Efficient Analysis of Probabilistic Programs with an Unbounded Counter

TOMÁŠ BRÁZDIL, Masaryk University
STEFAN KIEFER, University of Oxford
ANTONÍN KUČERA, Masaryk University

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. We start by establishing 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. Using these observations, we show that the expected termination time can be approximated up to an arbitrarily small relative error in polynomial time, and the same holds for the probability of all runs that satisfy a given ω-regular property encoded by a deterministic Rabin automaton.

Categories and Subject Descriptors: F.3.1 [**Logics and Meanings of Programs**]: Specifying and Verifying and Reasoning about Programs—*Mechanical verification*

General Terms: Algorithms, Performance, Verification

Additional Key Words and Phrases: Markov chains, model-checking, one-counter automata

ACM Reference Format:
Tomáš Brázdil, Stefan Kiefer, and Antonín Kučera. 2014. Efficient analysis of probabilistic programs with an unbounded counter. J. ACM 61, 6, Article 41 (November 2014), 35 pages.
DOI: http://dx.doi.org/10.1145/2629599

1. INTRODUCTION

In this article we aim at designing *efficient* algorithms for analyzing basic properties of probabilistic programs operating on unbounded data domains that can be abstracted into a nonnegative integer counter. Consider, for example, the recursive program *TreeEval* of Figure 1 which evaluates a given AND-OR tree, that is, a tree whose root is an AND node, all children of AND nodes are either leaves or OR nodes, and all children of OR nodes are either leaves or AND nodes. Note that the program *TreeEval* evaluates a subtree only when necessary. In general, we cannot say much about its expected termination time; if the input tree is infinite, the program may not even terminate, that is, 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 children on average;
—an OR node has about o children on average;
—the length of a branch is b on average;
—the probability that a leaf evaluates to 1 is z.

T. Brázdil and A. Kučera are supported by the research center Institute for Theoretical Computer Science (ITI), grant No. P202/12/G061. S. Kiefer is supported by a University Research Fellowship of the Royal Society.
Authors' addresses: T. Brázdil and A. Kučera, Faculty of Informatics, Masaryk University, Botanická 68a, 602 00 Brno, Czech Republic; email: {brazdil, kucera}@fi.muni.cz; S. Kiefer, Department of Computer Science, University of Oxford, United Kingdom; email: stefan.kiefer@cs.ox.ac.uk.

DOI: http://dx.doi.org/10.1145/2629599

```
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
```

Fig. 1. The program *TreeEval* for evaluating AND-OR trees.

Further, let us assume that the actual number of children 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 children is $(1 - x_a)^{n-1}x_a$ with $x_a = \frac{1}{a}$. Under these assumptions, the behaviour of *TreeEval* is well-defined in the probabilistic sense, and we may ask what its *expected complexity* is, that is, the expected termination time. Since *TreeEval* is recursive, this question is not trivial. By applying the generic results of Section 3, the answer can be produced automatically and efficiently.

Apart from the expected termination time, which is a fundamental characteristic of terminating runs, we also consider the properties of *nonterminating* runs in probabilistic programs, 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.

The abstract class of probabilistic programs considered in this article 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 nonnegative 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 whose states are the configurations and whose transitions are determined by the rules. For every pair of control states p, q, we use $Run(p{\downarrow}q)$ to denote the set of all runs initiated in $p(1)$ that reach $q(0)$ so that the counter value stays positive in all configurations preceding the visit to $q(0)$. The probability of $Run(p{\downarrow}q)$ is denoted by $[p{\downarrow}q]$, and the conditional expected number of transitions needed to reach $q(0)$ from $p(1)$, under the condition that a run of $Run(p{\downarrow}q)$ is performed, is denoted by $E(p{\downarrow}q)$. A probability of the form $[p{\downarrow}q]$ is called *termination probability*, and an expectation of the form $E(p{\downarrow}q)$ is called *expected termination time*. The runs initiated in $p(1)$ that do not visit a configuration with zero counter are called *diverging*, and the probability of all diverging runs initiated in $p(1)$ is denoted by $[p{\uparrow}]$ (clearly, $[p{\uparrow}] = 1 - \sum_q [p{\downarrow}q]$).

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 pOC model of Figure 2 with 6 control states and 12 positive rules (zero rules are irrelevant and hence not shown in Figure 2). 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. The initial configuration is $\langle and, init\rangle(1)$, and the

/* if we have a leaf, return 1 or 0 */

$\langle and, init\rangle \xrightarrow{y\cdot z,\,-1} \langle or, return, 1\rangle,$

$\langle and, init\rangle \xrightarrow{y\cdot(1-z),\,-1} \langle or, return, 0\rangle$

/* otherwise, call OR */

$\langle and, init\rangle \xrightarrow{(1-y),1} \langle or, init\rangle$

/* if OR returns 1, call another OR or return 1 */

$\langle and, return, 1\rangle \xrightarrow{(1-x_a),\,1} \langle or, init\rangle$

$\langle and, return, 1\rangle \xrightarrow{x_a,\,-1} \langle or, return, 1\rangle$

/* if OR returns 0, return 0 immediately */

$\langle and, return, 0\rangle \xrightarrow{1,\,-1} \langle or, return, 0\rangle$

/* if we have a leaf, return 1 or 0 */

$\langle or, init\rangle \xrightarrow{y\cdot z,\,-1} \langle and, return, 1\rangle,$

$\langle or, init\rangle \xrightarrow{y\cdot(1-z),\,-1} \langle and, return, 0\rangle$

/* otherwise, call AND */

$\langle or, init\rangle \xrightarrow{(1-y),1} \langle and, init\rangle$

/* if AND returns 0, call another AND or return 0 */

$\langle or, return, 0\rangle \xrightarrow{(1-x_o),\,1} \langle and, init\rangle$

$\langle or, return, 0\rangle \xrightarrow{x_o,\,-1} \langle and, return, 0\rangle$

/* if AND returns 1, return 1 immediately */

$\langle or, return, 1\rangle \xrightarrow{1,\,-1} \langle and, return, 1\rangle$

Fig. 2. A probabilistic one-counter automaton which models the program *TreeEval*.

pOC terminates either in $\langle or, return, 0\rangle(0)$ or $\langle or, return, 1\rangle(0)$, which corresponds to evaluating the input tree to 0 and 1, respectively. Hence, $E(\langle and, init\rangle{\downarrow}\langle or, return, 0\rangle)$ and $E(\langle and, init\rangle{\downarrow}\langle or, return, 1\rangle)$ are the conditional expected termination times under the condition that the input tree evaluates to 0 and 1, respectively.

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) [Esparza et al. 2004] or *recursive Markov chains* (RMC) [Etessami and Yannakakis 2005c]. However, as it is mentioned in this article, 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 represent a convenient tradeoff between expressiveness and tractability.

The relevance of pOC is not limited just to recursive programs. As observed in Etessami et al. [2008], 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 the late 60s (see, e.g., Neuts [1981]). QBDs are widely used in queuing theory, performance evaluation, etc., and the main algorithmic problems studied in this context concern the invariant probability distribution in ergodic QBDs. Very recently, games over (probabilistic) one-counter automata, also called "energy games", were considered in several independent works [Chatterjee and Doyen 2010; Chatterjee et al. 2010; Brázdil et al. 2010a, 2010b]. The study is motivated by optimizing the use of resources (such as energy) in modern computational devices.

Our contribution. We start by connecting the quantitative analysis of pOC to martingale theory (see, e.g., Billingsley [1995], Rosenthal [2006], and Williams [1991] for a general introduction to martingales). In Theorem 3.4, we show how to construct a suitable martingale for a given pOC. By analyzing this martingale, we obtain the following results.

(A) We characterize the case when $E(p{\downarrow}q) = \infty$, and we give an *upper bound* for a finite $E(p{\downarrow}q)$ in Theorem 3.2.
(B) We give a *lower bound* for a positive divergence probability $[p{\uparrow}]$ in Theorem 4.8.

These results have the following algorithmic consequences.

(1) The problem whether for given p, q the value $E(p{\downarrow}q)$ is finite is in **P** (see Corollary 3.3). Further, a finite $E(p{\downarrow}q)$ is computable up to an arbitrarily small relative error $\varepsilon > 0$ in time polynomial in the size of the underlying pOC and $\log(1/\varepsilon)$ (see Theorem 3.5). 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 initiated in a configuration $p(0)$ of 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 $\log(1/\varepsilon)$. Further, the problem whether this probability is equal to 1 is in **P** (see Theorem 4.1).

In our algorithms, we employ the techniques that have been invented for pPDA and RMC in Esparza et al. [2004, 2005] and Etessami and Yannakakis [2005c]. We also rely on the recent result of Stewart et al. [2013] where it is shown that the termination probabilities in pOC can be approximated up to a given relative error $\varepsilon > 0$ in time which is polynomial in the size of pOC and $\log(1/\varepsilon)$. More concretely, (1) is proven as follows.

—The problem whether $E(p{\downarrow}q)$ is infinite is shown to be in **P** by using the characterization in (A).

—For all finite $E(p{\downarrow}q)$, we construct a system of linear equations $\mathcal{L}$ based on the natural recursive dependency among all finite $E(p{\downarrow}q)$ such that the tuple of all finite $E(p{\downarrow}q)$ is the unique solution of $\mathcal{L}$. The coefficients in $\mathcal{L}$ are given only symbolically as fractions involving termination probabilities of the form $[p{\downarrow}q]$, and their values may be irrational. Using the upper bound for a finite $E(p{\downarrow}q)$ obtained in (A), we show that for a given $\varepsilon > 0$ there exists $\delta > 0$ computable in time polynomial in the size of the underlying pOC and $\log(1/\varepsilon)$ such that a system $\mathcal{L}'$ derived from $\mathcal{L}$ by approximating the coefficients up to the relative error δ still has a unique solution whose relative error (with respect to the solution of $\mathcal{L}$) is bounded by ε. Hence, it suffices to compute the system $\mathcal{L}'$, which can be done in polynomial time due to Stewart et al. [2013], and solve $\mathcal{L}'$ exactly.

The results of (2) are proven in several steps.

—For a pOC $\mathscr{A}$, a configuration $p(0)$ of $\mathscr{A}$, and a deterministic Rabin automaton $\mathcal{R}$, we construct a *finite* Markov chain $\mathcal{G}$ and a state p_0 of $\mathcal{G}$ such that the probability of all runs initiated in $p(0)$ and accepted by $\mathcal{R}$ is equal to the probability $\mathcal{P}(Run(p_0, good))$ of all runs initiated in p_0 and visiting a "good" bottom strongly connected component (BSCC) of $\mathcal{G}$. The set of states and the transition relation of $\mathcal{G}$ are computable in time polynomial in $|\mathscr{A}|$ and $|\mathcal{R}|$, and the defining condition of a "good" BSCC is also verifiable in polynomial time. However, the transition probabilities in $\mathcal{G}$ are specified only symbolically and may take irrational values.

—The problem whether $\mathcal{P}(Run(p_0, good)) = 1$ is shown to be in **P** by applying standard methods for finite-state Markov chains (here we do not need to compute/approximate the transition probabilities of $\mathcal{G}$; see, for example, Kemeny and Snell [1960]). Thus, we obtain the qualitative part of (2).

—We show that the transition probabilities of $\mathcal{G}$ can be approximated up to a given relative error $\delta > 0$ in time polynomial in $|\mathscr{A}|$ and $\log(1/\delta)$. This result crucially depends on the lower bound obtained in (B).

—We construct a system of linear equations $\mathcal{L}$ such that $\mathcal{P}(Run(p_0, good))$ is a component of the unique solution of $\mathcal{L}$. The variables of $\mathcal{L}$ correspond to the states of $\mathcal{G}$ that can reach a good BSCC of $\mathcal{G}$ with probability strictly between zero and one, and the coefficients of $\mathcal{L}$ correspond to the transition probabilities of $\mathcal{G}$. Using the lower bound of (B), we show that, for a given $\varepsilon > 0$, there exists $\delta > 0$ computable in time polynomial in $|\mathscr{A}|$ and $\log(1/\varepsilon)$ such that an approximated system $\mathcal{L}'$, where the relative error of all coefficients is bounded by δ, has a unique solution whose relative error is bounded by ε. Hence, it suffices to compute $\mathcal{L}'$ (here we again employ the procedure of Stewart et al. [2013] for approximating termination probabilities) and solve this system exactly. Thus, we obtain the quantitative part of (2).

Let us note that in the preliminary conference version of this article [Brázdil et al. 2011b], we used a procedure of Etessami et al. [2010] for approximating the termination probabilities $[p{\downarrow}q]$ up to a given relative error $\varepsilon > 0$. This procedure runs in polynomial time on the unit-cost rational arithmetic RAM, and the same computational model was adopted in Brázdil et al. [2011b] when formulating the results. In this article, we employ the improved procedure of Stewart et al. [2013], and hence we can use the standard Turing machine model.

Related work. In Esparza et al. [2004] and Etessami and Yannakakis [2005c], 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) [Canny 1988]. Due to the results of Etessami and Yannakakis [2005c], 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 [Kiefer et al. 2007]. Further, any nontrivial approximation of the nontermination probabilities is at least as hard as the SquareRootSum problem [Etessami and Yannakakis 2005c], whose exact complexity is a long-standing open question in exact numerical computations. The best-known upper bound for SquareRootSum is **CH** (counting hierarchy; see Corollary 1.4 in Allender et al. [2008]). 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 article, is provably infeasible because the termination probabilities can be doubly-exponentially small in the size of a given pPDA or RMC [Etessami and Yannakakis 2005c].

The expected termination time and the expected reward per transition in pPDA and RMC have been studied in Esparza et al. [2005]. 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 and fractions of) termination probabilities are used as coefficients. Hence, the equational system can be represented only symbolically, and the corresponding approximation algorithm employs the decision procedure for Tarski algebra (the system $\mathcal{L}$ used in the approximation algorithm of (1) can be seen as a special case of the system constructed in Esparza et al. [2005]). There are other results for pPDA and RMC, which concern model-checking problems for linear-time [Etessami and Yannakakis 2005a, 2005b] and branching-time [Brázdil et al. 2005b] logics, long-run average properties [Brázdil et al. 2005a], discounted properties of runs [Brázdil et al. 2008], etc. An overview of the existing results about pPDA and RMC can be found in Brázdil et al. [2013].

It has been shown in Etessami et al. [2010] that when the decomposed Newton's method is used to approximate the termination probabilities in pOC, it needs only *polynomially* many iterations before it starts to produce one bit of precision per iteration (cf., the corresponding result for pPDA mentioned previously). Consequently, termination probabilities in pOC can be approximated up to a given relative error $\varepsilon > 0$ using only a polynomial number of arithmetic operations. In other words, the approximation algorithm of Etessami et al. [2010] runs in polynomial time assuming the unit-cost rational arithmetic RAM model of computation. This algorithm has recently been modified in [Stewart et al. 2013] by rounding the intermediate results carefully so that it runs in polynomial time on the standard Turing machine model.

