSCC)* if for every $u \in U$ and every $u \to v$ we have that $v \in U$.

We assume familiarity with basic notions of probability theory, e.g., *probability space*, *random variable*, or the *expected value*. As usual, a *probability distribution* over a finite or countably infinite set X is a function $f : X \to [0,1]$ such that $\sum_{x \in X} f(x) = 1$. We call f *positive* if $f(x) > 0$ for every $x \in X$, and *rational* if $f(x) \in \mathbb{Q}$ for every $x \in X$.

Definition 1. *A* Markov chain *is a triple $\mathcal{M} = (S, \to, \mathit{Prob})$ where S is a finite or countably infinite set of* states*, $\to \subseteq S \times S$ is a total* transition relation*, and Prob is a function that assigns to each state $s \in S$ a positive probability distribution over the outgoing transitions of s. As usual, we write $s \xrightarrow{x} t$ when $s \to t$ and x is the probability of $s \to t$.*

A Markov chain $\mathcal{M}$ can be also represented by its *transition matrix* $M \in [0,1]^{S \times S}$, where $M_{s,t} = 0$ if $s \not\to t$, and $M_{s,t} = x$ if $s \xrightarrow{x} t$.

To every $s \in S$ we associate the probability space $(\mathit{Run}_{\mathcal{M}}(s), \mathcal{F}, \mathcal{P})$ of runs starting at s, where $\mathcal{F}$ is the σ-field generated by all *basic cylinders*, $\mathit{Run}_{\mathcal{M}}(w)$, where w is a finite path starting at s, and $\mathcal{P} : \mathcal{F} \to [0,1]$ is the unique probability measure such that $\mathcal{P}(\mathit{Run}_{\mathcal{M}}(w)) = \prod_{i=1}^{\mathit{length}(w)} x_i$ where $w(i-1) \xrightarrow{x_i} w(i)$ for every $1 \leq i \leq \mathit{length}(w)$. If $\mathit{length}(w) = 0$, we put $\mathcal{P}(\mathit{Run}_{\mathcal{M}}(w)) = 1$.

Definition 2. *A* probabilistic one-counter automaton (pOC) *is a tuple, $\mathscr{A} = (Q, \delta^{=0}, \delta^{>0}, P^{=0}, P^{>0})$, where*

- *Q is a finite set of* states*,*
- *$\delta^{>0} \subseteq Q \times \{-1, 0, 1\} \times Q$ and $\delta^{=0} \subseteq Q \times \{0, 1\} \times Q$ are the sets of* positive *and* zero rules *such that each $p \in Q$ has an outgoing positive rule and an outgoing zero rule;*
- *$P^{>0}$ and $P^{=0}$ are* probability assignments*, assigning to each $p \in Q$ a positive rational probability distribution over the outgoing rules in $\delta^{>0}$ and $\delta^{=0}$, resp., of p.*

In the following, we often write $p \xrightarrow{x,c}_{=0} q$ to denote that $(p,c,q) \in \delta^{=0}$ and $P^{=0}(p,c,q) = x$, and similarly $p \xrightarrow{x,c}_{>0} q$ to denote that $(p,c,q) \in \delta^{>0}$ and $P^{>0}(p,c,q) = x$. The size of $\mathscr{A}$, denoted by $|\mathscr{A}|$, is the length of the string which represents $\mathscr{A}$, where the probabilities of rules are written in binary. A *configuration* of $\mathscr{A}$ is an element of $Q \times \mathbb{N}_0$, written as $p(i)$. To $\mathscr{A}$ we associate an infinite-state Markov chain $\mathcal{M}_{\mathscr{A}}$ whose states are the configurations of $\mathscr{A}$, and for all $p, q \in Q$, $i \in \mathbb{N}$, and $c \in \mathbb{N}_0$ we have that $p(0) \xrightarrow{x} q(c)$ iff $p \xrightarrow{x,c}_{=0} q$, and $p(i) \xrightarrow{x} q(c)$ iff $p \xrightarrow{x,c-i}_{>0} q$. For all $p, q \in Q$, let

- $Run_{\mathscr{A}}(p{\downarrow}q)$ be the set of all runs in $\mathcal{M}_{\mathscr{A}}$ initiated in $p(1)$ that visit $q(0)$ and the counter stays positive in all configurations preceding this visit;
- $Run_{\mathscr{A}}(p{\uparrow})$ be the set of all runs in $\mathcal{M}_{\mathscr{A}}$ initiated in $p(1)$ where the counter never reaches zero.

We omit the "$\mathscr{A}$" in $Run_{\mathscr{A}}(p{\downarrow}q)$ and $Run_{\mathscr{A}}(p{\uparrow})$ when it is clear from the context, and we use $[p{\downarrow}q]$ and $[p{\uparrow}]$ to denote the probability of $Run(p{\downarrow}q)$ and $Run(p{\uparrow})$, respectively. Observe that $[p{\uparrow}] = 1 - \sum_{q \in Q} [p{\downarrow}q]$ for every $p \in Q$.

At various places in this paper we rely on the following proposition proven in [14] (recall that we adopt the unit-cost rational arithmetic RAM model of computation):

Proposition 3. *Let $\mathscr{A} = (Q, \delta^{=0}, \delta^{>0}, P^{=0}, P^{>0})$ be a pOC, and $p, q \in Q$.*

- *The problem whether $[p{\downarrow}q] > 0$ is decidable in polynomial time.*
- *If $[p{\downarrow}q] > 0$, then $[p{\downarrow}q] \geq x_{\min}^{|Q|^3}$, where $x_{\min}$ is the least (positive) probability used in the rules of $\mathscr{A}$.*
- *The probability $[p{\downarrow}q]$ can be approximated up to an arbitrarily small relative error $\varepsilon > 0$ in a time polynomial in $|\mathscr{A}|$ and $\log(1/\varepsilon)$.*

Due to Proposition 3, the set $T^{>0}$ of all pairs $(p,q) \in Q \times Q$ satisfying $[p{\downarrow}q] > 0$ is computable in polynomial time.

3 Expected Termination Time

In this section we give an efficient algorithm which approximates the expected termination time in pOC up to an arbitrarily small relative (or even absolute) error $\varepsilon > 0$.

For the rest of this section, we fix a pOC $\mathscr{A} = (Q, \delta^{=0}, \delta^{>0}, P^{=0}, P^{>0})$. For all $p, q \in Q$, let $R_{p{\downarrow}q} : Run(p(1)) \to \mathbb{N}_0$ be a random variable which to a given run w assigns either the least k such that $w(k) = q(0)$, or 0 if there is no such k. If $(p,q) \in T^{>0}$, we use $E(p{\downarrow}q)$ to denote the conditional expectation $\mathbb{E}[R_{p{\downarrow}q} \mid Run(p{\downarrow}q)]$. Note that $E(p{\downarrow}q)$ can be finite even if $[p{\downarrow}q] < 1$.

The first problem we have to deal with is that the expectation $E(p{\downarrow}q)$ can be infinite, as illustrated by the following example.

Example 4. Consider a simple pOC with only one control state p and two positive rules $(p,-1,p)$ and $(p,1,p)$ that are both assigned the probability $1/2$. Then $[p{\downarrow}p] = 1$, and due to results of [13], $E(p{\downarrow}p)$ is the least solution (in $\mathbb{R}^+ \cup \{\infty\}$) of the equation $x = 1/2 + 1/2(1 + 2x)$, which is ∞.

We proceed as follows. First, we show that the problem whether $E(p{\downarrow}q) = \infty$ is decidable in polynomial time (Section 3.1). Then, we eliminate all infinite expectations, and show how to approximate the finite values of the remaining $E(p{\downarrow}q)$ up to a given absolute (and hence also relative) error $\varepsilon > 0$ efficiently (Section 3.2).

3.1 Finiteness of the expected termination time

In this subsection we exhibit conditions that, given $(p,q) \in T^{>0}$, allow to decide in polynomial time whether $E(p{\downarrow}q)$ is finite. To state these conditions, we need some notions. Define sets $Pre^*(q(0))$ and $Post^*(p(1))$, where

- $Pre^*(q(0))$ consists of all $r(k)$ that can reach $q(0)$ along a run w in $\mathcal{M}_{\mathscr{A}}$ such that the counter stays positive in all configurations preceding the visit to $q(0)$;
- $Post^*(p(1))$ consists of all $r(k)$ that can be reached from $p(1)$ along a run w in $\mathcal{M}_{\mathscr{A}}$ where the counter stays positive in all configurations preceding the visit to $r(k)$.

Note that $q(0) \in Pre^*(q(0))$ and $p(1) \in Post^*(p(1))$. Further, define a finite-state Markov chain $\mathcal{X}$ with Q as set of states, and transition matrix $A \in [0,1]^{Q\times Q}$ given by $A_{p,q} = \sum_{(p,c,q)\in\delta^{>0}} P^{>0}(p,c,q)$. Given a BSCC $\mathscr{B}$ of $\mathcal{X}$, let $\boldsymbol{\alpha} \in (0,1]^{\mathscr{B}}$ be the *invariant distribution* of $\mathscr{B}$, i.e., the unique (row) vector satisfying $\boldsymbol{\alpha}A = \boldsymbol{\alpha}$ and $\boldsymbol{\alpha}\mathbf{1} = 1$ (see, e.g., [18, Theorem 5.1.2]). Further, we define the (column) vector $\boldsymbol{s} \in \mathbb{R}^{\mathscr{B}}$ of *expected counter changes* by $\boldsymbol{s}_p = \sum_{(p,c,q)\in\delta^{>0}} P^{>0}(p,c,q)\cdot c$ and the *trend* $t \in \mathbb{R}$ of $\mathscr{B}$ by $t = \boldsymbol{\alpha}\boldsymbol{s}$. Intuitively, the trend is the average counter increase per step. Note that t is easily computable in polynomial time. Now we can state the following theorem:

Theorem 5. *Let $(p,q) \in T^{>0}$. Let $x_{\min}$ denote the smallest nonzero probability in A. Then we have:*

(A) if q is not in a BSCC of $\mathcal{X}$, then $E(p{\downarrow}q) \le 5|Q| \,/\, x_{\min}^{|Q|+|Q|^3}$;
(B) if q is in a BSCC $\mathscr{B}$ of $\mathcal{X}$, then:
  (a) if $Pre^*(q(0)) \cap Post^*(p(1)) \cap \mathscr{B}\times\mathbb{N}$ is a finite set, then $E(p{\downarrow}q) \le 20|Q|^3/x_{\min}^{4|Q|^3}$;
  (b) if $Pre^*(q(0)) \cap Post^*(p(1)) \cap \mathscr{B}\times\mathbb{N}$ is an infinite set, then:
    (1) if $\mathscr{B}$ has trend $t \neq 0$, then $E(p{\downarrow}q) \le 85000|Q|^6/(x_{\min}^{5|Q|+|Q|^3}\cdot t^4)$;
    (2) if $\mathscr{B}$ has trend $t = 0$, then $E(p{\downarrow}q)$ is infinite.

One can check in polynomial time which case of Theorem 5 applies. In particular, due to [11], there are finite-state automata constructible in polynomial time recognizing the sets $Pre^*(q(0))$ and $Post^*(p(1))$. Hence, we can efficiently compute a finite-state automaton $\mathcal{F}$ recognizing the set $Pre^*(q(0)) \cap Post^*(p(1)) \cap \mathscr{B}\times\mathbb{N}$ and check whether the language accepted by $\mathcal{F}$ is finite. Thus we have the following corollary:

Corollary 6. *Let $(p,q) \in T^{>0}$. The problem whether $E(p{\downarrow}q)$ is finite is decidable in polynomial time.*

In the rest of this subsection we sketch a qualitative proof for Theorem 5; i.e., we sketch why $E(p{\downarrow}q)$ is infinite only in case (B.b.2). First assume case (A), i.e., q is not in a BSCC of $\mathcal{X}$. Then for all $s(\ell) \in Post^*(p(1))$, where $\ell \ge |Q|$, we have that $s(\ell)$ can reach a configuration outside $Pre^*(q(0))$ in at most $|Q|$ transitions. It follows that the probability of performing a path from $p(1)$ to $q(0)$ of length i decays exponentially in i, and hence $E(p{\downarrow}q)$ is finite.

Next assume case (B.a), i.e., $\mathscr{B}$ is a BSCC and $C := Pre^*(q(0)) \cap Post^*(p(1)) \cap \mathscr{B}\times\mathbb{N}$ is a finite set. It is easy to show that the expected time for a run in $Run(p{\downarrow}q)$ to reach $\mathscr{B}$ is finite. Once the run has reached $\mathscr{B}$ it basically moves within a Markov chain on C.

By assumption, C is finite (which implies, by a pumping argument, that $|C| \leq 3|Q|^3$). Consequently, after the run has reached $\mathscr{B}$, it reaches $q(0)$ in finite expected time.

Case (B.b) requires new non-trivial techniques. For the sake of simplicity, *from now on we assume that* $Q = \mathscr{B}$ (the general case requires only slight modifications of the arguments presented below). We employ a generic observation which connects the study of pOC to martingale theory. Recall that a stochastic process $m^{(0)}, m^{(1)}, \ldots$ is a martingale if, for all $i \in \mathbb{N}$, $\mathbb{E}(|m^{(i)}|) < \infty$, and $\mathbb{E}(m^{(i+1)} \mid m^{(1)}, \ldots, m^{(i)}) = m^{(i)}$ almost surely. Let us fix an initial configuration $r(c) \in Q \times \mathbb{N}$. Our aim is to construct a suitable martingale over $Run(r(c))$. Let $p^{(i)}$ and $c^{(i)}$ be random variables which to every run $w \in Run(r(c))$ assign the control state and the counter value of the configuration $w(i)$, respectively. Note that if the vector $\boldsymbol{s}$ of expected counter changes is constant, i.e., $\boldsymbol{s} = \mathbf{1} \cdot t$ where t is the trend of $\mathcal{X}$, then we can define a martingale $m^{(0)}, m^{(1)}, \ldots$ simply by

$$m^{(i)} = \begin{cases} c^{(i)} - i \cdot t & \text{if } c^{(j)} \geq 1 \text{ for all } 0 \leq j < i; \\ m^{(i-1)} & \text{otherwise.} \end{cases}$$

Since $\boldsymbol{s}$ is generally not constant, we might try to "compensate" the difference among the individual control states by a suitable vector $\boldsymbol{v} \in \mathbb{R}^Q$. The next proposition shows that this is indeed possible.