One-counter Markov decision processes and one-counter stochastic games, where the choice among the outgoing transitions of a given configuration can be either stochastic (as in pOC) or nondeterministic, have been studied in Brázdil et al. [2010b, 2011a, 2010a, 2012]. Let us note that the martingale construction for pOC introduced in this article turned out to be applicable also to these generalized models.

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, nonnegative integers, rational numbers, and real numbers, respectively. The symbol ∞ in treated in the standard way (in particular, $x < \infty$ and $x+\infty = \infty+x = \infty$ for all $x \in \mathbb{R}$), and we also adopt the standard notation for intervals (e.g., $[0, \infty]$ denotes $\{x \in \mathbb{R} \mid x \geq 0\} \cup \{\infty\}$).

Let $\delta > 0$, $x \in \mathbb{Q}$, and $y \in \mathbb{R}$. We say that x is a *relative δ-approximation* of y if either $y \neq 0$ and $|x - y|/|y| \leq \delta$, or $x = y = 0$. Further, we say that x is an *absolute δ-approximation* of y if $|x - y| \leq \delta$.

Given a finite set Q, we regard elements of $\mathbb{R}^Q$ as vectors over Q. We use boldface symbols like $\mathbf{u}, \mathbf{v}$ 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. All vectors are considered as column vectors in matrix multiplications, unless otherwise stated (an example of a frequently used *row* vector is the invariant distribution β introduced in Section 3.1).

Let $\mathcal{V} = (V, \rightarrow)$, where V is a nonempty set of vertices and $\rightarrow \; \subseteq V \times V$ a *total* relation (i.e., for every $v \in V$ there is some $u \in V$ such that $v \rightarrow u$). A *finite path* in $\mathcal{V}$ of *length* $k \geq 0$ is a finite sequence of vertices $v_0, \ldots, v_k$, where $v_i \rightarrow v_{i+1}$ for all $0 \leq i < k$. The length of a finite path w is denoted by *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 for all $u, v \in U$ there is a finite path from u to v. Further, we say that U is a *strongly connected component (SCC)* if U is a maximal strongly connected subset of V. A *bottom SCC (BSCC)* is a SCC U such that for every $u \in U$ and every $u \rightarrow v$ we have that $v \in U$.

The class of problems solvable by a deterministic Turing machine in polynomial time is denoted by $\mathbf{P}$. Whenever we say that *X is computable in polynomial time*, we mean that X is computable by a deterministic Turing machine in polynomial time.

2.1. Markov Chains

We assume familiarity with basic notions of probability theory, for example, *probability space*, *random variable*, or the *expected value*. As usual, a *probability distribution* over a finite or countably infinite set X is a function $f : X \rightarrow [0, 1]$ 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 2.1. A *Markov chain* is a triple $\mathcal{M} = (S, \rightarrow, \mathit{Prob})$ where S is a finite or countably infinite set of *states*, $\rightarrow \; \subseteq S \times S$ 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 $s \rightarrow t$ and x is the probability of $s \rightarrow t$.

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 $s \not\rightarrow t$, 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 $\mathcal{P} : \mathcal{F} \rightarrow [0, 1]$ is the unique probability measure such that $\mathcal{P}(Run_{\mathcal{M}}(w)) = \prod_{i=1}^{length(w)} x_i$ where $w(i-1) \xrightarrow{x_i} w(i)$ for every $1 \leq i \leq length(w)$. If $length(w) = 0$, we put $\mathcal{P}(Run_{\mathcal{M}}(w)) = 1$.

2.2. Probabilistic One-Counter Automata

One-counter automata are abstract computational devices equipped with a finite control unit and an unbounded counter which can store nonnegative integers. Each transition can either increment, decrement, or leave unchanged the current counter value. Further, the counter can be "tested for zero" in the sense that there can be special transitions enabled only in configurations with zero counter. The probabilistic variant of one-counter automata is obtained by assigning positive probabilities to transitions. A formal definition follows.

Definition 2.2. 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 $\delta^{=0}$, respectively, 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 as fractions of binary numbers.

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 $p \xrightarrow{x,c-i}_{>0} q$.

We say that a finite path $p_0(\ell_0), \ldots, p_m(\ell_m)$ in $\mathcal{M}_{\mathscr{A}}$ is *zero-safe* if $\ell_i > 0$ for all $0 \leq i < m$ (in particular, observe that $p(0)$ is a zero-safe finite path of length 0 from $p(0)$ to $p(0)$, and there is no other zero-safe finite path initiated in $p(0)$). Further, 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 start with a zero-safe finite path from $p(1)$ to $q(0)$. The runs of $\bigcup_{q \in Q} Run_{\mathscr{A}}(p{\downarrow}q)$ are called the *terminating runs of* $p(1)$;
—$Run_{\mathscr{A}}(p{\uparrow})$ be the set of all *diverging* 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 article, we rely on the following proposition.

PROPOSITION 2.3. *Let $\mathscr{A} = (Q, \delta^{=0}, \delta^{>0}, P^{=0}, P^{>0})$ be a pOC, and $p, q \in Q$.*

(*A*) *The problem whether $[p{\downarrow}q] > 0$ is in **P**.*
(*B*) *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 $\mathscr{A}$.*
(*C*) *The probability $[p{\downarrow}q]$ can be approximated up to an arbitrarily small relative error $\varepsilon > 0$ in time polynomial in $|\mathscr{A}|$ and $\log(1/\varepsilon)$.*

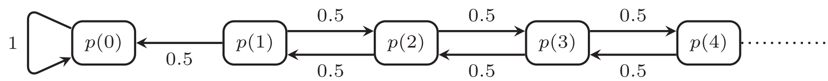

Fig. 3. The Markov chain $\mathcal{M}_{\mathscr{A}}$.

The problem considered in Part (A) of Proposition 2.3 is a special case of the standard reachability problem for pushdown automata [Hopcroft and Ullman 1979] which is known to be in **P** (see also the proof of Lemma 5.3 where a more general result of Esparza et al. [2000] related to the reachability problem for pushdown automata is recalled). Parts (B) and (C) of Proposition 2.3 are proven in Etessami et al. [2008, Corollary 6] and Stewart et al. [2013, Theorem 5], respectively. Let us note that a variant of Part (C) valid for the unit-cost rational arithmetic RAM model of computation was established already in Etessami et al. [2008, Theorem 14]. In our approximation algorithms (see Theorem 3.5 and Theorem 4.1), we use the procedure of Part (C) to compute the coefficients in certain systems of linear equations which are then solved exactly.

Let $T^{>0}$ be the set of all pairs $(p, q) \in Q \times Q$ satisfying $[p{\downarrow}q] > 0$. Note that $T^{>0}$ is computable in polynomial time due to Proposition 2.3(A). Further, for every $r(j) \in Q \times \mathbb{N}_0$, we define the sets $\mathit{Pre}^*(r(j))$ and $\mathit{Post}^*(r(j))$ where

—$\mathit{Pre}^*(r(j))$ consists of all configurations $t(\ell)$ such that there exists a zero-safe finite path from $t(\ell)$ to $r(j)$;
—$\mathit{Post}^*(r(j))$ consists of all configurations $t(\ell)$ such that there exists a zero-safe finite path from $r(j)$ to $t(\ell)$.

Note that $r(0) \in \mathit{Pre}^*(r(0))$ and $r(1) \in \mathit{Post}^*(r(1))$.

3. EXPECTED TERMINATION TIME

In this section, we give an efficient algorithm for approximating 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 $\mathscr{A} = (Q, \delta^{=0}, \delta^{>0}, P^{=0}, P^{>0})$, and we use $x_{\min}$ to denote the least (positive) probability used in the rules of $\mathscr{A}$. For all $p, q \in Q$, let $R_{p{\downarrow}q} : \mathit{Run}(p(1)) \to \mathbb{N}_0 \cup \{\infty\}$ be a random variable which to a given run w assigns either the k such that $w(0), \ldots, w(k)$ is a zero-safe finite path from $p(1)$ to $q(0)$, or ∞ if there is no such k. If $(p, q) \in T^{>0}$, we use $E(p{\downarrow}q)$ to denote the conditional expectation $\mathbb{E}[R_{p{\downarrow}q} \mid \mathit{Run}(p{\downarrow}q)]$.

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

Example 3.1. Consider a simple pOC with only one control state p, one zero rule $(p, 0, p)$, and two positive rules $(p, -1, p)$ and $(p, 1, p)$ that are both assigned the probability 1/2. The Markov chain $\mathcal{M}_{\mathscr{A}}$ is shown in Figure 3. Note that $[p{\downarrow}p]$ has to satisfy the equation $x = \frac{1}{2} + \frac{1}{2}x^2$, and hence $[p{\downarrow}p] = 1$. Further, $E(p{\downarrow}p)$ has to satisfy the equation $x = \frac{1}{2} + \frac{1}{2}(1 + 2x)$, which means $E(p{\downarrow}p) = \infty$ because the equation has no other nonnegative solution. See Esparza et al. [2005] and Section 3.2 for more details.

We proceed as follows. First, we show that if $E(p{\downarrow}q) < \infty$, then $E(p{\downarrow}q)$ is at most exponential in $|\mathscr{A}|$, and the problem whether $E(p{\downarrow}q) = \infty$ is in **P** (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. Size and Finiteness Of The Expected Termination Time

Let $\mathcal{X}$ be a finite-state Markov chain with Q as set of states and transition matrix $A \in [0, 1]^{Q\times Q}$ given by

$$A_{p,q} = \sum_{p \xrightarrow{x,c}_{>0} q} x.$$

Given a BSCC $\mathscr{B}$ of $\mathcal{X}$, let $B \in [0, 1]^{\mathscr{B}\times\mathscr{B}}$ be the restriction of A to the elements of $\mathscr{B}\times\mathscr{B}$, and let $\beta \in (0, 1]^{\mathscr{B}}$ be the *invariant distribution* of $\mathscr{B}$, that is, the unique row vector satisfying $\beta B = \beta$ and $\beta\mathbf{1} = 1$ (see, e.g., Kemeny and Snell [1960, Theorem 5.1.2]). Now we define

—the vector $\mathbf{s} \in \mathbb{R}^{\mathscr{B}}$ of *expected counter changes* by $\mathbf{s}_p = \sum_{p \xrightarrow{x,c}_{>0} q} x \cdot c$,
—the *trend* $t \in \mathbb{R}$ of $\mathscr{B}$ by $t = \beta\mathbf{s}$.

Intuitively, the trend is the average counter change per transition. Note that t is easily computable in time polynomial in $|\mathscr{A}|$ (hence, the binary length of t is also polynomial in $|\mathscr{A}|$). Our aim is to prove the following theorem.

THEOREM 3.2. *Let $(p, q) \in T^{>0}$. Then we have the following.*

(A) If q is not in a BSCC of $\mathcal{X}$, then $E(p{\downarrow}q) \leq 5|Q|/x_{\min}^{|Q|+|Q|^3}$.
(B) If q is in a BSCC $\mathscr{B}$ of $\mathcal{X}$, then
  (a) if $Pre^*(q(0)) \cap Post^*(p(1)) \cap \mathscr{B} \times \mathbb{N}_0$ is a finite set, then $E(p{\downarrow}q) \leq 20|Q|^3/x_{\min}^{4|Q|^3}$;
  (b) if $Pre^*(q(0)) \cap Post^*(p(1)) \cap \mathscr{B} \times \mathbb{N}_0$ is an infinite set, then
    (1) if $\mathscr{B}$ has trend $t \neq 0$, then $E(p{\downarrow}q) \leq 85000|Q|^6/(x_{\min}^{5|Q|+|Q|^3} \cdot t^4)$;
    (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 3.2 applies. In particular, due to Esparza et al. [2000], 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 $Pre^*(q(0)) \cap Post^*(p(1)) \cap \mathscr{B} \times \mathbb{N}_0$ and check whether the language accepted by $\mathcal{F}$ is finite (cf. Lemma 5.3). Thus, we have the following corollary.

COROLLARY 3.3. *Let $(p, q) \in T^{>0}$. The problem whether $E(p{\downarrow}q)$ is finite is in $\boldsymbol{P}$.*

The rest of this section is devoted to the proof of Theorem 3.2. In particular, we establish a powerful link between pOC and martingale theory which is also used in Section 4. For the sake of readability, we concentrate mainly on explaining the underlying ideas, and postpone the technical details to Section 5.1.

First assume case (A), that is, q is not in a BSCC of $\mathcal{X}$. Then, for all $s(\ell) \in Post^*(p(1))$, we have that $s(\ell)$ can reach either $q(0)$ or a configuration outside $Pre^*(q(0))$ in at most $|Q| - 1$ transitions. It follows that the probability of performing a zero-safe finite path from $p(1)$ to $q(0)$ of length i decays exponentially in i, and hence $E(p{\downarrow}q)$ is finite. The upper bound of case (A) is proven by standard methods (see Lemma 5.2).

Next assume case (B), that is, $q \in \mathscr{B}$ for some BSCC $\mathscr{B}$ of $\mathcal{X}$. It is easy to show that the expected time for a run in $Run(p{\downarrow}q)$ to reach $\mathscr{B}$ is finite. If we further assume that $\mathscr{C} := Pre^*(q(0)) \cap Post^*(p(1)) \cap \mathscr{B} \times \mathbb{N}$ is a finite set (case (B)(a)), then every run basically moves within a finite-state Markov chain on $\mathscr{C}$ after reaching $\mathscr{B}$. By assumption, $\mathscr{C}$ is finite which implies, by a pumping argument, that $|\mathscr{C}| \leq 3|Q|^3$ (see Lemma 5.3). Consequently, after a run of $Run(p{\downarrow}q)$ has reached $\mathscr{B}$, it reaches $q(0)$ in finite expected time which can be estimated due to the upper bound on the size of $\mathscr{C}$. Thus, we obtain the upper bound of case (B)(a) in Lemma 5.4.

Case (B)(b) requires new nontrivial techniques. We employ a generic observation which connects the study of pOC to martingale theory. Recall that a stochastic process $m^{(0)}, m^{(1)}, \dots$ is a martingale if, for all $i \in \mathbb{N}_0$,

—$\mathbb{E}[|m^{(i)}|] < \infty$,
—$\mathbb{E}[m^{(i+1)} \mid m^{(0)}, \dots, m^{(i)}] = m^{(i)}$ almost surely.

Two generic results about martingales that are used in this paper are *Azuma's inequality* and the *optional stopping theorem* (see, e.g., Billingsley [1995], Rosenthal [2006], and Williams [1991]). Let $m^{(0)}, m^{(1)}, \dots$ be a martingale such that $|m^{(k)} - m^{(k-1)}| \leq d$ almost surely for all $k \in \mathbb{N}$, and let $\tau : \Omega \to \mathbb{N}_0 \cup \{\infty\}$ be a random variable over the underlying probability space of $m^{(0)}, m^{(1)}, \dots$ such that $\mathbb{E}[\tau]$ is finite and τ is a stopping time, that is, for all $k \in \mathbb{N}_0$ the occurrence of the event $\tau = k$ depends only on the values $m^{(0)}, \dots, m^{(k)}$. Then Azuma's inequality states that for all $b > 0$ and $i \in \mathbb{N}$ we have that both $\mathcal{P}(m^{(i)} - m^{(0)} \geq b)$ and $\mathcal{P}(m^{(i)} - m^{(0)} \leq -b)$ are bounded by

$$\exp\left(\frac{-b^2}{2id^2}\right),$$

and the optional stopping theorem guarantees that $\mathbb{E}[m^{(\tau)}] = \mathbb{E}[m^{(0)}]$.

Let us fix an initial configuration $r(c) \in \mathscr{B} \times \mathbb{N}$. 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 $\mathbf{s}$ of expected counter changes is constant, that is, $\mathbf{s} = \mathbf{1} \cdot t$ where t is the trend of $\mathscr{B}$, then we can define a martingale $m^{(0)}, m^{(1)}, \dots$ 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 $\mathbf{s}$ is generally not constant, we might try to "compensate" the difference among the individual control states by a suitable vector $\mathbf{v} \in \mathbb{R}^{\mathscr{B}}$. The next proposition shows that this is indeed possible (a proof is postponed to Section 5.1).

THEOREM 3.4. *There is a vector $\mathbf{v} \in [0, \infty)^{\mathscr{B}}$ such that the stochastic process $m^{(0)}, m^{(1)}, \dots$ defined by*

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

is a martingale, where t is the trend of $\mathscr{B}$. Moreover, the vector $\mathbf{v}$ satisfies $0 \leq \mathbf{v}_p \leq 2|\mathscr{B}|/x_{\min}^{|\mathscr{B}|}$ for every $p \in \mathscr{B}$.

Due to Theorem 3.4, powerful results of martingale theory, such as the aforementioned optional stopping theorem and Azuma's inequality, become applicable to pOC. In this paper, we use the constructed martingale to establish case (B)(b) of Theorem 3.2, and to prove the crucial *divergence gap theorem* in Section 4. The range of possible applications of Theorem 3.4 is of course wider.

Assume now case (B)(b)(1), that is, $t \neq 0$. For simplicity, let us first assume that $p \in \mathscr{B}$. For every $i \in \mathbb{N}$, let $Run(p{\downarrow}q, i) = \{w \in Run(p{\downarrow}q) \mid R_{p{\downarrow}q}(w) = i\}$ be the set of all runs initiated in $p(1)$ that reach $q(0)$ in exactly 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 $[p{\downarrow}q, i] \leq a^i$. This immediately implies that $E(p{\downarrow}q)$ is finite, and the bound on $E(p{\downarrow}q)$ can be obtained by analyzing the size of a and h.

Consider the martingale $m^{(0)}, m^{(1)}, \ldots$ over $Run(p(1))$ as defined in Theorem 3.4, and let $\delta_{\mathbf{v}} := \mathbf{v}_{\max} - \mathbf{v}_{\min}$, where $\mathbf{v}_{\max}$ and $\mathbf{v}_{\min}$ are the maximal and the minimal components of $\mathbf{v}$, respectively. Realize that, for every $w \in Run(p{\downarrow}q, i)$, we have that

$$\big(m^{(i)} - m^{(0)}\big)(w) = \mathbf{v}_q - \mathbf{v}_p - i \cdot t.$$

Hence, $[p{\downarrow}q, i] \leq \mathcal{P}(m^{(i)} - m^{(0)} = \mathbf{v}_q - \mathbf{v}_p - i \cdot t)$. A simple computation reveals that, for a sufficiently large $h \in \mathbb{N}$ and all $i \geq h$ we have the following.

—If $t < 0$, then $[p{\downarrow}q, i] \leq \mathcal{P}(m^{(i)} - m^{(0)} \geq (i/2) \cdot (-t))$.
—If $t > 0$, then $[p{\downarrow}q, i] \leq \mathcal{P}(m^{(i)} - m^{(0)} \leq (i/2) \cdot (-t))$.

In each step, the martingale value changes by at most $\delta_{\mathbf{v}} + |t| + 1$, where $\delta_{\mathbf{v}}$ is defined above. Hence, by applying Azuma's inequality, we obtain the following for all $t \neq 0$ and $i \geq h$:

$$[p{\downarrow}q, i] \leq \exp\left(-\frac{(i/2)^2 t^2}{2i(\delta_{\mathbf{v}} + |t| + 1)^2}\right) = a^i.$$

Here $a := \exp(-t^2 \,/\, 8(\delta_{\mathbf{v}} + |t| + 1)^2) < 1$. In the general case, when p does not necessarily belong to $\mathscr{B}$, the analysis is slightly more complicated, and we also need to re-use the upper bound on the expected time to reach $\mathscr{B}$. The details are given in the proof of Lemma 5.6.

Finally, consider case (B)(b)(2), that is, $t = 0$. We need to show that $E(p{\downarrow}q) = \infty$. Since $Pre^*(q(0)) \cap Post^*(p(1)) \cap \mathscr{B} \times \mathbb{N}_0$ is infinite, for an arbitrarily large $k \in \mathbb{N}$ there is a configuration $r(k) \in Pre^*(q(0)) \cap Post^*(p(1))$. We show that if k is sufficiently large, then the expected number of transitions needed to decrease the counter by some fixed constant b is infinite. This is achieved by analyzing the martingale $m^{(0)}, m^{(1)}, \ldots$ for $r(k)$, but this time we use the optional stopping theorem to show that the probability of performing a finite path of length i which decreases the counter by b decays sufficiently slowly to make the expected length of this path infinite. It follows that $E(p{\downarrow}q)$ is also infinite. See Lemma 5.7 for details.

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 $E(p{\downarrow}q) < \infty$. Note that due to Corollary 3.3, the set $T^{>0}_{<\infty}$ is computable in polynomial time. Our aim is to prove the following theorem.

THEOREM 3.5. *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 an absolute ε-approximation of $E(p{\downarrow}q)$ (where $0 < \varepsilon < 1$) is also a relative ε-approximation of $E(p{\downarrow}q)$ because $E(p{\downarrow}q) \geq 1$.

Our proof of Theorem 3.5 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 are fractions of termination probabilities. Hence, the coefficients may take irrational values, but can be efficiently approximated up to an arbitrarily small relative error due to Proposition 2.3(C). The main problem is to determine a sufficient precision for the coefficients so that the solution of the perturbed system is sufficiently close to the vector of all $E(p{\downarrow}q)$. Here we use the bounds of Theorem 3.2.

Let us start by setting up the system of linear equations for $E(p{\downarrow}q)$. For all $(p, q) \in T^{>0}_{<\infty}$, 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, and all summands with zero coefficients are immediately deleted:

$$V(p{\downarrow}q) \;=\; \sum_{p \xrightarrow{x,-1}_{>0} q} \frac{x}{[p{\downarrow}q]} + \sum_{p \xrightarrow{x,0}_{>0} t} \frac{x \cdot [t{\downarrow}q]}{[p{\downarrow}q]} \cdot (1 + V(t{\downarrow}q))$$

$$+ \sum_{p \xrightarrow{x,1}_{>0} t} \sum_{r \in Q} \frac{x \cdot [t{\downarrow}r] \cdot [r{\downarrow}q]}{[p{\downarrow}q]} \cdot (1 + V(t{\downarrow}r) + V(r{\downarrow}q)).$$

Note that since $E(p{\downarrow}q) < \infty$, the equation for $V(p{\downarrow}q)$ in $\mathcal{L}$ cannot employ any variable $V(r{\downarrow}t)$ such that $E(r{\downarrow}t) = \infty$. Further, the vector $\vec{E}$ of all $E(p{\downarrow}q)$, where $(p,q) \in T^{>0}_{<\infty}$, is the least solution of $\mathcal{L}$ in $[0,\infty)^m$ with respect to component-wise ordering,[1] where $m := |T^{>0}_{<\infty}|$. Observe that $\vec{E} \geq \mathbf{1}$. We show that $\mathcal{L}$ has *no other solution in* $\mathbb{R}^m$. Assume the converse, that is, $\mathcal{L}$ has another solution $\vec{F} \in \mathbb{R}^m$. Then, $\vec{E}+c(\vec{E}-\vec{F})$ is also a solution of $\mathcal{L}$ for an arbitrarily small $c > 0$. Since $\vec{E} \neq \vec{F}$, there is a sufficiently small $\hat{c} > 0$ such that $\hat{c}(\vec{E}-\vec{F}) \neq \mathbf{0}$ and $|\hat{c}(\vec{E}-\vec{F})_i| \leq \frac{1}{2}$ for all $i \in \{1,\ldots,m\}$. Then, $\vec{E} + \hat{c}(\vec{E}-\vec{F}) \geq \mathbf{0}$ and since $\vec{E}$ is the least solution of $\mathcal{L}$ in $[0,\infty)^m$, we have that $\vec{E} \leq \vec{E} + \hat{c}(\vec{E}-\vec{F})$. From this we get $\hat{c}(\vec{E}-\vec{F}) \geq \mathbf{0}$. Since $\vec{E} - \hat{c}(\vec{E}-\vec{F}) \geq \mathbf{0}$ is also a solution of $\mathcal{L}$, we have a contradiction with the minimality of $\vec{E}$.

If we rewrite $\mathcal{L}$ into the standard matrix form, we obtain the system $\vec{V} = H \cdot \vec{V} + \mathbf{b}$, where H is a nonsingular nonnegative matrix, $\vec{V}$ is the vector of variables in $\mathcal{L}$, and $\mathbf{b}$ is a constant vector. Further, we have that $\mathbf{b} = \mathbf{1}$, which follows from the following equality (see Esparza et al. [2004] and Etessami and Yannakakis [2005c]):

$$[p{\downarrow}q] = \sum_{p \xrightarrow{x,-1}_{>0} q} x + \sum_{p \xrightarrow{x,0}_{>0} t} x \cdot [t{\downarrow}q] + \sum_{p \xrightarrow{x,1}_{>0} t} \sum_{r \in Q} x \cdot [t{\downarrow}r] \cdot [r{\downarrow}q]. \tag{1}$$

Hence, $\mathcal{L}$ takes the form $\vec{V} = H \cdot \vec{V} + \mathbf{1}$. As we already mentioned, the entries of H can take irrational values, but can be efficiently approximated up to an arbitrarily small relative error due to Proposition 2.3(C). Denote by G an approximated version of H. We aim at bounding the error of the solution of the "perturbed" system $\vec{V} = G \cdot \vec{V} + \mathbf{1}$ 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 $\mathbf{u}$, we have that $\|\mathbf{u}\| = \max_i |\mathbf{u}_i|$, and for a matrix M, we have $\|M\| = \max_i \sum_j |M_{ij}|$. Hence, $\|M\| = \|M \cdot \mathbf{1}\|$ if M is nonnegative. The next proposition is obtained by applying standard results of numerical analysis (see Section 5.2 for details).

PROPOSITION 3.6. *Let* $b \geq \max\{E(p{\downarrow}q) \mid (p,q) \in T^{>0}_{<\infty}\}$. *For every* ε *such that* $0 < \varepsilon < 1$, *let* $\delta = \varepsilon/(12 \cdot b^2)$. *If* $\|G - H\| \leq \delta$, *then the perturbed system* $\vec{V} = G \cdot \vec{V} + \mathbf{1}$ *has a unique solution* $\vec{F}$ *such that* $|E(p{\downarrow}q) - \vec{F}_{pq}| \leq \varepsilon$ *for all* $(p,q) \in T^{>0}_{<\infty}$. *Here,* $\vec{F}_{pq}$ *is the component of* $\vec{F}$ *corresponding to the variable* $V(p{\downarrow}q)$.

The value of b in Proposition 3.6 can be estimated as follows. By Theorem 3.2, for all $(p,q) \in T^{>0}_{<\infty}$, we have that

$$E(p{\downarrow}q) \leq 85000 \cdot |Q|^6 / \big(x_{\min}^{6|Q|^3} \cdot t_{\min}^4\big), \tag{2}$$

[1]This claim can be seen as a special case of a more general result achieved in Esparza et al. [2005] for probabilistic pushdown automata, and it can also be found in the standard QBD literature; see, for example, Neuts [1981].

where $t_{\min} = \min\{|t| \neq 0 \mid t \text{ is the trend of a BSCC of } \mathcal{X}\}$. Although b appears large, it is really the value of $\log(1/b)$ which matters, and it is still reasonable. Theorem 3.5 now follows by combining Propositions 3.6 and Inequality (2), because the approximated matrix G can be computed in time polynomial in $|\mathscr{A}|$ and $\log(1/\varepsilon)$.

4. QUALITATIVE AND QUANTITATIVE MODEL-CHECKING OF ω-REGULAR PROPERTIES

In this section, we show that for every ω-regular property encoded by a deterministic Rabin automaton,[2] 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 quantitative model-checking algorithm for pOC and ω-regular properties. The algorithm is based on similar ideas[3] as the algorithms designed for pPDA and RMC in Esparza et al. [2004] and Etessami and Yannakakis [2005a, 2005b]. The crucial new observation underpinning its functionality is the *divergence gap theorem* (i.e., Theorem 4.8), which bounds a positive probability of the form $[p\uparrow]$ away from zero. In the proof of Theorem 4.8, we use the martingale of Section 3 and apply the optional stopping theorem to derive certain lower bounds.

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 $w \in Run_{\mathscr{A}}(p(i))$ 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 4.1. *For every $p \in Q$, let $Run_{\mathscr{A}}(p(0), \mathcal{R})$ be the set of all $w \in Run_{\mathscr{A}}(p(0))$ that are accepted by $\mathcal{R}$. The problem whether $\mathcal{P}(Run_{\mathscr{A}}(p(0), \mathcal{R})) = 1$ is in **P**. Further, $\mathcal{P}(Run_{\mathscr{A}}(p(0), \mathcal{R}))$ can be approximated up to an arbitrarily small relative error $\varepsilon > 0$ in time polynomial in $|\mathscr{A}|$, $|\mathcal{R}|$, and* $\log(1/\varepsilon)$.

Since $\mathcal{R}$ is deterministic, it can be simulated on-the-fly in the finite control of $\mathscr{A}$. The resulting pOC has $|Q| \cdot |R|$ control states, where R is the set of control states of $\mathcal{R}$, and behaves in the same way as $\mathscr{A}$. Thus, we can translate the problem of Theorem 4.1 into an equivalent but technically simpler problem of *computing the probability of all accepting runs in pOC with Rabin acceptance condition*, which is formally defined in the following.

Let $\mathscr{A} = (Q, \delta^{=0}, \delta^{>0}, P^{=0}, P^{>0})$ be a pOC. A *Rabin acceptance condition for* $\mathscr{A}$ is a 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

[2]Recall that deterministic Rabin automata can encode an arbitrary ω-regular language [Thomas 1991] and there are quite efficient translations from various LTL fragments to deterministic Rabin automata [Křetínský and Ledesma-Garza 2013], although the complexity of this translation is exponential in general.