Proposition 7. *There is a vector* $\boldsymbol{v} \in \mathbb{R}^Q$ *such that the stochastic process* $m^{(0)}, m^{(1)}, \ldots$ *defined by*

$$m^{(i)} = \begin{cases} c^{(i)} + \boldsymbol{v}_{p^{(i)}} - i \cdot t & \textit{if } c^{(j)} \geq 1 \textit{ for all } 0 \leq j < i; \\ m^{(i-1)} & \textit{otherwise} \end{cases}$$

is a martingale, where t *is the trend of* $\mathcal{X}$.

Moreover, the vector $\boldsymbol{v}$ *satisfies* $\boldsymbol{v}_{\max} - \boldsymbol{v}_{\min} \leq 2|Q|/x_{\min}^{|Q|}$, *where* $x_{\min}$ *is the smallest positive transition probability in* $\mathcal{X}$, *and* $\boldsymbol{v}_{\max}$ *and* $\boldsymbol{v}_{\min}$ *are the maximal and the minimal components of* $\boldsymbol{v}$, *respectively.*

Due to Proposition 7, powerful results of martingale theory become applicable to pOC. In this paper, we use the constructed martingale to establish statements (iii) and (iv) of Theorem 5, by employing Azuma's inequality and the optional stopping theorem (see [20, 21]). We also use the martingale to prove the crucial *divergence gap theorem* in Section 4. The range of possible applications of Proposition 7 is of course wider.

Assume now case (B.b.1), i.e., $t \neq 0$. For every $i \in \mathbb{N}$, let $Run(p{\downarrow}q, i)$ be the set of all $w \in Run(p{\downarrow}q)$ that visit $q(0)$ in i transitions, and let $[p{\downarrow}q, i]$ be the probability of $Run(p{\downarrow}q, i)$. We first show that there are $0 < a < 1$ and $h \in \mathbb{N}$ such that for all $i \geq h$ we have that $[p{\downarrow}q, i] \leq a^i$. Consider the martingale $m^{(0)}, m^{(1)}, \ldots$ over $Run(p(1))$ as defined in Proposition 7. A relatively straightforward computation reveals that for sufficiently large $h \in \mathbb{N}$ and all $i \geq h$ we have the following: If $t < 0$, then $[p{\downarrow}q, i] \leq \mathcal{P}\left(m^{(i)} - m^{(0)} \geq (i/2) \cdot (-t)\right)$, and if $t > 0$, then $[p{\downarrow}q, i] \leq \mathcal{P}\left(m^{(0)} - m^{(i)} \geq (i/2) \cdot t\right)$. In each step, the martingale value changes by at most $\boldsymbol{v}_{\max} - \boldsymbol{v}_{\min} + t + 1$, where $\boldsymbol{v}$ is from Proposition 7. Hence, by applying Azuma's inequality (see [21]) we obtain the following (for all $t \neq 0$ and $i \geq h$):

$$[p{\downarrow}q, i] \quad \leq \quad \exp\left(-\frac{(i/2)^2 t^2}{2i(\boldsymbol{v}_{\max} - \boldsymbol{v}_{\min} + t + 1)^2}\right) \quad = \quad a^i$$

Here $a = \exp\left(-t^2 \,/\, 8(\boldsymbol{v}_{\max} - \boldsymbol{v}_{\min} + t + 1)^2\right)$. It follows that

$$E(p{\downarrow}q) \quad = \quad \sum_{i=1}^{\infty} i \cdot \frac{[p{\downarrow}q, i]}{[p{\downarrow}q]} \quad \leq \quad \frac{1}{[p{\downarrow}q]} \left(\sum_{i=1}^{h-1} i \cdot [p{\downarrow}q, i] + \sum_{i=h}^{\infty} i \cdot a^i \right) \quad < \quad \infty \,.$$

Finally assume case (B.b.2), i.e., $t = 0$. We need to show that $E(p{\downarrow}q) = \infty$. Let us introduce some notation. For every $k \in \mathbb{N}_0$, let $Q(k)$ be the set of all configurations where the counter value equals k. Let $p, q \in Q$ and $\ell, k \in \mathbb{N}_0$, where $\ell > k$. An *honest path from* $p(\ell)$ *to* $q(k)$ is a finite path w from $p(\ell)$ to $q(k)$ such that the counter stays above k in all configurations of w except for the last one. We use $hpath(p(\ell), Q(k))$ to denote the set of all honest paths from $p(\ell)$ to some $q(k) \in Q(k)$. For a given $P \subseteq hpath(p(\ell), Q(k))$, the *expected length of an honest path in* P is defined as $\sum_{w \in P} \mathcal{P}(Run(w)) \cdot length(w)$. Using the martingale from Proposition 7 we show the following:

Proposition 8. *If $Pre^*(q(0))$ is infinite, then almost all runs initiated in an arbitrary configuration reach $Q(0)$. Moreover, there is $k_1 \in \mathbb{N}$ such that, for all $\ell \geq k_1$, the expected length of an honest path from $r(\ell)$ to $Q(0)$ is infinite.*

Proof (Sketch). Assume that $Pre^*(q(0))$ is infinite. The fact that almost all runs initiated in an arbitrary configuration reach $Q(0)$ follows from results of [4].

Consider an initial configuration $r(\ell)$ with $\ell + \boldsymbol{v}_r > \boldsymbol{v}_{\max}$. We will show that the expected length of an honest path from $r(\ell)$ to $Q(0)$ is infinite; i.e., we can take $k_1 := \lceil \boldsymbol{v}_{\max} - \boldsymbol{v}_{\min} + 1 \rceil$. Consider the martingale $m^{(0)}, m^{(1)}, \ldots$ defined in Proposition 7 over $Run(r(\ell))$. Note that as $t = 0$, the term $i \cdot t$ vanishes from the definition of the martingale.