[3]In principle, we could use the algorithms developed for pPDA and RMC and apply them to pOC. The main ingredient of these algorithms is the construction of a finite-state Markov chain (called X_Δ in Esparza et al. [2004] or "summary chain" in Etessami and Yannakakis [2005a, 2005b]) which captures the behavior of infinite runs. To make this article self-contained, we design a simplified model-checking algorithm tailored specifically for pOC where the infinite runs are analyzed using a finite-state Markov chain $\mathcal{G}$. The chain $\mathcal{G}$ is simpler than X_Δ and the associated analysis leads to better estimates in Lemma 5.11.

run $w \in Run_{\mathscr{A}}$, let Q-inf(w) be the set of all $p \in Q$ visited infinitely often along w. We use $Run_{\mathscr{A}}(p(0), acc)$ to denote the set of all *accepting runs* $w \in Run_{\mathscr{A}}(p(0))$ such that Q-inf$(w) \cap \mathcal{E}_i = \emptyset$ and Q-inf$(w) \cap \mathcal{F}_i \neq \emptyset$ for some $i \leq k$. Sometimes we also write $Run_{\mathscr{A}}(p(0), rej)$ to denote the set $Run_{\mathscr{A}}(p(0)) \setminus Run_{\mathscr{A}}(p(0), acc)$ of *rejecting* runs. Theorem 4.1 is obtained as a direct corollary to the following proposition.

PROPOSITION 4.2. *Let $\mathscr{A} = (Q, \delta^{=0}, \delta^{>0}, P^{=0}, P^{>0})$ be a pOC and $(\mathcal{E}_1, \mathcal{F}_1), \ldots, (\mathcal{E}_k, \mathcal{F}_k)$ a Rabin acceptance condition for $\mathscr{A}$. For every $p \in Q$, the problem whether $\mathcal{P}(Run_{\mathscr{A}}(p(0), acc)) = 1$ is in **P**. Further, $\mathcal{P}(Run_{\mathscr{A}}(p(0), acc))$ can be approximated up to an arbitrarily small relative error $\varepsilon > 0$ in time polynomial in $|\mathscr{A}|$, k, and $\log(1/\varepsilon)$.*

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 $\mathscr{A}$. Our proof of Proposition 4.2 consists of two steps.

(1) 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 $\mathcal{G}$.
(2) We show how to approximate the probability of reaching a "good" BSCC in $\mathcal{G}$ up to a relative error $\varepsilon > 0$ in time polynomial in $|\mathscr{A}|$, k, and $\log(1/\varepsilon)$.

In Step (2), we re-use the martingale introduced in Section 3 to prove the aforementioned *divergence gap theorem* (Theorem 4.8).

Step (1) 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.

The correspondence between the runs of $Run_{\mathscr{A}}(p(0))$ and $Run_{\mathcal{G}}(p_0)$ is formally captured by a function $\Phi : Run_{\mathscr{A}}(p(0)) \to Run_{\mathcal{G}}(p_0) \cup \{\perp\}$, where $\Phi(w)$ is obtained from a given $w \in Run_{\mathscr{A}}(p(0))$ as follows.

—First, each *maximal* zero-safe subpath in w of the form $r(1), \ldots, q(0)$ is replaced with a single transition $r_1 \to q_0$.
—Then, all of the remaining configurations $s(0)$ with zero counter are replaced with s_0. Note that if w contained infinitely many configurations with zero counter, then the resulting sequence is a run of $Run_{\mathcal{G}}(p_0)$, and thus we obtain our $\Phi(w)$. Otherwise, the resulting sequence takes the form $v\,\hat{w}$, where $v \in FPath_{\mathcal{G}}(p_0)$ and $\hat{w}$ is a suffix of w initiated in a configuration $r(1)$. Then, we distinguish three possibilities.
 —If $\hat{w}$ is accepting and $\mathcal{P}(Run_{\mathscr{A}}(r(1), acc) \cap Run_{\mathscr{A}}(r{\uparrow})) > 0$, we put $\Phi(w) = v\, r_1\, acc^{\omega}$.
 —If $\hat{w}$ is rejecting and $\mathcal{P}(Run_{\mathscr{A}}(r(1), rej) \cap Run_{\mathscr{A}}(r{\uparrow})) > 0$, we put $\Phi(w) = v\, r_1\, rej^{\omega}$.
 —Otherwise, we put $\Phi(w) = \perp$.

LEMMA 4.3. *For every measurable subset $A \subseteq Run_{\mathcal{G}}(p_0)$, we have that $\Phi^{-1}(A)$ is measurable and $\mathcal{P}(A) = \mathcal{P}(\Phi^{-1}(A))$.*

A proof of Lemma 4.3 is straightforward (it suffices to check that the lemma holds for all basic cylinders $Run_{\mathcal{G}}(w)$ where $w \in FPath_{\mathcal{G}}(p_0)$). Note that Lemma 4.3 implies $\mathcal{P}(\Phi{=}\perp) = 0$.

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 $j \in \mathbb{N}$. A BSCC of $\mathcal{G}$ which is not good is *bad*. Note that every BSCC of $\mathcal{G}$ can be effectively classified as good or bad in polynomial time (see the remarks after Theorem 3.2). Now observe the following.

LEMMA 4.4. *Let B be a BSCC of $\mathcal{G}$, and let $Run_{\mathcal{G}}(p_0, B)$ be the set of all $w \in Run_{\mathcal{G}}(p_0)$ such that w hits B. If B is good/bad, then almost all runs of $\Phi^{-1}(Run_{\mathcal{G}}(p_0, B))$ are accepting/rejecting, respectively.*

A proof of Lemma 4.4 is straightforward—if $B = \{acc\}$ or $B = \{rej\}$, then all runs of $\Phi^{-1}(Run_{\mathcal{G}}(p_0, B))$ are accepting or rejecting, because all of them have an accepting or rejecting suffix, respectively. Otherwise, it suffices to realize that for almost all $w \in \Phi^{-1}(Run_{\mathcal{G}}(p_0, B))$ we have that $Q\text{-inf}(w) = Q(B)$.

Since almost all runs of $Run_{\mathcal{G}}(p_0)$ hit a BSCC of $\mathcal{G}$, our next proposition is a direct consequence of Lemma 4.3 and Lemma 4.4.

PROPOSITION 4.5. *Let $p \in Q$, and let $Run_{\mathcal{G}}(p_0, good)$ be the set of all $w \in Run_{\mathcal{G}}(p_0)$ that hit a good BSCC of $\mathcal{G}$. Then $\mathcal{P}(Run_{\mathscr{A}}(p(0), acc)) = \mathcal{P}(Run_{\mathcal{G}}(p_0, good))$.*

Step 2. Due to Proposition 4.5, the problem of approximating $\mathcal{P}(Run_{\mathscr{A}}(p(0), acc))$ reduces to the problem of approximating the probability of hitting a good BSCC in the finite-state Markov chain $\mathcal{G}$. Note that the probabilities associated to transitions of the form $r_1 \xrightarrow{x} q_0$, $r_1 \xrightarrow{x} acc$, and $r_1 \xrightarrow{x} rej$ in $\mathcal{G}$ may take irrational values. In the last two cases, it is even not clear how to decide whether such a transition exists in $\mathcal{G}$, that is, whether the associated probability x is positive (see the definition of $\mathcal{G}$). We show the following.

(a) The transition relation of $\mathcal{G}$ is computable in polynomial time.

(b) The probability of a transition $r_1 \xrightarrow{x} q_0$ in $\mathcal{G}$ satisfies $x \geq x_{\min}^{|Q|^3}$, and the probability of a transition $r_1 \xrightarrow{x} acc$ or $r_1 \xrightarrow{x} rej$ in $\mathcal{G}$ satisfies

$$x \geq \frac{x_{\min}^{4|Q|^2} \cdot t_{\min^+}^3}{7000 \cdot |Q|},$$

where $t_{\min^+} = \min\{t > 0 \mid t \text{ is the trend of a BSCC of } \mathcal{X}\}$. If there is no BSCC of $\mathcal{X}$ with positive trend,[4] we put $t_{\min^+} = 1$. Moreover, all transition probabilities of $\mathcal{G}$ can be approximated up to an arbitrarily small *relative* error $\varepsilon > 0$ in time polynomial in $|\mathscr{A}|$ and $\log(1/\varepsilon)$.

Note that if Claim (a) holds, we can efficiently compute the sets S_0 and S_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. This proves the "qualitative part" of Proposition 4.2. The "quantitative part" of Proposition 4.2 is obtained from Claim (b) as follows. Let $S_?$ be the set of all states of $\mathcal{G}$ that are not contained in $S_0 \cup S_1$, and let G be the stochastic matrix of $\mathcal{G}$. For every $s \in S_?$ we fix a fresh variable V_s and the equation

$$V_s = \sum_{t \in S_?} G(s, t) \cdot V_t + \sum_{t \in S_1} G(s, t).$$

[4]As we shall see in Section 5.3, transitions of the form $r_1 \xrightarrow{x} acc$ or $r_1 \xrightarrow{x} rej$ do not exist in $\mathcal{G}$ if the trends of all BSCCs of $\mathcal{X}$ are negative; however, they may exist if there is a BSCC with zero trend and no BSCC with positive trend.

Thus, we obtain a system of linear equations $\vec{V} = A\vec{V} + \mathbf{b}$ whose unique solution $\vec{V}^*$ in $\mathbb{R}^{|S_?|}$ is the vector of probabilities of reaching a good BSCC from the states of $S_?$. Since the elements of A and $\mathbf{b}$ correspond to (sums of) transition probabilities in $\mathcal{G}$, for every $\varepsilon > 0$ it suffices to compute the transition probabilities of $\mathcal{G}$ with a sufficiently small relative error $\delta > 0$ so that the approximate A and $\mathbf{b}$ produce an approximate solution where the relative error of each component is bounded by ε. Using standard results of numerical analysis and the lower bound on transition probabilities given in Claim (b), we show that δ can be chosen so that $\log(1/\delta)$ is bounded by a polynomial in $|\mathscr{A}|$ and $\log(1/\varepsilon)$. Now is suffices to apply the second part of Claim (b). The details are postponed to Section 5.3 (see Lemma 5.11).

In the rest of this section, we indicate how to prove Claims (a) and (b). Due to Proposition 2.3, we only need to consider transitions of the form $r_1 \xrightarrow{x} acc$ and $r_1 \xrightarrow{y} rej$, and the respective probabilities x and y. Recall that x and y are the probabilities of all $w \in Run_{\mathscr{A}}(r{\uparrow})$ that are accepting and rejecting, respectively. A simple but important observation is that almost all $w \in Run_{\mathscr{A}}(r{\uparrow})$ still behave accordingly with the underlying finite-state Markov chain $\mathcal{X}$ of $\mathscr{A}$ (see Section 3.1 for the definition of $\mathcal{X}$). More precisely, we have the following lemma.

LEMMA 4.6. *Let $r \in Q$. For almost all $w \in Run_{\mathscr{A}}(r{\uparrow})$, we have that w visits a BSCC $\mathscr{B}$ of $\mathcal{X}$ after finitely many transitions, and then it visits all states of $\mathscr{B}$ infinitely often.*

In fact, Lemma 4.6 is a variant of the standard result saying that almost all runs in a finite-state Markov chain $\mathcal{M}$ hit a BSCC B of $\mathcal{M}$ and then visit all states of B infinitely often (see, e.g., Kemeny and Snell [1960]). A proof of Lemma 4.6 does not require any new insights.

A BSCC $\mathscr{B}$ of $\mathcal{X}$ is *consistent* with the considered Rabin acceptance condition if there is $i \leq k$ such that $\mathscr{B} \cap \mathcal{E}_i = \emptyset$ and $\mathscr{B} \cap \mathcal{F}_i \neq \emptyset$. If $\mathscr{B}$ is not consistent, it is *inconsistent*. An immediate corollary to Lemma 4.6 is the following.

COROLLARY 4.7. *Let $Run_{\mathscr{A}}(r(1), cons)$ and $Run_{\mathscr{A}}(r(1), inco)$ be the sets of all $w \in Run_{\mathscr{A}}(r(1))$ such that w visits a control state of some consistent and inconsistent BSCC of $\mathcal{X}$, respectively. Then*

—$\mathcal{P}(Run_{\mathscr{A}}(r(1), acc) \cap Run_{\mathscr{A}}(r{\uparrow})) = \mathcal{P}(Run_{\mathscr{A}}(r(1), cons) \cap Run_{\mathscr{A}}(r{\uparrow}))$,
—$\mathcal{P}(Run_{\mathscr{A}}(r(1), rej) \cap Run_{\mathscr{A}}(r{\uparrow})) = \mathcal{P}(Run_{\mathscr{A}}(r(1), inco) \cap Run_{\mathscr{A}}(r{\uparrow}))$.

Let $\mathscr{A}_{cons}$ be a pOC which is the same as $\mathscr{A}$ except that for each control state q of an *inconsistent* BSCC of $\mathcal{X}$, all *positive* outgoing rules of q are replaced with $q \xrightarrow{1,-1}_{>0} q$ (the outgoing zero rules of q are irrelevant and may stay unchanged). Thus, almost all runs of $Run_{\mathscr{A}}(r{\uparrow})$ which were *rejecting* become *terminating* (i.e., visit a configuration with zero counter) in $\mathscr{A}_{cons}$. Hence,

$$\mathcal{P}(Run_{\mathscr{A}}(r(1), acc) \cap Run_{\mathscr{A}}(r{\uparrow})) = \mathcal{P}(Run_{\mathscr{A}}(r(1), cons) \cap Run_{\mathscr{A}}(r{\uparrow})) = \mathcal{P}(Run_{\mathscr{A}_{cons}}(r{\uparrow})).$$

Similarly, we construct a pOC $\mathscr{A}_{inco}$ which are the same as $\mathscr{A}$ except that for each control state q of a *consistent* BSCC of $\mathcal{X}$, all positive outgoing rules of q are replaced with $q \xrightarrow{1,-1}_{>0} q$. Then, $\mathcal{P}(Run_{\mathscr{A}}(r(1), rej) \cap Run_{\mathscr{A}}(r{\uparrow})) = \mathcal{P}(Run_{\mathscr{A}_{inco}}(r{\uparrow}))$.

Due to these observations, the problem of computing the probability of a transition $r_1 \xrightarrow{x} acc$ (or $r_1 \xrightarrow{y} rej$) in $\mathcal{G}$ reduces to the problem of computing the probability $[r{\uparrow}]$ in an efficiently constructible pOC $\mathscr{A}_{cons}$ (or $\mathscr{A}_{inco}$, respectively). Since the problem whether $[r{\uparrow}] > 0$ for a given control state r of a given pOC is solvable in polynomial time [Brázdil et al. 2010b], we obtain Claim (a).

To prove Claim (b), we need to establish a sufficiently large lower bound on $[r\uparrow]$ in $\mathscr{A}_{cons}$ and $\mathscr{A}_{inco}$. This bound is given in the following "divergence gap theorem".

THEOREM 4.8. *Let $\hat{\mathscr{A}}$ be a pOC with Q as the set of control states where the least positive probability used in the rules of $\hat{\mathscr{A}}$ is at least $x_{\min}$. For all $p, q \in Q$, let $[p, q]$ be the probability of all $w \in Run_{\hat{\mathscr{A}}}(p(1))$ such that w starts with a zero-safe finite path from $p(1)$ to $q(k)$, where $k \geq 1$. Let $p \in Q$ such that the probability $[p\uparrow]$ (considered in $\hat{\mathscr{A}}$) is positive. Then there are two possibilities.*

(1) *There is $q \in Q$ such that $[p, q] > 0$ and $[q\uparrow] = 1$. Hence, $[p\uparrow] \geq [p, q] \geq x_{\min}^{|Q|^2}$.*
(2) *There exists a BSCC $\mathscr{B}$ of the underlying finite-state Markov chain $\hat{\mathcal{X}}$ of $\hat{\mathscr{A}}$ and a state q of $\mathscr{B}$ such that the trend t of $\mathscr{B}$ is positive, $[p, q] > 0$, and $\mathbf{v}_q = \mathbf{v}_{\max}$. Here $\mathbf{v}$ is the vector of Theorem 3.4, and $\mathbf{v}_{\max}$ is the maximal component of $\mathbf{v}$; all of these are considered in $\mathscr{B}$. Further,*

$$[p\uparrow] \;\geq\; \frac{x_{\min}^{4|Q|^2} \cdot t^3}{7000 \cdot |Q|^3}\,.$$

A proof of Theorem 4.8 is obtained by analyzing the martingale of Section 3; see Section 5.3 for details.

Note that $\mathscr{A}_{cons}$ and $\mathscr{A}_{inco}$ have the same set of control states as $\mathscr{A}$, the least positive rule probability in $\mathscr{A}_{cons}$ and $\mathscr{A}_{inco}$ is at least $x_{\min}$, and if $\mathscr{B}$ is a BSCC of $\mathcal{X}_{cons}$ (or $\mathcal{X}_{inco}$) with a positive trend t, then $\mathscr{B}$ is also a BSCC of $\mathcal{X}$ with the same trend. Hence, Theorem 4.8 gives a lower bound on $[r\uparrow]$ in $\mathscr{A}_{cons}$ and $\mathscr{A}_{inco}$ and thus we obtain the first part of Claim (b).

The second part of Claim (b) is a trivial consequence of Theorem 4.8 and Proposition 2.3. Recall that $[r\uparrow] = 1 - \sum_{q\in Q}[r\downarrow q]$, and hence we can approximate $[r\uparrow]$ up to an arbitrarily small *absolute* error $\delta > 0$ efficiently by applying Proposition 2.3(C). Using the bound of Theorem 4.8, we can efficiently compute $\delta > 0$ such that $\log(1/\delta)$ is polynomial in $\log(1/\varepsilon)$ and the size of $\mathscr{A}_{cons}$ (or $\mathscr{A}_{inco}$), and every *absolute* δ-approximation of $[r\uparrow]$ is also a *relative* ε-approximation of $[r\uparrow]$.

5. PROOFS

In this section, we give proofs that were only sketched or completely omitted in the previous sections.

5.1. Proofs of Section 3.1

Recall that we assume a fixed pOC $\mathscr{A} = (Q, \delta^{=0}, \delta^{>0}, P^{=0}, P^{>0})$, where $x_{\min}$ denotes the least positive probability used in the rules of $\mathscr{A}$. Also recall the definition of the finite-state Markov chain $\mathcal{X}$.

For a given initial configuration $r(j) \in Q \times \mathbb{N}$ and a set of target configurations $F \subseteq Q \times \mathbb{N}_0$, we define a random variable T_F over the runs of $\mathcal{M}_{\mathscr{A}}$ initiated in $r(j)$ where $T_F(w)$ returns either the least $k \in \mathbb{N}_0$ such that $w(0), \ldots, w(k)$ is a zero-safe finite path from $r(j)$ to a configuration of F, or ∞ if there is no such k. Further, we use $Pre^*(F)$ to denote the set $\bigcup_{t(\ell)\in F} Pre^*(t(\ell))$. We start by establishing a simple tail bound for T_F.

LEMMA 5.1. *Let $r(j) \in Q \times \mathbb{N}$ be an initial configuration and $F \subseteq Q \times \mathbb{N}_0$ a set of target configurations. Further, let $n \in \mathbb{N}$ be a constant such that for every configuration $t(\ell) \in Post^*(r(j))$ there is a zero-safe finite path of length strictly less than n from $t(\ell)$ to a configuration which either belongs to F or is not contained in $Pre^*(F)$. Then, for all $k \geq n$, we have that $\mathcal{P}(k \leq T_F < \infty) \leq 2c^k$ where $c := \exp(-x_{\min}^n/n)$.*

Proof. Let us first consider the case when $x_{\min} = 1$. Then, there is a *unique* run w in $\mathcal{M}_{\mathscr{A}}$ initiated in $r(j)$. If $T_F(w) = \infty$, then $\mathcal{P}(T_F < \infty) = 0$. If $T_F(w) = m$ for some $m \in \mathbb{N}_0$, then the least constant n satisfying the assumptions of our lemma is equal to $m + 1$. Clearly, for every $k \geq m+1$ we have that $\mathcal{P}(k \leq T_F < \infty) = 0$.

Now assume $x_{\min} < 1$, and let $n \in \mathbb{N}$ be a constant satisfying the assumptions of our lemma. Since for each control state the sum of the probabilities of the outgoing (zero or positive) rules is 1, we must have $x \leq 1/2$. For every initial configuration $t(\ell) \in Post^*(r(j))$, call *crash* the event of performing a zero-safe finite path of length at most $n - 1$ ending in a configuration which either belongs to F or does not belong to $Pre^*(F)$. The probability of crash is at least $x_{\min}^{n-1} \geq x_{\min}^{n}$, regardless of the initial configuration $t(\ell)$. Let $k \geq n$. For the event $k \leq T_F < \infty$, a crash has to be avoided at least $\lfloor \frac{k-1}{n-1} \rfloor$ times, that is,

$$\mathcal{P}(k \leq T_F < \infty) \leq \left(1 - x_{\min}^{n}\right)^{\lfloor \frac{k-1}{n-1} \rfloor} .$$

As $\lfloor \frac{k-1}{n-1} \rfloor \geq \frac{k-1}{n-1} - 1 \geq \frac{k}{n} - 1$ and $\left(1 - x_{\min}^{n}\right)^{-1} \leq 2$ (recall $x_{\min} \leq 1/2$), we obtain

$$\begin{aligned}\mathcal{P}(k \leq T_F < \infty) &\leq \frac{1}{1 - x_{\min}^{n}} \cdot \left(\left(1 - x_{\min}^{n}\right)^{1/n}\right)^k \leq 2 \cdot \left(\left(1 - x_{\min}^{n}\right)^{1/n}\right)^k \\ &= 2 \cdot \left(\exp\left(\frac{1}{n} \log\left(1 - x_{\min}^{n}\right)\right)\right)^k \leq 2 \cdot \left(\exp\left(\frac{1}{n} \cdot \left(-x_{\min}^{n}\right)\right)\right)^k = 2 \cdot c^k .\end{aligned}$$

For the last inequality, recall that $\log(1 - y) \leq -y$ for all $y \in (0, 1)$. □

Now we can easily prove the following lemma.

Lemma 5.2 [*CASE (A) OF THEOREM* 3.2]. *Let $p, q \in Q$ such that $[p{\downarrow}q] > 0$ and q is not in a BSCC of $\mathcal{X}$. Then $E(p{\downarrow}q) \leq 5|Q| \,/\, x_{\min}^{|Q|+|Q|^3}$.*

Proof. Observe that if q is not in a BSCC of $\mathcal{X}$, then, for every configuration $t(\ell) \in Post^*(p(1))$, there is a zero-safe finite path of length at most $|Q| - 1$ initiated in $t(\ell)$ which ends either in $q(0)$ or in a configuration not contained in $Pre^*(q(0))$. Hence, we can apply Lemma 5.1 for the initial configuration $p(1)$, $F = \{q(0)\}$, and $n = |Q|$. Thus, we obtain

$$\begin{aligned}E(p{\downarrow}q) \cdot [p{\downarrow}q] &= \sum_{k \in \mathbb{N}} \mathcal{P}(k \leq R_{p{\downarrow}q} < \infty) = \sum_{k \in \mathbb{N}} \mathcal{P}(k \leq T_F < \infty) \\ &\leq \sum_{k=1}^{|Q|} 1 + \sum_{k=0}^{\infty} 2c^k = |Q| + \frac{2}{1-c}, \qquad \text{(Lemma 5.1)}\end{aligned}$$

where $c := \exp(-x_{\min}^{|Q|}/|Q|)$. We have $1 - c = 1 - \exp(-x_{\min}^{|Q|}/|Q|) \geq x_{\min}^{|Q|}/(2|Q|)$, hence

$$E(p{\downarrow}q) \cdot [p{\downarrow}q] \leq |Q| + \frac{4|Q|}{x_{\min}^{|Q|}} \leq \frac{5|Q|}{x_{\min}^{|Q|}} .$$

As $[p{\downarrow}q] \geq x_{\min}^{|Q|^3}$ by Proposition 2.3 (B), it follows $E(p{\downarrow}q) \leq 5|Q| \,/\, x_{\min}^{|Q|+|Q|^3}$. □

Now consider case (B) of Theorem 3.2. From now on, we assume that q belongs to some (fixed) BSCC $\mathscr{B}$ of the finite-state Markov chain $\mathcal{X}$. Every run $w \in Run(p{\downarrow}q)$ starts with a zero-safe finite path from $p(1)$ to some $r(k)$ where $r \in \mathscr{B}$, followed by a zero-safe finite path from $r(k)$ to $q(0)$ which visits only the configurations of $\mathscr{B} \times \mathbb{N}_0$. Our upper bounds for $E(p{\downarrow}q)$ are obtained as sums of upper bounds for the expected number of transitions needed to perform the two finite subpaths.

More precisely, we define the following random variables over $Run_{\mathcal{M}_{\mathscr{A}}}(p(1))$.

—$R^{(1)}(w)$ is the least $k \in \mathbb{N}_0$ such that $w(0), \ldots, w(k)$ is a zero-safe finite path from $p(1)$ to a configuration of $\mathscr{B} \times \mathbb{N}_0$; if there is no such k, we put $R^{(1)}(w) = \infty$.
—$R^{(2)}(w)$ is the ℓ such that $w(R^{(1)}(w)), \ldots, w(R^{(1)}(w){+}\ell)$ is a zero-safe finite path from $w(R^{(1)}(w))$ to $q(0)$; if there is no such ℓ (which includes the case when $R^{(1)}(w) = \infty$), we put $R^{(2)}(w) = \infty$. Intuitively, ℓ is the number of transitions needed to reach $q(0)$ after hitting $\mathscr{B}$.
—$Con(w)$ is the configuration $w(k)$ where $k = R^{(1)}(w)$. If $R^{(1)}(w) = \infty$, then $Con(w) = \bot$. That is, $Con(w)$ is the first configuration of w which hits $\mathscr{B}$.

Note that $R_{p\downarrow q}(w) = R^{(1)}(w) + R^{(2)}(w)$ for all runs w initiated in $p(1)$. Further, we have the following[5]:

$$
\begin{aligned}
E(p{\downarrow}q) \cdot [p{\downarrow}q] &= \sum_{k \in \mathbb{N}_0} \mathcal{P}\big(R_{p\downarrow q} = k\big) \cdot k = \sum_{k \in \mathbb{N}_0} \mathcal{P}\big(R^{(1)} + R^{(2)} = k\big) \cdot k \\
&= \sum_{k_1, k_2 \in \mathbb{N}_0} \mathcal{P}\big(R^{(1)} = k_1 \ \wedge\ R^{(2)} = k_2\big) \ \cdot (k_1 + k_2) \\
&= \sum_{k_1, k_2 \in \mathbb{N}_0} \mathcal{P}\big(R^{(1)} = k_1\big) \cdot \mathcal{P}\big(R^{(2)} = k_2 \mid R^{(1)} = k_1\big) \cdot (k_1 + k_2) \\
&= E_1 + E_2 \,,
\end{aligned}
$$

where

$$
\begin{aligned}
E_1 &:= \sum_{k_1, k_2 \in \mathbb{N}_0} \mathcal{P}\big(R^{(1)} = k_1\big) \cdot \mathcal{P}\big(R^{(2)} = k_2 \mid R^{(1)} = k_1\big) \cdot k_1 \qquad \text{and} \\
E_2 &:= \sum_{k_1, k_2 \in \mathbb{N}_0} \mathcal{P}\big(R^{(1)} = k_1\big) \cdot \mathcal{P}\big(R^{(2)} = k_2 \mid R^{(1)} = k_1\big) \cdot k_2 \,.
\end{aligned}
$$

Now observe that for every configuration $t(\ell) \in Post^*(p(1))$, there is a zero-safe finite path of length at most $|Q| - 1$ from $t(\ell)$ to a configuration which either belongs to $\mathscr{B} \times \mathbb{N}_0$ or is not contained in $Pre^*(\mathscr{B} \times \mathbb{N}_0)$. Hence, we can apply Lemma 5.1 for the initial configuration $p(1)$, $F = \mathscr{B} \times \mathbb{N}_0$, and $n = |Q|$. Thus, we obtain the following upper bound on E_1:

$$
\begin{aligned}
E_1 &= \sum_{k_1 \in \mathbb{N}_0} \mathcal{P}\big(R^{(1)} = k_1\big) \cdot k_1 \cdot \sum_{k_2 \in \mathbb{N}_0} \mathcal{P}\big(R^{(2)} = k_2 \mid R^{(1)} = k_1\big) \\
&\le \sum_{k_1 \in \mathbb{N}_0} \mathcal{P}\big(R^{(1)} = k_1\big) \cdot k_1 = \sum_{k_1 \in \mathbb{N}} \mathcal{P}(T_F = k_1) \cdot k_1 = \sum_{k_1 \in \mathbb{N}} \mathcal{P}(k_1 \le T_F < \infty) \\
&\le \sum_{k_1 = 1}^{|Q|} 1 + \sum_{k=0}^{\infty} 2c^k = |Q| + \frac{2}{1 - c} \qquad \text{(Lemma 5.1)}\,.
\end{aligned}
$$

We have $1 - c = 1 - \exp(-x_{\min}^{|Q|}/|Q|) \ge x_{\min}^{|Q|}/(2|Q|)$, hence

$$
E_1 \le \frac{5|Q|}{x_{\min}^{|Q|}}. \tag{3}
$$

[5]To simplify our notation, we adopt the convention that $\mathcal{P}(A \mid B)$ denotes 0 whenever $\mathcal{P}(B) = 0$.

Establishing an upper bound for E_2 is more difficult and it is done separately for case (B)(a) and case (B)(b) of Theorem 3.2.

Now we aim at proving the upper bound of case (B)(a). First, we show that if the set $Pre^*(q(0)) \cap Post^*(p(1)) \cap \mathscr{B} \times \mathbb{N}_0$ is finite, then its size is bounded by $3|Q|^3$. Intuitively, we just observe that the sets $Pre^*(q(0))$ and $Post^*(p(1))$ are recognizable by finite-state automata of "small" size, and hence the same holds for the product automaton recognizing the intersection. Then, we apply the standard pumping argument and conclude that if the product automaton accepted a "long" word, it would necessarily accept an infinite language.