Now let us fix $k \in \mathbb{N}$ such that $\ell + \boldsymbol{v}_r < \boldsymbol{v}_{\max} + k$ and define a *stopping time* τ (see e.g. [21]) which returns the first point in time in which either $m^{(\tau)} \geq \boldsymbol{v}_{\max} + k$, or $m^{(\tau)} \leq \boldsymbol{v}_{\max}$. A routine application of optional stopping theorem gives us the following

$$\mathcal{P}(m^{(\tau)} \geq \boldsymbol{v}_{\max} + k) \quad \geq \quad \frac{\ell + \boldsymbol{v}_r - \boldsymbol{v}_{\max}}{k + M} \,. \tag{1}$$

Denote by T the number of steps to hit $Q(0)$. Note that $m^{(\tau)} \geq \boldsymbol{v}_{\max} + k$ implies $c^{(\tau)} = m^{(\tau)} - \boldsymbol{v}_{p^{(\tau)}} \geq \boldsymbol{v}_{\max} + k - \boldsymbol{v}_{p^{(\tau)}} \geq k$, and thus also $T \geq k$, as at least k steps are required to decrease the counter value from k to 0. It follows that $\mathcal{P}(m^{(\tau)} \geq \boldsymbol{v}_{\max} + k) \leq \mathcal{P}(T \geq k)$. By putting this inequality together with the inequality (1) we obtain

$$\mathbb{E}[T] \;=\; \sum_{k \in \mathbb{N}} \mathcal{P}(T \geq k) \;\geq\; \sum_{k=\ell+1}^{\infty} \mathcal{P}(T \geq k) \;\geq\; \sum_{k=\ell+1}^{\infty} \frac{\ell + \boldsymbol{v}_r - \boldsymbol{v}_{\max}}{k + M} \;=\; \infty \,. \qquad \square$$

Further, we need the following observation about the structure of $\mathcal{M}_{\mathscr{A}}$, which holds also for non-probabilistic one-counter automata:

Proposition 9. *There is $k_2 \in \mathbb{N}$ such that for every configuration $r(\ell) \in Pre^*(q(0))$, where $\ell \geq k_2$, we have that if $r(\ell) \to r'(\ell')$, then $r'(\ell') \in Pre^*(q(0))$.*

To show that $E(p{\downarrow}q) = \infty$, it suffices to identify a subset $W \subseteq R(p{\downarrow}q)$ such that $\mathcal{P}(W) > 0$ and $\mathbb{E}[R_{p{\downarrow}q} \mid W] = \infty$. Now observe that if $Pre^*(q(0)) \cap Post^*(p(1))$ is infinite, there is a configuration $r(\ell) \in Pre^*(q(0))$ reachable from $p(1)$ along a finite path u such that $\ell \geq k_1 + k_2$, where k_1 and k_2 are the constants of Propositions 8 and 9.

Due to Proposition 8, the expected length of an honest path from $r(\ell - k_2)$ to $Q(0)$ is infinite. However, then also the expected length of an honest path from $r(\ell)$ to $Q(k_2)$ is infinite. This means that there is a state $s \in Q$ such that the expected length of an honest path from $r(\ell)$ to $s(k_2)$ in infinite. Further, it follows directly from Proposition 9 that $s(k_2) \in Pre^*(q(0))$ because there is an honest path from $r(\ell)$ to $s(k_2)$.

Now consider the set W of all runs w initiated in $p(1)$ that start with the finite path u, then follow an honest path from $r(\ell)$ to $s(k_2)$, and then follow an honest path from $s(k_2)$ to $q(0)$. Obviously, $\mathcal{P}(W) > 0$, and $\mathbb{E}[R_{p\downarrow q} \mid W] = \infty$ because the expected length of the middle subpath is infinite. Hence, $E(p\downarrow q) = \infty$ as needed.

3.2 Efficient approximation of finite expected termination time

Let us denote by $T^{>0}_{<\infty}$ the set of all pairs $(p, q) \in T^{>0}$ satisfying $E(p\downarrow q) < \infty$. Our aim is to prove the following:

Theorem 10. *For all $(p, q) \in T^{>0}_{<\infty}$, the value of $E(p\downarrow q)$ can be approximated up to an arbitrarily small absolute error $\varepsilon > 0$ in time polynomial in $|\mathscr{A}|$ and* $\log(1/\varepsilon)$.

Note that if y approximates $E(p\downarrow q)$ up to an absolute error $1 > \varepsilon > 0$, then y approximates $E(p\downarrow q)$ also up to the relative error ε because $E(p\downarrow q) \geq 1$.

The proof of Theorem 10 is based on the fact that the vector of all $E(p\downarrow q)$, where $(p, q) \in T^{>0}_{<\infty}$, is the unique solution of a system of linear equations whose coefficients can be efficiently approximated (see below). Hence, it suffices to approximate the coefficients, solve the approximated equations, and then bound the error of the approximation using standard arguments from numerical analysis.

Let us start by setting up the system of linear equations for $E(p\downarrow q)$. For all $p, q \in T^{>0}$, we fix a fresh variable $V(p\downarrow q)$, and construct the following system of linear equations, $\mathcal{L}$, where the termination probabilities are treated as constants:

$$V(p\downarrow q) \;=\; \sum_{(p,-1,q)\in\delta^{>0}} \frac{P^{>0}(p,-1,q)}{[p\downarrow q]} \;+\; \sum_{(p,0,t)\in\delta^{>0}} \frac{P^{>0}(p,0,t)\cdot[t\downarrow q]}{[p\downarrow q]} \cdot \Big(1 + V(t\downarrow q)\Big)$$

$$+ \sum_{(p,1,t)\in\delta^{>0}} \sum_{r\in Q} \frac{P^{>0}(p,1,t)\cdot[t\downarrow r]\cdot[r\downarrow q]}{[p\downarrow q]} \cdot \Big(1 + V(t\downarrow r) + V(r\downarrow q)\Big)$$

It has been shown in [13] that the tuple of all $E(p\downarrow q)$, where $(p, q) \in T^{>0}$, is the least solution of $\mathcal{L}$ in $\mathbb{R}^+ \cup \{\infty\}$ with respect to component-wise ordering (where ∞ is treated according to the standard conventions). Due to Corollary 6, we can further simplify the system $\mathcal{L}$ by erasing the defining equations for all $V(p\downarrow q)$ such that $E(p\downarrow q) = \infty$ (note that if $E(p\downarrow q) < \infty$, then the defining equation for $V(p\downarrow q)$ in $\mathcal{L}$ cannot contain any variable $V(r\downarrow t)$ such that $E(r\downarrow t) = \infty$).