LEMMA 5.3. *Let $p, q \in Q$. If the set $Pre^*(q(0)) \cap Post^*(p(1)) \cap \mathscr{B} \times \mathbb{N}_0$ is finite, then it has at most $3|Q|^3$ elements.*

PROOF. We use the notions and results of Esparza et al. [2000] which show how to compute the sets of all predecessors/successors of a regular set of configurations of a pushdown automaton. Note that $\mathscr{A}$ naturally determines a pushdown automaton Δ where Q is the set of control states, $\{I\}$ is the stack alphabet, and $tI \hookrightarrow rI^{c+1}$ is a transition rule of Δ iff $(t, c, r) \in \delta^{>0}$ (here I^k denotes the word consisting of k copies of the symbol I; in particular, I^0 is the empty string ε). Every configuration $p(k)$ of $\mathscr{A}$ then corresponds to the configuration pI^k of Δ, and there is a natural one-to-one correspondence between zero-safe finite paths initiated in $p(k)$ and finite paths initiated in pI^k (by definition, pushdown automata get stuck when the stack is emptied).

A $\mathcal{P}$*-automaton* is a nondeterministic finite-state automaton $\mathcal{F}$ over the alphabet $\{I\}$ such that the set of control states of $\mathcal{F}$ subsumes Q. A configuration rI^k of $\mathscr{A}$ is *recognized* by $\mathcal{F}$ if $\mathcal{F}$ accepts the word I^k from the initial state r. In Esparza et al. [2000], it has been shown that for every $\mathcal{F}$, one can compute another $\mathcal{P}$-automaton $\mathcal{F}_{pre^*}$ recognizing all predecessors of all configurations recognized by $\mathcal{F}$. Further, the automaton $\mathcal{F}_{pre^*}$ has the same set of control states as $\mathcal{F}$. In our case, $\mathcal{F}$ recognizes just $q\varepsilon$ and hence it has only $|Q|$ control states. So, $\mathcal{F}_{pre^*}$ has also $|Q|$ control states.

Similarly, one can construct a $\mathcal{P}$-automaton $\mathcal{F}_{post^*}$ recognizing all successors of all configurations recognized by $\mathcal{F}$. In our case, $\mathcal{F}$ recognizes just pI, and hence the resulting $\mathcal{F}_{post^*}$ has at most $|Q| + 2$ states (see Esparza et al. [2000]).

Using the standard product construction, we obtain a $\mathcal{P}$-automaton $\mathcal{F}$ with at most $|Q| \cdot (|Q| + 2)$ states recognizing the intersection of the sets recognized by $\mathcal{F}_{pre^*}$ and $\mathcal{F}_{post^*}$. That is, $\mathcal{F}$ recognizes all configurations rI^k such that $r(k) \in Pre^*(q(0)) \cap Post^*(p(1))$. Now note that if $Pre^*(q(0)) \cap Post^*(p(1)) \cap \mathscr{B} \times \mathbb{N}_0$ is finite, then a standard pumping argument for finite-state automata implies that the length of every word accepted by $\mathcal{F}$ from an initial state $r \in \mathscr{B}$ is bounded by $|Q| \cdot (|Q| + 2)$. It follows that there are at most $|Q|^2 \cdot (|Q| + 2)$ configurations in $Pre^*(q(0)) \cap Post^*(p(1)) \cap \mathscr{B} \times \mathbb{N}_0$ where the counter is positive, and at most one configuration with zero counter. Hence, the size of $Pre^*(q(0)) \cap Post^*(p(1)) \cap \mathscr{B} \times \mathbb{N}_0$ is bounded by $|Q|^3 + 2|Q|^2 + 1$. Note that $3|Q|^3 > |Q|^3 + 2|Q|^2 + 2$ for $|Q| \geq 2$; and in the special case when $|Q| = 1$ we have that $Pre^*(q(0)) \cap Post^*(p(1)) \cap \mathscr{B} \times \mathbb{N}_0$ is either empty (which contradicts $[p\downarrow q] > 0$) or infinite. □

LEMMA 5.4 (CASE (B)(A) OF THEOREM 3.2). *Let $p, q \in Q$ such that $[p\downarrow q] > 0$, $q \in \mathscr{B}$, and $Pre^*(q(0)) \cap Post^*(p(1)) \cap \mathscr{B} \times \mathbb{N}_0$ is finite. Then, $E(p\downarrow q) \leq 20|Q|^3/x_{\min}^{4|Q|^3}$.*

PROOF. Let $\mathscr{C} := Pre^*(q(0)) \cap Post^*(p(1)) \cap \mathscr{B} \times \mathbb{N}_0$. Note that if $Con(w) \neq \bot$ for a given run $w \in Run_{\mathcal{M}_{\mathscr{A}}}(p(1))$, then $Con(w) \in \mathscr{C}$.

Since $E(p{\downarrow}q) \cdot [p{\downarrow}q] = E_1 + E_2$ and we already have an upper bound on E_1 due to Inequality (3), it suffices to establish an upper bound on E_2. We have that

$$\begin{aligned}
E_2 &= \sum_{k_1,k_2\in\mathbb{N}_0} \mathcal{P}\big(R^{(1)} = k_1\big) \cdot \mathcal{P}\big(R^{(2)} = k_2 \mid R^{(1)} = k_1\big) \cdot k_2 \\
&= \sum_{k_2\in\mathbb{N}_0} k_2 \cdot \sum_{k_1\in\mathbb{N}_0} \mathcal{P}\big(R^{(2)} = k_2 \wedge R^{(1)} = k_1\big) \\
&= \sum_{k_2\in\mathbb{N}_0} k_2 \cdot \sum_{k_1\in\mathbb{N}_0} \sum_{r(j)\in\mathscr{C}} \mathcal{P}\big(R^{(2)} = k_2 \wedge R^{(1)} = k_1 \wedge Con = r(j)\big) \\
&= \sum_{k_2\in\mathbb{N}_0} k_2 \cdot \sum_{r(j)\in\mathscr{C}} \mathcal{P}\big(R^{(2)} = k_2 \wedge Con = r(j)\big) \\
&= \sum_{r(j)\in\mathscr{C}} \mathcal{P}\big(Con = r(j)\big) \cdot \sum_{k_2\in\mathbb{N}_0} k_2 \cdot \mathcal{P}\big(R^{(2)} = k_2 \mid Con = r(j)\big) \\
&\leq \max_{r(j)\in\mathscr{C}} \sum_{k_2\in\mathbb{N}_0} k_2 \cdot \mathcal{P}\big(R^{(2)} = k_2 \mid Con = r(j)\big).
\end{aligned}$$

Hence, all we need is an upper bound on $\sum_{k_2\in\mathbb{N}_0} k_2 \cdot \mathcal{P}(R^{(2)} = k_2 \mid Con = r(j))$ in the case when $\mathcal{P}(Con = r(j)) > 0$. Observe that for every configuration $t(\ell) \in Post^*(r(j))$ there is a zero-safe finite path of length at most $|Q|^3 + 2|Q|^2$ initiated in $t(\ell)$ which ends either in $q(0)$ or in a configuration not contained in $Pre^*(q(0))$ (here we use Lemma 5.3). By applying Lemma 5.1 for the initial configuration $r(j)$, $F = \{q(0)\}$, and $n = 3|Q|^3$, we get the following:

$$\begin{aligned}
\sum_{k_2\in\mathbb{N}_0} k_2 \cdot \mathcal{P}\big(R^{(2)} = k_2 \mid Con = r(j)\big) &= \sum_{k_2\in\mathbb{N}_0} k_2 \cdot \mathcal{P}(T_F = k_2) = \sum_{k_2\in\mathbb{N}_0} \mathcal{P}(k_2 \leq T_F < \infty) \\
&\leq 3|Q|^3 + \frac{12|Q|^3}{x_{\min}^{3|Q|^3}} \leq \frac{15|Q|^3}{x_{\min}^{3|Q|^3}}.
\end{aligned}$$

Hence, $E_2 \leq 15|Q|^3/x_{\min}^{3|Q|^3}$, and thus we have

$$[p{\downarrow}q] \cdot E(p{\downarrow}q) = E_1 + E_2 \leq \frac{5|Q|}{x_{\min}^{|Q|}} + \frac{15|Q|^3}{x_{\min}^{3|Q|^3}} \leq \frac{20|Q|^3}{x_{\min}^{3|Q|^3}}.$$

Since $[p{\downarrow}q] \geq x_{\min}^{|Q|^3}$ by Proposition 2.3(B), we finally obtain $E(p{\downarrow}q) \leq 20|Q|^3/x_{\min}^{4|Q|^3}$. □

It remains to prove case (B)(b) of Theorem 3.2. Recall the following notions: B denotes the transition matrix of $\mathscr{B}$, $\mathbf{s} \in \mathbb{R}^{\mathscr{B}}$ is the vector of expected counter changes defined by $\mathbf{s}_p = \sum_{p \xrightarrow{x,c}_{>0} q} x \cdot c$, and $t = \beta\mathbf{s}$ is the trend of $\mathscr{B}$, where β is the invariant distribution of $\mathscr{B}$. First, we restate and prove Theorem 3.4.

THEOREM 3.4. *There is a vector $\mathbf{v} \in [0,\infty)^{\mathscr{B}}$ such that the stochastic process $m^{(0)}$, $m^{(1)}, \ldots$ defined by*

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

is a martingale, where t is the trend of $\mathscr{B}$. Moreover, the vector $\mathbf{v}$ satisfies $0 \leq \mathbf{v}_p \leq 2|\mathscr{B}|/x_{\min}^{|\mathscr{B}|}$ for every $p \in \mathscr{B}$.

PROOF. A *potential* is a vector $\mathbf{v} \in \mathbb{R}^{\mathscr{B}}$ that satisfies $\mathbf{s} + B\mathbf{v} = \mathbf{v} + \mathbf{1}t$. The intuitive meaning of a potential $\mathbf{v}$ is that, starting in any state $r \in \mathscr{B}$, the expected counter increase after i steps for large i is $i \cdot t + \mathbf{v}_r$. Given a potential $\mathbf{v}$, we use $\mathbf{v}_{\max}$ and $\mathbf{v}_{\min}$ to denote the maximal and the minimal component of $\mathbf{v}$, respectively. First, we prove the following.

(a) Let $W := \mathbf{1}\beta$, that is, W is a square matrix where each row equals β. Let $Z := (I - B + W)^{-1}$. The matrix Z exists and the vector $Z\mathbf{s}$ is a potential.
(b) Denote by $x_{\min}$ the least positive coefficient of B. There exists a potential $\mathbf{v}$ with $0 \leq \mathbf{v}_p \leq 2|\mathscr{B}|/x_{\min}^{|\mathscr{B}|}$ for every $p \in \mathscr{B}$.

A proof of Claim (a). The matrix $Z := (I - B + W)^{-1}$ exists[6] by Kemeny and Snell [1960, Theorem 5.1.3]. Furthermore, by Kemeny and Snell [1960, Theorem 5.1.3(d)] the matrix Z satisfies $I + BZ = Z + W$. Multiplying with $\mathbf{s}$ and setting $\mathbf{u} := Z\mathbf{s}$, we obtain $\mathbf{s} + B\mathbf{u} = \mathbf{u} + \mathbf{1}\beta\mathbf{s}$; that is, $\mathbf{u}$ is a potential.

A Proof of Claim (b). Let $\mathbf{u}$ be the potential from Claim (a); that is, we have

$$(I - B)\mathbf{u} = \mathbf{s} - \mathbf{1}t\,. \tag{4}$$

By the Perron-Frobenius theorem for strongly connected matrices, there exists a positive vector $\mathbf{d} \in (0,1]^{\mathscr{B}}$ with $B\mathbf{d} = \mathbf{d}$; that is, $(I - B)\mathbf{d} = \mathbf{0}$. Observe that $\mathbf{u} + \kappa\mathbf{d}$ is a potential for all $\kappa \in \mathbb{R}$. Choose κ such that $\mathbf{v} := \mathbf{u} + \kappa\mathbf{d}$ satisfies $\mathbf{v}_{\max} = 2|\mathscr{B}|/x_{\min}^{|\mathscr{B}|}$. It suffices to prove $\mathbf{v}_{\min} \geq 0$. Let $q \in \mathscr{B}$ such that $\mathbf{v}_q = \mathbf{v}_{\max}$. Define the *distance* of a state $p \in \mathscr{B}$, denoted by η_p, as the distance of p from q in the graph induced by B. Note that $\eta_q = 0$ and all states of $\mathscr{B}$ have distance at most $|\mathscr{B}| - 1$, as $\mathscr{B}$ is strongly connected. We prove by induction that a state p with distance i satisfies $\mathbf{v}_p \geq \mathbf{v}_{\max} - 2i/x_{\min}^i$. The claim is obvious for the induction base ($i = 0$). For the induction step, let p be a state such that $\eta_p = i + 1$. Then, there is a state r such that $B_{r,p} > 0$ and $\eta_r = i$. We have

$$\begin{aligned}
\mathbf{v}_r &= (B\mathbf{v})_r + \mathbf{s}_r - t && \text{(as } \mathbf{v} \text{ is a potential)}\\
&\leq (B\mathbf{v})_p + 2 && \text{(as } \mathbf{s}_p, t \in [-1,1])\\
&= \left(B_{r,p} \cdot \mathbf{v}_p + \sum_{p' \neq p} B_{r,p'} \cdot \mathbf{v}_{p'} \right) + 2\\
&\leq B_{r,p} \cdot \mathbf{v}_p + (1 - B_{r,p}) \cdot \mathbf{v}_{\max} + 2\,.
\end{aligned}$$

By rewriting the last inequality and applying induction hypothesis to $\mathbf{v}_r$ we obtain

$$\mathbf{v}_p \geq \mathbf{v}_{\max} - \frac{\mathbf{v}_{\max} - \mathbf{v}_r + 2}{B_{r,p}} \geq \mathbf{v}_{\max} - \frac{\mathbf{v}_{\max} - \left(\mathbf{v}_{\max} - 2i/x_{\min}^i\right) + 2}{x_{\min}} \geq \mathbf{v}_{\max} - \frac{2(i+1)}{x_{\min}^{i+1}}\,.$$

This completes the induction step. Hence, we have $\mathbf{v}_{\min} \geq 0$ as desired.

It remains to show that the sequence $m^{(0)}, m^{(1)}, \ldots$ is indeed a martingale, where $\mathbf{v}$ is chosen as above. Let us fix some $i \in \mathbb{N}_0$. Obviously, $m^{(i)} \leq c^{(0)} + i + \mathbf{v}_{\max} - i \cdot t$, and hence $\mathbb{E}[m^{(i)}]$ is finite. Further, we need to prove that $\mathbb{E}[m^{(i+1)} \mid m^{(0)}, \ldots, m^{(i)}] = m^{(i)}$ almost surely. Since the values of $m^{(0)}, \ldots, m^{(i)}$ depend only on the configurations $p^{(0)}c^{(0)}, \ldots, p^{(i)}c^{(i)}$, it suffices to show that for every finite path u of length i initiated in $p^{(0)}c^{(0)}$ we have that $\mathbb{E}[m^{(i+1)} \mid Run(u)] = m^{(i)}$. Let us fix such a path u. If $c^{(j)} = 0$ for some $0 \leq j \leq i$, then for every run $w \in Run(u)$ we have $m^{(i+1)}(w) = m^{(i)}(w)$ which

[6]The matrix Z is sometimes called the *fundamental matrix* of the finite-state Markov chain induced by B.

implies $\mathbb{E}[m^{(i+1)} \mid Run(u)] = m^{(i)}$. Now assume that $c^{(j)} \geq 1$ for all $0 \leq j \leq i$. Then,

$$\begin{aligned}
\mathbb{E}\big[m^{(i+1)} \mid Run(u)\big] &= \mathbb{E}\big[c^{(i+1)} + \mathbf{v}_{p^{(i+1)}} - (i+1)\cdot t \mid Run(u)\big] \\
&= c^{(i)} + \sum_{p^{(i)} \xrightarrow{x,a}_{>0} q} x \cdot a + \sum_{p^{(i)} \xrightarrow{x,a}_{>0} q} x \cdot \mathbf{v}_q - (i+1)\cdot t \\
&= c^{(i)} + \mathbf{s}_{p^{(i)}} + (B\mathbf{v})_{p^{(i)}} - (i+1)\cdot t \\
&= m^{(i)} + \mathbf{s}_{p^{(i)}} + (B\mathbf{v})_{p^{(i)}} - \mathbf{v}_{p^{(i)}} - t \\
&= m^{(i)},
\end{aligned}$$

where the last equality holds because $\mathbf{v}$ is a potential. □

Now we have all tools needed to prove the two remaining subcases of Theorem 3.2.

LEMMA 5.6 (CASE (B)(b)(1) OF THEOREM 3.2). *Let $p, q \in Q$ such that $q \in \mathscr{B}$, $Pre^*(q(0)) \cap Post^*(p(1)) \cap \mathscr{B} \times \mathbb{N}_0$ is infinite, and the trend t of $\mathscr{B}$ satisfies $t \neq 0$. Then*

$$E(p{\downarrow}q) \leq 85000 \cdot \frac{|Q|^6}{x_{\min}^{5|Q|+|Q|^3} \cdot t^4}\,.$$

PROOF. Recall that $E(p{\downarrow}q) \cdot [p{\downarrow}q] = E_1 + E_2$, and we already have an upper bound on E_1 due to Inequality (3). Hence, we need to establish an upper bound on E_2. Assume that a run of $Run(p{\downarrow}q)$ reaches $\mathscr{B}$ in a configuration $r(k)$. First, we show that the probability of performing a terminating path of length i from $r(k)$ decays exponentially in i, and we give an explicit upper bound on this probability.

Let $r(k)$ be a configuration where $r \in \mathscr{B}$. Let $Run(r(k){\downarrow}, i)$ be the set of all runs that start with a zero-safe finite path of length i from $r(k)$ to a configuration with zero counter. Let $\mathbf{v}$ be the vector of Theorem 3.4, $\delta_{\mathbf{v}} := \mathbf{v}_{\max} - \mathbf{v}_{\min}$, and

$$a := \exp\left(-\frac{t^2}{8(\delta_{\mathbf{v}} + |t| + 1)^2}\right). \tag{5}$$

Note that $0 < a < 1$. Further, let h denote either $2(-\delta_{\mathbf{v}} - k)/t$ or $2(\delta_{\mathbf{v}} - k)/t$, depending on whether $t < 0$ or $t > 0$, respectively. We show that, for all $i \geq h$ we have that $\mathcal{P}(Run(r(k){\downarrow}, i)) \leq a^i$. Observe that all runs in $Run(r(k){\downarrow}, i)$ satisfy $m^{(i)} = \mathbf{v}_{p^{(i)}} - i \cdot t$ and hence

$$m^{(0)} - m^{(i)} = k + \mathbf{v}_r - \mathbf{v}_{p^{(i)}} + i \cdot t\,. \tag{6}$$

If $t < 0$, then, for all $i \geq h$, we obtain the following:

$$\begin{aligned}
\mathcal{P}(Run(r(k){\downarrow}, i)) &= \mathcal{P}(Run(r(k){\downarrow}, i) \ \wedge\ m^{(i)} - m^{(0)} = -k - \mathbf{v}_r + \mathbf{v}_{p^{(i)}} - i \cdot t) \\
&\leq \mathcal{P}\big(m^{(i)} - m^{(0)} = -k - \mathbf{v}_r + \mathbf{v}_{p^{(i)}} - i \cdot t\big) \\
&\leq \mathcal{P}\big(m^{(i)} - m^{(0)} \geq -k - \delta_{\mathbf{v}} - i \cdot t\big) \\
&= \mathcal{P}\big(m^{(i)} - m^{(0)} \geq (i - h/2) \cdot (-t)\big) \\
&\leq \mathcal{P}\big(m^{(i)} - m^{(0)} \geq (i/2) \cdot (-t)\big)\,.
\end{aligned}$$

Similarly, if $t > 0$, then, for all $i \geq h$, we obtain that

$$\mathcal{P}(Run(r(k){\downarrow}, i)) \leq \mathcal{P}\big(m^{(0)} - m^{(i)} \geq (i/2) \cdot t\big)\,.$$

In each step, the martingale value changes by at most $\delta_{\mathbf{v}} + |t| + 1$. Hence, Azuma's inequality asserts for $t \neq 0$ and $i \geq h$ the following:

$$\mathcal{P}(Run(r(k)\downarrow, i)) \leq \exp\left(-\frac{(i/2)^2 t^2}{2i(\delta_{\mathbf{v}} + |t| + 1)^2}\right) = a^i\,.$$

Now we derive an upper on E_2. Recall that

$$E_2 := \sum_{k_1,k_2 \in \mathbb{N}_0} \mathcal{P}\big(R^{(1)} = k_1\big) \cdot \mathcal{P}\big(R^{(2)} = k_2 \mid R^{(1)} = k_1\big) \cdot k_2\,.$$

By applying Lemma 5.1, we obtain

$$\mathcal{P}\big(R^{(1)} = k_1\big) \;\leq\; 2c^{k_1} \tag{7}$$

for all $k_1 \geq |Q|$ where $c := \exp(-x_{\min}^{|Q|}/|Q|)$. Let us fix some $k_1 \in \mathbb{N}_0$. Then

$$\begin{aligned}
&\sum_{k_2 \in \mathbb{N}_0} \mathcal{P}\big(R^{(2)} = k_2 \mid R^{(1)} = k_1\big) \cdot k_2 \\
&= \sum_{r \in \mathscr{B}} \sum_{j=0}^{k_1+1} \sum_{k_2 \in \mathbb{N}_0} \mathcal{P}\big(R^{(2)} = k_2 \mid R^{(1)} = k_1,\ Con = r(j)\big) \cdot k_2 \cdot \mathcal{P}\big(Con = r(j) \mid R^{(1)} = k_1\big) \\
&= \sum_{r \in \mathscr{B}} \sum_{j=0}^{k_1+1} \sum_{k_2 \in \mathbb{N}_0} \mathcal{P}\big(R^{(2)} = k_2 \mid Con = r(j)\big) \cdot k_2 \cdot \mathcal{P}\big(Con = r(j) \mid R^{(1)} = k_1\big)
\end{aligned}$$

In these equalities, we used the fact that in each step the counter value can increase by at most 1, thus $R^{(1)} = k_1$ and $Con(w) = r(j)$ imply $j \leq k_1 + 1$. Denote by $Con(k_1) \in \mathscr{B} \times \{0, \ldots, k_1 + 1\}$ the value of Con that maximizes $\sum_{k_2 \in \mathbb{N}_0} \mathcal{P}(R^{(2)} = k_2 \mid Con = r(j)) \cdot k_2$. Then we can continue:

$$\begin{aligned}
&\leq \sum_{k_2 \in \mathbb{N}_0} \mathcal{P}\big(R^{(2)} = k_2 \mid Con = Con(k_1)\big) \cdot k_2 \cdot \sum_{r \in \mathscr{B}} \sum_{j=0}^{k_1+1} \mathcal{P}\big(Con = r(j) \mid R^{(1)} = k_1\big) \\
&= \sum_{k_2 \in \mathbb{N}_0} \mathcal{P}\big(R^{(2)} = k_2 \mid Con = Con(k_1)\big) \cdot k_2\,.
\end{aligned}$$

Let $h(k_1) := 2(\delta_{\mathbf{v}} + k_1 + 1)/|t|$. Observe that $h(k_1) \geq 2(-\delta_{\mathbf{v}} - (k_1{+}1))/t$ or $h(k_1) \geq 2(\delta_{\mathbf{v}} - (k_1{+}1))/t$, depending on whether $t < 0$ or $t > 0$, respectively, which means that for all $k_2 \geq h(k_1)$ we have that $\mathcal{P}(R^{(2)} = k_2 \mid Con = Con(k_1)) \leq a^{k_2}$ with a defined by (5). So we can continue:

$$\leq \sum_{k_2=0}^{\lfloor h(k_1) \rfloor} k_2 + \sum_{k_2=\lceil h(k_1) \rceil}^{\infty} a^{k_2} \cdot k_2 \;\leq\; h(k_1)^2 + \frac{a}{(1-a)^2} \;=\; \frac{4(\delta_{\mathbf{v}} + k_1 + 1)^2}{t^2} + \frac{a}{(1-a)^2}\,.$$

By combining this inequality with Inequality (7), we get a bound on E_2:

$$\begin{aligned}
E_2 &= \sum_{k_1 \in \mathbb{N}_0} \mathcal{P}\big(R^{(1)} = k_1\big) \cdot \sum_{k_2 \in \mathbb{N}_0} \mathcal{P}\big(R^{(2)} = k_2 \mid R^{(1)} = k_1\big) \cdot k_2 \\
&\le \sum_{k_1=0}^{|Q|-1} \left(\frac{4(\delta_{\mathbf{v}} + k_1 + 1)^2}{t^2} + \frac{a}{(1-a)^2} \right) + \sum_{k_1=0}^{\infty} 2c^{k_1} \frac{a}{(1-a)^2} + \sum_{k_1=0}^{\infty} 2c^{k_1} \frac{4(\delta_{\mathbf{v}} + k_1 + 1)^2}{t^2} \\
&\le \frac{4|Q|(\delta_{\mathbf{v}} + |Q|)^2}{t^2} + \frac{2|Q|}{(1-c)(1-a)^2} + \frac{8}{t^2} \sum_{k_1=0}^{\infty} c^{k_1} (\delta_{\mathbf{v}} + k_1 + 1)^2,
\end{aligned}$$

where $c := \exp(-x_{\min}^{|Q|}/|Q|)$. The last series can be bounded as follows:

$$\begin{aligned}
\sum_{k_1=0}^{\infty} c^{k_1} (\delta_{\mathbf{v}} + k_1 + 1)^2 &\le \sum_{k_1=0}^{\lfloor \delta_{\mathbf{v}}+1 \rfloor} (2(\delta_{\mathbf{v}} + 1))^2 + \sum_{k_1=\lfloor \delta_{\mathbf{v}}+1 \rfloor + 1}^{\infty} c^{k_1} \cdot (2k_1)^2 \\
&\le 4(\delta_{\mathbf{v}} + 2)^3 + 4 \sum_{k_1=0}^{\infty} c^{k_1} \cdot k_1^2 = 4(\delta_{\mathbf{v}} + 2)^3 + 4 \frac{c(c+1)}{(1-c)^3} \\
&\le 4(\delta_{\mathbf{v}} + 2)^3 + \frac{8}{(1-c)^3}
\end{aligned}$$

It follows:

$$E_2 \le \frac{4|Q|(\delta_{\mathbf{v}} + |\mathscr{B}|)^2}{t^2} + \frac{2|Q|}{(1-c)(1-a)^2} + \frac{32}{t^2} \left((\delta_{\mathbf{v}} + 2)^3 + \frac{2}{(1-c)^3} \right). \tag{8}$$

Recall the following bounds:

$$\begin{aligned}
\delta_{\mathbf{v}} &\le 2|\mathscr{B}|/x_{\min}^{|\mathscr{B}|} && \text{(Theorem 3.4)}, \\
1 - c = 1 - \exp\big(-x_{\min}^{|Q|}/|Q|\big) &\ge x_{\min}^{|Q|}/(2|Q|) && \text{(Lemma 5.1)}, \\
1 - a \ge 1 - \exp\big(-t^2/\big(8(\delta_{\mathbf{v}} + 2)^2\big)\big) &\ge t^2/\big(16(\delta_{\mathbf{v}} + 2)^2\big) && \text{(by } |t| \le 1), \\
[p{\downarrow}q] &\ge x_{\min}^{|Q|^3} && \text{(Proposition 2.3 (B))}.
\end{aligned}$$

After plugging these bounds into Inequality (8) we obtain

$$E_2 \;\le\; 84356 \frac{|Q|^6}{x_{\min}^{5|Q|} \cdot t^4}. \tag{9}$$

Hence, by combining Inequalities (3) and (9), we finally obtain

$$E(p{\downarrow}q) \;=\; \frac{E_1 + E_2}{[p{\downarrow}q]} \;\le\; 85000 \cdot \frac{|Q|^6}{x_{\min}^{5|Q| + |Q|^3} \cdot t^4}. \quad \square$$

LEMMA 5.7 (CASE (B)(b)(2) OF THEOREM 3.2). *Let $p, q \in Q$ such that $q \in \mathscr{B}$, $Pre^*(q(0)) \cap Post^*(p(1)) \cap \mathscr{B} \times \mathbb{N}_0$ is infinite, and the trend t of $\mathscr{B}$ satisfies $t = 0$. Then, $E(p{\downarrow}q) = \infty$.*

PROOF. We start by introducing some notation. For every configuration $r(k) \in \mathscr{B} \times \mathbb{N}$ and every $\ell \in \mathbb{N}_0$ such that $k > \ell$, let $Run(r(k){\Downarrow}\ell)$ be the set of all runs initiated in $r(k)$ that visit a configuration with counter value equal to ℓ. Further, let $R_{r(k)\Downarrow\ell}$ be a random variable which for every $w \in Run(r(k))$ returns either the least i such that the counter value in $w(i)$ is equal to ℓ, or $\perp$ if there is no such i.We use $E(r(k){\Downarrow}\ell)$ to denote the conditional expected value $\mathbb{E}[R_{r(k)\Downarrow\ell} \mid Run(r(k)\Downarrow\ell)]$. The set of all runs $w \in Run(r(k))$ that start with a zero-safe finite path from $r(k)$ to $q(0)$ is denoted by $Run(r(k)\downarrow q)$.