Thus, we obtain the system $\mathcal{L}'$. It is straightforward to show that the vector of all finite $E(p\downarrow q)$ is the *unique* solution of the system $\mathcal{L}'$ (see, e.g., Lemma 6.2.3 and Lemma 6.2.4 in [1]). If we rewrite $\mathcal{L}'$ into a standard matrix form, we obtain a system $\boldsymbol{V} = H \cdot \boldsymbol{V} + \boldsymbol{b}$, where H is a nonsingular nonnegative matrix, $\boldsymbol{V}$ is the vector of variables

in $\mathcal{L}'$, and $\boldsymbol{b}$ is a vector. Further, we have that $\boldsymbol{b} = \mathbf{1}$, i.e., the constant coefficients are all 1. This follows from the following equality (see [12, 17]):

$$[p{\downarrow}q] = \sum_{(p,-1,q)\in\delta^{>0}} P^{>0}(p,-1,q) \;+\; \sum_{(p,0,t)\in\delta^{>0}} P^{>0}(p,0,t)\cdot[t{\downarrow}q] \;+\; \sum_{(p,1,t)\in\delta^{>0}} \sum_{r\in Q} P^{>0}(p,1,t)\cdot[t{\downarrow}r]\cdot[r{\downarrow}q] \tag{2}$$

Hence, $\mathcal{L}'$ takes the form $\boldsymbol{V} = H \cdot \boldsymbol{V} + \mathbf{1}$. Unfortunately, the entries of H can take irrational values and cannot be computed precisely in general. However, they can be approximated up to an arbitrarily small relative error using Proposition 3. Denote by G an approximated version of H. We aim at bounding the error of the solution of the "perturbed" system $\boldsymbol{V} = G \cdot \boldsymbol{V} + \mathbf{1}$ in terms of the error of G. To measure these errors, we use the l_∞ norm of vectors and matrices, defined as follows: For a vector $\boldsymbol{V}$ we have that $\|\boldsymbol{V}\| = \max_i |\boldsymbol{V}_i|$, and for a matrix M we have $\|M\| = \max_i \sum_j |M_{ij}|$. Hence, $\|M\| = \|M \cdot \mathbf{1}\|$ if M is nonnegative. We show the following:

Proposition 11. *Let* $b \geq \max\left\{E(p{\downarrow}q) \mid (p,q) \in T^{>0}_{<\infty}\right\}$. *Then for each* ε, *where* $0 < \varepsilon < 1$, *let* $\delta = \varepsilon/(12 \cdot b^2)$. *If* $\|G - H\| \leq \delta$, *then the perturbed system* $\boldsymbol{V} = G \cdot \boldsymbol{V} + \mathbf{1}$ *has a unique solution* $\boldsymbol{F}$, *and in addition, we have that*

$$|E(p{\downarrow}q) - \boldsymbol{F}_{pq}| \quad \leq \quad \varepsilon \qquad \textit{for all } (p,q) \in T^{>0}_{<\infty}.$$

Here $\boldsymbol{F}_{pq}$ *is the component of* $\boldsymbol{F}$ *corresponding to the variable* $V(p{\downarrow}q)$.

The proof of Proposition 11 is based on estimating the size of the condition number $\kappa = \|1 - H\| \cdot \|(1 - H)^{-1}\|$ and applying standard results of numerical analysis.

The value of b in Proposition 11 can be estimated as follows: By Theorem 5, we have

$$E(p{\downarrow}q) \quad \leq \quad 85000 \cdot |Q|^6 / \left(x_{\min}^{6|Q|^3} \cdot t_{\min}^4\right) \qquad \text{for all } (p,q) \in T^{>0}_{<\infty},$$

where $t_{\min} = \min\{|t| \neq 0 \mid t \text{ is the trend in a BSCC of } \mathcal{X}\}$. Although b appears large, it is really the value of $\log(1/b)$ which matters, and it is still reasonable. Theorem 10 now follows by combining Propositions 11 and 3, because the approximated matrix G can be computed using a number of arithmetical operations which is polynomial in $|\mathcal{A}|$ and $\log(1/\varepsilon)$.

4 Quantitative Model-Checking of ω-regular Properties

In this section, we show that for every ω-regular property encoded by a deterministic Rabin automaton, the probability of all runs in a given pOC that satisfy the property can be approximated up to an arbitrarily small relative error $\varepsilon > 0$ in polynomial time. This is achieved by designing and analyzing a new quantitative model-checking algorithm for pOC and ω-regular properties, which is *not* based on techniques developed for pPDA and RMC in [12, 15, 16].

Recall that a deterministic Rabin automaton (DRA) over a finite alphabet Σ is a deterministic finite-state automaton $\mathcal{R}$ with total transition function and *Rabin acceptance*

condition $(E_1, F_1), \ldots, (E_k, F_k)$, where $k \in \mathbb{N}$, and all E_i, F_i are subsets of control states of $\mathcal{R}$. For a given infinite word w over Σ, let $\inf(w)$ be the set of all control states visited infinitely often along the unique run of $\mathcal{R}$ on w. The word w is accepted by $\mathcal{R}$ if there is $i \leq k$ such that $\inf(w) \cap E_i = \emptyset$ and $\inf(w) \cap F_i \neq \emptyset$.

Let Σ be a finite alphabet, $\mathcal{R}$ a DRA over Σ, and $\mathscr{A} = (Q, \delta^{=0}, \delta^{>0}, P^{=0}, P^{>0})$ a pOC. A *valuation* is a function ν which to every configuration $p(i)$ of $\mathscr{A}$ assigns a unique letter of Σ. For simplicity, we assume that $\nu(p(i))$ depends only on the control state p (note that a "bounded" information about the current counter value can be encoded and maintained in the finite control of $\mathscr{A}$). Intuitively, the letters of Σ correspond to collections of predicates that are valid in a given configuration of $\mathscr{A}$. Thus, every run $w \in \mathit{Run}_{\mathscr{A}}(p(i))$ determines a unique infinite word $\nu(w)$ over Σ which is either accepted by $\mathcal{R}$ or not. The main result of this section is the following theorem:

Theorem 12. *For every $p \in Q$, the probability of all $w \in \mathit{Run}_{\mathscr{A}}(p(0))$ such that $\nu(w)$ is accepted by $\mathcal{R}$ can be approximated up to an arbitrarily small relative error $\varepsilon > 0$ in time polynomial in $|\mathscr{A}|$, $|\mathcal{R}|$, and $\log(1/\varepsilon)$.*

Our proof of Theorem 12 consists of three steps:

1. We show that the problem of our interest is equivalent to the problem of computing the probability of all accepting runs in pOC with Rabin acceptance condition.
2. We introduce a finite-state Markov chain $\mathcal{G}$ (with possibly irrational transition probabilities) such that the probability of all accepting runs in $\mathcal{M}_{\mathscr{A}}$ is equal to the probability of reaching a "good" BSCC in $\mathcal{G}$.
3. We show how to compute the probability of reaching a "good" BSCC in $\mathcal{G}$ with relative error at most ε in time polynomial in $|\mathscr{A}|$ and $\log(1/\varepsilon)$.

Let us note that Steps 1 and 2 are relatively simple, but Step 3 requires several insights. In particular, we cannot solve Step 3 without bounding a positive non-termination probability in pOC (i.e., a positive probability of the form $[p\uparrow]$) away from zero. This is achieved in our "divergence gap theorem" (i.e., Theorem 18), which is based on applying Azuma's inequality to the martingale constructed in Section 3.

Step 1. Let $\mathscr{A} = (Q, \delta^{=0}, \delta^{>0}, P^{=0}, P^{>0})$ be a pOC. A *Rabin acceptance condition for* $\mathscr{A}$ is finite sequence $(\mathcal{E}_1, \mathcal{F}_1), \ldots, (\mathcal{E}_k, \mathcal{F}_k)$, where $\mathcal{E}_i, \mathcal{F}_i \subseteq Q$ for all $1 \leq i \leq k$. For every run $w \in \mathit{Run}_{\mathscr{A}}$, let $Q\text{-inf}(w)$ be the set of all $p \in Q$ visited infinitely often along w. We use $\mathit{Run}_{\mathscr{A}}(p(0), \mathit{acc})$ to denote the set of all *accepting runs* $w \in \mathit{Run}_{\mathscr{A}}(p(0))$ such that $Q\text{-inf}(w) \cap \mathcal{E}_i = \emptyset$ and $Q\text{-inf}(w) \cap \mathcal{F}_i \neq \emptyset$ for some $i \leq k$. Sometimes we also write $\mathit{Run}_{\mathscr{A}}(p(0), \mathit{rej})$ to denote the set $\mathit{Run}_{\mathscr{A}}(p(0)) \setminus \mathit{Run}_{\mathscr{A}}(p(0), \mathit{acc})$ of *rejecting* runs. Our next proposition says that the problem of computing/approximating the probability of all runs w in a given pOC that are accepted by a given DRA is efficiently reducible to the problem of computing/approximating the probability of all accepting runs in a given pOC with Rabin acceptance condition. The proof is simple (we just "synchronize" a given pOC with a given DRA).

Proposition 13. *Let Σ be a finite alphabet, $\mathscr{A}$ a pOC, ν a valuation, $\mathcal{R}$ a DRA over Σ, and $p(0)$ a configuration of $\mathscr{A}$. Then there is a pOC $\mathscr{A}'$ with Rabin acceptance condition and a configuration $p'(0)$ of $\mathscr{A}'$ constructible in polynomial time such that the probability of all $w \in \mathit{Run}_{\mathscr{A}}(p(0))$ where $\nu(w)$ is accepted by $\mathcal{R}$ is equal to the probability of all accepting $w \in \mathit{Run}_{\mathscr{A}'}(p'(0))$.*

For the rest of this section, we fix a pOC $\mathscr{A} = (Q, \delta^{=0}, \delta^{>0}, P^{=0}, P^{>0})$ and a Rabin acceptance condition $(\mathcal{E}_1, \mathcal{F}_1), \ldots, (\mathcal{E}_k, \mathcal{F}_k)$ for $\mathscr{A}$. We show how to approximate the probability of $Run_{\mathscr{A}}(p(0), acc)$.

Step 2. Let $\mathcal{G}$ be a finite-state Markov chain, where $Q \times \{0,1\} \cup \{acc, rej\}$ is the set of states (the elements of $Q \times \{0,1\}$ are written as $q(i)$, where $i \in \{0,1\}$), and the transitions of $\mathcal{G}$ are defined as follows:

- $r(0) \xrightarrow{x} q(j)$ is a transition of $\mathcal{G}$ iff $r(0) \xrightarrow{x} q(j)$ is a transition of $\mathcal{M}_{\mathscr{A}}$;
- $r(1) \xrightarrow{x} q(0)$ iff $x = [r{\downarrow}q] > 0$;
- $r(1) \xrightarrow{x} acc$ iff $x = \mathcal{P}(Run_{\mathscr{A}}(r(1), acc) \cap Run_{\mathscr{A}}(r{\uparrow})) > 0$;
- $r(1) \xrightarrow{x} rej$ iff $x = \mathcal{P}(Run_{\mathscr{A}}(r(1), rej) \cap Run_{\mathscr{A}}(r{\uparrow})) > 0$;
- $acc \xrightarrow{1} acc$, $rej \xrightarrow{1} rej$;
- there are no other transitions.

Note that almost every $w \in Run_{\mathscr{A}}(p(0))$ has its "twin" $w' \in Run_{\mathcal{G}}(p(0))$, which is obtained from w as follows: each honest subpath in w of the form $r(1), \ldots, q(0)$ is replaced with a single transition $r(1) \to q(0)$ in w'; and if the counter is decreased to zero only finitely many times along w, then the last transition of the form $r(0) \to q(1)$ in w is replaced either with $r(0) \to acc$ or $r(0) \to rej$ in w', depending on whether w is accepting or rejecting (the rest of w is then replaced with loops on acc or rej).

A BSCC B of $\mathcal{G}$ is *good* if either $B = \{acc\}$, or there is $i \leq k$ such that $\mathcal{E}_i \cap Q(B) = \emptyset$ and $\mathcal{F}_i \cap Q(B) \neq \emptyset$, where $Q(B)$ consists of all $r \in Q$ such that either $r(j) \in B$ for some $j \in \{0,1\}$, or there are $t(1), q(0) \in B$ such that $t(1) \to q(0)$ is a transition in $\mathcal{G}$ and $r(j) \in Pre^*(q(0)) \cap Post^*(t(1))$ for some $j \in \mathbb{N}_0$. For every $p \in Q$, let $Run_{\mathcal{G}}(p(0), good)$ be the set of all $w \in Run_{\mathcal{G}}(p(0))$ that visit a good BSCC of $\mathcal{G}$. The next proposition is obtained by a careful case analysis of accepting runs in $\mathcal{M}_{\mathscr{A}}$.

Proposition 14. *For every $p \in Q$ we have $\mathcal{P}(Run_{\mathscr{A}}(p(0), acc)) = \mathcal{P}(Run_{\mathcal{G}}(p(0), good))$.*

Step 3. Due to Proposition 14, the problem of our interest reduces to the problem of approximating the probability of visiting a good BSCC in the finite-state Markov chain $\mathcal{G}$. Since the termination probabilities in $\mathscr{A}$ can be approximated efficiently (see Proposition 3), the only problem with $\mathcal{G}$ is approximating the probabilities x and y in transitions of the form $p(1) \xrightarrow{x} acc$ and $p(1) \xrightarrow{y} rej$. Recall that x and y are the probabilities of all $w \in Run_{\mathscr{A}}(p{\uparrow})$ that are accepting and rejecting, respectively. A crucial observation is that almost all $w \in Run_{\mathscr{A}}(p{\uparrow})$ still behave accordingly with the underlying finite-state Markov chain $\mathcal{X}$ of $\mathscr{A}$ (see Section 3). More precisely, we have the following:

Proposition 15. *Let $p \in Q$. For almost all $w \in Run_{\mathscr{A}}(p{\uparrow})$ we have that w visits a BSCC B of $\mathcal{X}$ after finitely many transitions, and then it visits all states of B infinitely often.*

A BSCC B of $\mathcal{X}$ is *consistent* with the considered Rabin acceptance condition if there is $i \leq k$ such that $B \cap \mathcal{E}_i = \emptyset$ and $B \cap \mathcal{F}_i \neq \emptyset$. If B is not consistent, it is *inconsistent*. An immediate corollary to Proposition 15 is the following:

Corollary 16. *Let $Run_{\mathscr{A}}(p(1), cons)$ and $Run_{\mathscr{A}}(p(1), inco)$ be the sets of all $w \in Run_{\mathscr{A}}(p(1))$ such that w visit a control state of some consistent and inconsistent BSCC of $\mathcal{X}$, respectively. Then*

- $\mathcal{P}(Run_{\mathscr{A}}(p(1), acc) \cap Run_{\mathscr{A}}(p\uparrow)) = \mathcal{P}(Run_{\mathscr{A}}(p(1), cons) \cap Run_{\mathscr{A}}(p\uparrow))$
- $\mathcal{P}(Run_{\mathscr{A}}(p(1), rej) \cap Run(p\uparrow)) = \mathcal{P}(Run_{\mathscr{A}}(p(1), inco) \cap Run_{\mathscr{A}}(p\uparrow))$

Due to Corollary 16, we can reduce the problem of computing the probabilities of transitions of the form $p(1) \xrightarrow{x} acc$ and $p(1) \xrightarrow{y} rej$ to the problem of computing the divergence probability in pOC. More precisely, we construct pOC's $\mathscr{A}_{cons}$ and $\mathscr{A}_{inco}$ which are the same as $\mathscr{A}$, except that for each control state q of an inconsistent (or consistent, resp.) BSCC of $\mathcal{X}$, all positive outgoing rules of q are replaced with $q \xrightarrow{1,-1}_{>0} q$. Then $x = \mathcal{P}(Run_{\mathscr{A}_{cons}}(p\uparrow))$ and $y = \mathcal{P}(Run_{\mathscr{A}_{inco}}(p\uparrow))$.

Due to [4], the problem whether a given divergence probability is positive (in a given pOC) is decidable in polynomial time. This means that the underlying graph of $\mathcal{G}$ is computable in polynomial time, and hence the sets G_0 and G_1 consisting of all states s of $\mathcal{G}$ such that $\mathcal{P}(Run_{\mathcal{G}}(s, good))$ is equal to 0 and 1, respectively, are constructible in polynomial time. Let G be the set of all states of $\mathcal{G}$ that are not contained in $G_0 \cup G_1$, and let $X_{\mathcal{G}}$ be the stochastic matrix of $\mathcal{G}$. For every $s \in G$ we fix a fresh variable V_s and the equation

$$V_s \quad = \quad \sum_{s' \in G} X_{\mathcal{G}}(s, s') \cdot V_{s'} \quad + \quad \sum_{s' \in G_1} X_{\mathcal{G}}(s, s')$$

Thus, we obtain a system of linear equations $\boldsymbol{V} = A\boldsymbol{V} + \boldsymbol{b}$ whose unique solution $\boldsymbol{V}^*$ in $\mathbb{R}$ is the vector of probabilities of reaching a good BSCC from the states of G. This system can also be written as $(I - A)\boldsymbol{V} = \boldsymbol{b}$. Since the elements of A and $\boldsymbol{b}$ correspond to (sums of) transition probabilities in $\mathcal{G}$, it suffices to compute the transition probabilities of $\mathcal{G}$ with a sufficiently small relative error so that the approximate A and $\boldsymbol{b}$ produce an approximate solution where the relative error of each component is bounded by the ε. By combining standard results for finite-state Markov chains with techniques of numerical analysis, we show the following:

Proposition 17. *Let $c = 2|Q|$. For every $s \in G$, let R_s be the probability of visiting a BSCC of $\mathcal{G}$ from s in at most c transitions, and let $R = \min\{R_s \mid s \in G\}$. Then $R > 0$ and if all transition probabilities in $\mathcal{G}$ are computed with relative error at most $\varepsilon R^3/8(c+1)^2$, then the resulting system $(I - A')\boldsymbol{V} = \boldsymbol{b}'$ has a unique solution $\boldsymbol{U}^*$ such that $|\boldsymbol{V}^*_s - \boldsymbol{U}^*_s|/\boldsymbol{V}^*_s \leq \varepsilon$ for every $s \in G$.*

Note that the constant R of Proposition 17 can be bounded from below by $x_t^{|Q|-1} \cdot x_n$, where

- $x_t = \min\{X_{\mathcal{G}}(s, s') \mid s, s' \in G\}$, i.e., x_t is the minimal probability that is either explicitly used in $\mathscr{A}$, or equal to some positive termination probability in $\mathscr{A}$;
- $x_n = \min\{X_{\mathcal{G}}(s, s') \mid s \in G, s' \in G_1\}$, i.e., x_n is the minimal probability that is either a positive termination probability in $\mathscr{A}$, or a positive non-termination probability in the pOC's $\mathscr{A}_{cons}$ and $\mathscr{A}_{inco}$ constructed above.

Now we need to employ the promised divergence gap theorem, which bounds a positive non-termination probability in pOC away from zero (for all $p, q \in Q$, we use $[p, q]$ to denote the probability of all runs w initiated in $p(1)$ that visit a configuration $q(k)$, where $k \geq 1$ and the counter stays positive in all configurations preceding this visit).

Theorem 18. *Let $\mathscr{A} = (Q, \delta^{=0}, \delta^{>0}, P^{=0}, P^{>0})$ be a pOC and $\mathcal{X}$ the underlying finite-state Markov chain of $\mathscr{A}$. Let $p \in Q$ such that $[p\uparrow] > 0$. Then there are two possibilities:*

1. *There is* $q \in Q$ *such that* $[p,q] > 0$ *and* $[q\uparrow] = 1$. *Hence,* $[p\uparrow] \geq [p,q]$.
2. *There is a BSCC* $\mathscr{B}$ *of* $\mathcal{X}$ *and a state* q *of* $\mathscr{B}$ *such that* $[p,q] > 0$, $t > 0$, *and* $\boldsymbol{v}_q = \boldsymbol{v}_{\max}$ *(here* t *is the trend,* $\boldsymbol{v}$ *is the vector of Proposition 7, and* $\boldsymbol{v}_{\max}$ *is the maximal component of* $\boldsymbol{v}$*; all of these are considered in* $\mathscr{B}$*). Further,* $[p\uparrow] \geq [p,q]t^3/12(2(\boldsymbol{v}_{\max} - \boldsymbol{v}_{\min}) + 4)^3$.

Hence, denoting the relative precision $\varepsilon R^3/8(c+1)^2$ of Proposition 17 by δ, we obtain that $\log(1/\delta)$ is bounded by a polynomial in $|\mathscr{A}|$ and $\log(1/\varepsilon)$. Further, the transition probabilities of $\mathcal{G}$ can be approximated up to the relative error δ in time polynomial in $|\mathscr{A}|$ and $\log(1/\varepsilon)$ by approximating the termination probabilities of $\mathscr{A}$ (see Proposition 3). This proves Theorem 12.

5 Experimental results, future work

We have implemented a prototype tool in the form of a Maple worksheet[3], which allows to compute the termination probabilities of pOC and the conditional expected termination times. Our tool employs Newton's method to approximate the termination probabilities within a sufficient accuracy so that the expected termination time is computed with absolute error (at most) one by solving the linear equation system from Section 3.2.

We applied our tool to the pOC model of the program *TreeEval* (see Section 1) for various values of the parameters. The following table shows the results. We also show the associated termination probabilities, rounded to three digits. We write $[a\downarrow 0]$ etc. to abbreviate $[(and,init)\downarrow(or,return,0)]$ etc., and $[a\downarrow]$ for $[a\downarrow 0] + [a\downarrow 1]$.

	$[a\downarrow]$	$[a\downarrow 0]$	$[a\downarrow 1]$	$E[a\downarrow 0]$	$E[a\downarrow 1]$
$z = 0.5, y = 0.4, x_a = 0.2, x_o = 0.2$	0.800	0.500	0.300	11.000	7.667
$z = 0.5, y = 0.4, x_a = 0.2, x_o = 0.4$	0.967	0.667	0.300	104.750	38.917
$z = 0.5, y = 0.4, x_a = 0.2, x_o = 0.6$	1.000	0.720	0.280	20.368	5.489
$z = 0.5, y = 0.4, x_a = 0.2, x_o = 0.8$	1.000	0.732	0.268	10.778	2.758
$z = 0.5, y = 0.5, x_a = 0.1, x_o = 0.1$	0.861	0.556	0.306	11.400	5.509
$z = 0.5, y = 0.5, x_a = 0.2, x_o = 0.1$	0.931	0.556	0.375	23.133	20.644
$z = 0.5, y = 0.5, x_a = 0.3, x_o = 0.1$	1.000	0.546	0.454	83.199	111.801
$z = 0.5, y = 0.5, x_a = 0.4, x_o = 0.1$	1.000	0.507	0.493	12.959	21.555
$z = 0.2, y = 0.4, x_a = 0.2, x_o = 0.2$	0.810	0.696	0.115	7.827	6.266
$z = 0.3, y = 0.4, x_a = 0.2, x_o = 0.2$	0.811	0.636	0.175	8.928	6.783
$z = 0.4, y = 0.4, x_a = 0.2, x_o = 0.2$	0.808	0.571	0.236	10.005	7.258
$z = 0.5, y = 0.4, x_a = 0.2, x_o = 0.2$	0.800	0.500	0.300	11.000	7.667

We believe that other interesting quantities and numerical characteristics of pOC, related to both finite paths and infinite runs, can also be efficiently approximated using the methods developed in this paper. An efficient implementation of the associated algorithms would result in a verification tool capable of analyzing an interesting class of infinite-state stochastic programs, which is beyond the scope of currently available tools limited to finite-state systems only.

[3] Available at `http://www.comlab.ox.ac.uk/people/stefan.kiefer/pOC.mws`.

References

1. T. Brázdil. *Verification of Probabilistic Recursive Sequential Programs*. PhD thesis, Masaryk University, Faculty of Informatics, 2007.
2. T. Brázdil, V. Brožek, J. Holeček, and A. Kučera. Discounted properties of probabilistic pushdown automata. In *Proceedings of LPAR 2008*, volume 5330 of *LNCS*, pages 230–242. Springer, 2008.
3. T. Brázdil, V. Brožek, and K. Etessami. One-counter stochastic games. In *Proceedings of FST&TCS 2010*, volume 8 of *LIPIcs*, pages 108–119. Schloss Dagstuhl, 2010.
4. T. Brázdil, V. Brožek, K. Etessami, A. Kučera, and D. Wojtczak. One-counter Markov decision processes. In *Proceedings of SODA 2010*, pages 863–874. SIAM, 2010.
5. T. Brázdil, J. Esparza, and A. Kučera. Analysis and prediction of the long-run behavior of probabilistic sequential programs with recursion. In *Proceedings of FOCS 2005*, pages 521–530. IEEE, 2005.
6. T. Brázdil, S. Kiefer, and A. Kučera. Efficient analysis of probabilistic programs with an unbounded counter. *CoRR*, abs/1102.2529, 2011.
7. T. Brázdil, A. Kučera, and O. Stražovský. On the decidability of temporal properties of probabilistic pushdown automata. In *Proceedings of STACS 2005*, volume 3404 of *LNCS*, pages 145–157. Springer, 2005.
8. J. Canny. Some algebraic and geometric computations in PSPACE. In *Proceedings of STOC'88*, pages 460–467. ACM Press, 1988.
9. K. Chatterjee and L. Doyen. Energy parity games. In *Proceedings of ICALP 2010, Part II*, volume 6199 of *LNCS*, pages 599–610. Springer, 2010.
10. K. Chatterjee, L. Doyen, T. Henzinger, and J.-F. Raskin. Generalized mean-payoff and energy games. In *Proceedings of FST&TCS 2010*, volume 8 of *LIPIcs*, pages 505–516. Schloss Dagstuhl, 2010.
11. J. Esparza, D. Hansel, P. Rossmanith, and S. Schwoon. Efficient algorithms for model checking pushdown systems. In *Proceedings of CAV 2000*, volume 1855 of *LNCS*, pages 232–247. Springer, 2000.
12. J. Esparza, A. Kučera, and R. Mayr. Model-checking probabilistic pushdown automata. In *Proceedings of LICS 2004*, pages 12–21. IEEE, 2004.
13. J. Esparza, A. Kučera, and R. Mayr. Quantitative analysis of probabilistic pushdown automata: Expectations and variances. In *Proceedings of LICS 2005*, pages 117–126. IEEE, 2005.
14. K. Etessami, D. Wojtczak, and M. Yannakakis. Quasi-birth-death processes, tree-like QBDs, probabilistic 1-counter automata, and pushdown systems. In *Proceedings of 5th Int. Conf. on Quantitative Evaluation of Systems (QEST'08)*. IEEE, 2008.
15. K. Etessami and M. Yannakakis. Algorithmic verification of recursive probabilistic systems. In *Proceedings of TACAS 2005*, volume 3440 of *LNCS*, pages 253–270. Springer, 2005.
16. K. Etessami and M. Yannakakis. Checking LTL properties of recursive Markov chains. In *Proceedings of 2nd Int. Conf. on Quantitative Evaluation of Systems (QEST'05)*, pages 155–165. IEEE, 2005.
17. K. Etessami and M. Yannakakis. Recursive Markov chains, stochastic grammars, and monotone systems of non-linear equations. In *Proceedings of STACS 2005*, volume 3404 of *LNCS*, pages 340–352. Springer, 2005.
18. J.G. Kemeny and J.L. Snell. *Finite Markov Chains*. D. Van Nostrand Company, 1960.
19. S. Kiefer, M. Luttenberger, and J. Esparza. On the convergence of Newton's method for monotone systems of polynomial equations. In *Proceedings of STOC 2007*, pages 217–226. ACM Press, 2007.
20. J.S. Rosenthal. *A first look at rigorous probability theory*. World Scientific Publishing, 2006.
21. D. Williams. *Probability with Martingales*. Cambridge University Press, 1991.