We prove the following:

(a) There is $b_1 \in \mathbb{N}$ such that for every $r(k) \in \mathscr{B} \times \mathbb{N}$ where $k \geq b_1$ we have that $E(r(k)\Downarrow 0) = \infty$. Consequently, $E(r(m{+}b_1)\Downarrow m) = \infty$ for every $m \in \mathbb{N}_0$.
(b) There is $b_2 \in \mathbb{N}$ such that for every configuration $r(k) \in Pre^*(q(0)) \cap \mathscr{B} \times \mathbb{N}$ where $k \geq b_2$ we have that if $r(k) \to s(\ell)$, then $s(\ell) \in Pre^*(q(0)) \cap \mathscr{B} \times \mathbb{N}$.

First, we show that Claims (a) and (b) together imply $E(p\downarrow q) = \infty$. Since $Pre^*(q(0)) \cap Post^*(p(1)) \cap \mathscr{B} \times \mathbb{N}_0$ is infinite, there is $r(k) \in Pre^*(q(0)) \cap Post^*(p(1)) \cap \mathscr{B} \times \mathbb{N}_0$ such that $k = b_1 + b_2$, where b_1 and b_2 are the bounds of Claim (a) and (b), respectively. Our aim is to prove that

$$\sum_{k_2 \in \mathbb{N}_0} \mathcal{P}\big(R^{(2)} = k_2 \mid Con = r(k)\big) \cdot k_2 \;=\; \infty.$$

This immediately implies $E_2 = \infty$ and hence also $E(p\downarrow q) = (E_1 + E_2)/[p\downarrow q] = \infty$. Let $D^{(1)}, D^{(2)}$ be random variables over $Run(r(k))$ defined as follows:

—$D^{(1)}(w)$ is the least $\ell \in \mathbb{N}_0$ such that the counter value in $w(\ell)$ is equal to b_2; if there is no such ℓ, we put $D^{(1)}(w) = \perp$;
—$D^{(2)}(w)$ is the least $\ell \in \mathbb{N}_0$ such that $w(D^{(1)}(w)), \ldots, w(D^{(1)}(w){+}\ell)$ is a zero-safe finite path from $w(D^{(1)}(w))$ to $q(0)$; if there is no such ℓ or $D^{(1)}(w) = \perp$, we put $D^{(2)}(w) = \perp$.

Further, let C be the set of all $t \in \mathscr{B}$ such that there is a finite path from $r(k)$ to $t(b_2)$ where the counter stays strictly above b_2 before the visit to $t(b_2)$. Note that for each $t \in C$, we have $t(b_2) \in Pre^*(q(0))$ by Claim (b); hence, $M(b_2) := \min_{t \in C} \mathcal{P}(Run(t(b_2)\downarrow q))$ is positive. Then, we have the following:

$$\begin{aligned}
\sum_{k_2 \in \mathbb{N}_0} \mathcal{P}\big(R^{(2)}{=}k_2 \mid Con{=}r(k)\big) \cdot k_2 &= \sum_{\ell_1, \ell_2 \in \mathbb{N}_0} \mathcal{P}\big(D^{(1)}{=}\ell_1 \wedge D^{(2)}{=}\ell_2 \mid Run(r(k)\downarrow q)\big) \cdot (\ell_1 + \ell_2) \\
&\geq \sum_{\ell_1 \in \mathbb{N}_0} \mathcal{P}\big(D^{(1)}{=}\ell_1 \mid Run(r(k)\downarrow q)\big) \cdot \ell_1 \\
&\geq \sum_{\ell_1 \in \mathbb{N}_0} \frac{\mathcal{P}(D^{(1)}{=}\ell_1) \cdot M(b_2)}{\mathcal{P}(Run(r(k)\downarrow q))} \cdot \ell_1 \\
&= \frac{M(b_2)}{\mathcal{P}(Run(r(k)\downarrow q))} \cdot \sum_{\ell_1 \in \mathbb{N}_0} \mathcal{P}\big(D^{(1)}{=}\ell_1\big) \cdot \ell_1 \\
&= \frac{M(b_2)}{\mathcal{P}(Run(r(k)\downarrow q))} \cdot E(r(k)\Downarrow b_2) = \infty.
\end{aligned}$$

In the last step, we use Claim (a) to conclude $E(r(k)\Downarrow b_2) = \infty$.

So, it remains to prove Claims (a) and (b). For Claim (a), let us first realize that *every* configuration $r(k) \in \mathscr{B} \times \mathbb{N}$ satisfies $\mathcal{P}(Run(r(k)\Downarrow 0)) > 0$. Since $Pre^*(q(0)) \cap \mathscr{B} \times \mathbb{N}_0$ is infinite, there is $s \in \mathscr{B}$ such that $s(i) \in Pre^*(q(0))$ for infinitely many i's, which means that $\mathcal{P}(Run(s(i)\Downarrow 0)) > 0$ for *every* $i \in \mathbb{N}_0$. Since $\mathscr{B}$ is strongly connected, we can fix the shortest path from r to s in the finite-state Markov chain $\mathcal{X}$, and follow this path from $r(k)$. Thus, we either visit a configuration with zero counter, or enter a configuration $s(i)$ for some i. In both cases, we have that $\mathcal{P}(Run(r(k)\Downarrow 0)) > 0$.

Now let $b_1 := \lceil \mathbf{v}_{\max} - \mathbf{v}_{\min} + 1 \rceil$, where $\mathbf{v}_{\max}$ and $\mathbf{v}_{\min}$ are the constants introduced in Theorem 3.4. Let us fix some $r(k) \in \mathscr{B} \times \mathbb{N}$ where $k \geq b_1$, and consider the martingale $m^{(0)}, m^{(1)}, \ldots$ defined in Theorem 3.4 over $Run(r(k))$ (note that as $t = 0$, the term $i \cdot t$ vanishes from the definition of the martingale). For every $n \geq k+1$, we define a *stopping time* τ which returns the first point in time in which either $m^{(\tau)} \geq \mathbf{v}_{\max} + n$, or $m^{(\tau)} \leq \mathbf{v}_{\max}$. For the remaining runs of $U := \{w \in Run(r(k)) \mid \mathbf{v}_{\max} < m^{(i)}(w) < \mathbf{v}_{\max} + n \text{ for all } i \in \mathbb{N}_0\}$, the variable τ returns ∞. Observe that $\mathcal{P}(U) = 0$, that is, $\mathbb{E}[\tau] < \infty$ as required by the optional stopping theorem. To see this, realize that all runs of U visit only configurations $s(\ell)$ such that $\ell < \delta_{\mathbf{v}} + n$ (where $\delta_{\mathbf{v}} := \mathbf{v}_{\max} - \mathbf{v}_{\min}$), because the martingale value would reach at least $\mathbf{v}_{\max} + n$ otherwise. Since $\mathcal{P}(Run(s(\ell){\Downarrow}0)) > 0$ (see the preceding text), we obtain that almost all runs of U visit a configuration with zero counter, where the martingale value is at most $\mathbf{v}_{\max}$. Thus, we obtain $\mathcal{P}(U) = 0$.

By applying the optional stopping theorem, we obtain that $\mathbb{E}(m^{(\tau)}) = \mathbb{E}(m^{(0)}) = \mathbf{v}_r + k$. Since the martingale value changes by at most $\delta_{\mathbf{v}} + 1$ in each step (where $\delta_{\mathbf{v}} := \mathbf{v}_{\max} - \mathbf{v}_{\min}$), we further obtain that if $m^{(\tau)} \geq \mathbf{v}_{\max} + n$, then $m^{(\tau)} \leq \mathbf{v}_{\max} + n + \delta_{\mathbf{v}} + 1$. Hence,

$$\mathbf{v}_r + k \;=\; \mathbb{E}\big(m^{(\tau)}\big) \;\leq\; \mathcal{P}\big(m^{(\tau)} \geq \mathbf{v}_{\max} + n\big) \cdot (\mathbf{v}_{\max} + n + \delta_{\mathbf{v}} + 1) + \mathcal{P}\big(m^{(\tau)} \leq \mathbf{v}_{\max}\big) \cdot \mathbf{v}_{\max}.$$

From this, we get

$$\mathcal{P}\big(m^{(\tau)} \geq \mathbf{v}_{\max} + n\big) \geq \frac{\mathbf{v}_r + k - \mathbf{v}_{\max}}{n + \mathbf{v}_{\max} + \delta_{\mathbf{v}} + 1}\,. \tag{10}$$

Note that $m^{(\tau)} \geq \mathbf{v}_{\max} + n$ implies $c^{(\tau)} = m^{(\tau)} - \mathbf{v}_{p^{(\tau)}} \geq \mathbf{v}_{\max} + n - \mathbf{v}_{p^{(\tau)}} \geq n$, and thus also $R_{r(k){\Downarrow}0} \geq n$, because at least n steps are required to decrease the counter value from n to 0. It follows that $\mathcal{P}(m^{(\tau)} \geq \mathbf{v}_{\max} + n) \leq \mathcal{P}(R_{r(k){\Downarrow}0} \geq n)$. By combining this inequality with Inequality (10), we have

$$\sum_{n \in \mathbb{N}} \mathcal{P}\big(R_{r(k){\Downarrow}0} \geq n\big) \geq \sum_{n=k+1}^{\infty} \mathcal{P}\big(R_{r(k){\Downarrow}0} \geq n\big) \geq \sum_{n=k+1}^{\infty} \frac{\mathbf{v}_r + k - \mathbf{v}_{\max}}{n + \mathbf{v}_{\max} + \delta_{\mathbf{v}} + 1} = \infty\,. \tag{11}$$

Using Inequality (11), we finally obtain

$$E(r(k){\Downarrow}0) = \sum_{n \in \mathbb{N}} \mathcal{P}\big(R_{r(k){\Downarrow}0} \geq n \mid Run(r(k){\Downarrow}0)\big) = \frac{\sum_{n \in \mathbb{N}} \mathcal{P}(R_{r(k){\Downarrow}0} \geq n)}{\mathcal{P}(Run(r(k){\Downarrow}0))} = \infty\,.$$

Now we prove Claim (b). We start by observing that $Pre^*(q(0))$ has an “ultimately periodic” structure. For every $i \in \mathbb{N}_0$, let $Pre(i) = \{s \in \mathscr{B} \mid s(i) \in Pre^*(q(0))\}$. Note that, if $Pre(i) = Pre(j)$ for some $i, j \in \mathbb{N}_0$, then also $Pre(i{+}1) = Pre(j{+}1)$. Let m_1 be the least index such that $Pre(m_1) = Pre(j)$ for some $j > m_1$, and let m_2 be the least $j > m_1$ such that $Pre(m_1) = Pre(j)$. Further, we put $m = m_2 - m_1$. Observe that $m_1, m_2 \leq 2^{|\mathscr{B}|}$, and for every $\ell \geq m_2$, we have that $Pre(\ell) = Pre(\ell{+}m)$.

For every configuration $s(i) \in \mathscr{B} \times \mathbb{N}_0$, let $C(s(i))$ be the set of all configurations $s(i+j)$ such that $0 \leq j < m$ and $s \in Pre(i{+}j)$. Note that $C(s(i))$ has at most m elements, and we define the *index* of $s(i)$ as the cardinality of $C(s(i))$. Due the periodicity of $Pre^*(q(0))$, we immediately obtain that for every $s(i)$ and $j \in \mathbb{N}_0$, where $i \geq m_1$, the index of $s(i)$ is the same as the index of $s(i{+}j)$.

Let $b_2 := m_1 + |\mathscr{B}| + 1$, and assume that there is a transition $r(k) \to s(\ell)$ such that $r \in Pre(k)$, $s \notin Pre(\ell)$, and $k \geq b_2$. Then, $r(k{+}j) \to s(\ell{+}j)$ for all $0 \leq j < m$. Obviously, if $s \in Pre(\ell{+}j)$, then also $r \in Pre(k{+}j)$, which means that the index of $s(\ell)$ is *strictly smaller* than the index of $r(k)$. Since $\mathscr{B}$ is strongly connected, there is a finite path from $s(\ell)$ to $r(n)$ of length at most $|\mathscr{B}|$, where $n \geq m_1$. This means that there is a finite path from $s(\ell+j)$ to $r(n+j)$ for every $0 \leq j < m$. Hence, the index of $s(\ell)$ is at least as large as the index of $r(n)$. Since the indexes of $r(n)$ and $r(k)$ are the same, we have a contradiction. □

5.2. Proofs of Section 3.2

We start by recalling a standard result of numerical analysis (see, e.g., Isaacson and Keller [1966][7]).

THEOREM 5.8. *Consider a system of linear equations, $C \cdot \vec{U} = \mathbf{c}$, where $C \in \mathbb{R}^{n\times n}$ and $\mathbf{c} \in \mathbb{R}^n$. Suppose that C is nonsingular and $\mathbf{c} \neq \mathbf{0}$. Let $\vec{U}^* = C^{-1} \cdot \mathbf{c}$ be the unique solution of this system (note that $\vec{U}^* \neq \mathbf{0}$). Denote by $\kappa(C) = \|C\| \cdot \|C^{-1}\|$ the condition number of C. Consider a system of equations $(C + \Delta) \cdot \vec{U} = \mathbf{c} + \mathbf{d}$ where $\Delta \in \mathbb{R}^{n\times n}$ and $\mathbf{d} \in \mathbb{R}^n$. If $\|\Delta\| < \frac{1}{\|C^{-1}\|}$, then the system $(C + \Delta) \cdot \vec{U} = \mathbf{c} + \mathbf{d}$ has a unique solution $\vec{U}^*_p$. Moreover, for every $\delta > 0$ satisfying $\frac{\|\Delta\|}{\|C\|} \leq \delta$ and $\frac{\|\mathbf{d}\|}{\|\mathbf{c}\|} \leq \delta$ and $4 \cdot \delta \cdot \kappa(C) < 1$, the solution $\vec{U}^*_p$ satisfies*

$$\frac{\|\vec{U}^* - \vec{U}^*_p\|}{\|\vec{U}^*\|} \leq 4 \cdot \delta \cdot \kappa(C).$$

Using Theorem 5.8, we prove the following proposition.

PROPOSITION 5.9. *Consider a system of linear equations, $B \cdot \vec{V} = \mathbf{b}$, where $B \in \mathbb{R}^{n\times n}$ and $\mathbf{b} \in \mathbb{R}^n$. Suppose that B is nonsingular and $\mathbf{b} \neq \mathbf{0}$. Let $\vec{V}^* = B^{-1} \cdot \mathbf{b}$ be the unique solution of this system. Consider a system $(B+\mathcal{E})\cdot\vec{V} = \mathbf{b}$ where $\mathcal{E} \in \mathbb{R}^{n\times n}$. Let $\|B\| \leq u \geq 1$ and $\|B^{-1}\| \leq v \geq 1$. If $\|\mathcal{E}\| < 1/v$, then the system $(B + \mathcal{E}) \cdot \vec{V} = \mathbf{b}$ has a unique solution $\vec{W}^*$. Moreover, if $\|\mathcal{E}\| \leq \delta < 1/(4uv)$, then $\vec{W}^*$ satisfies*

$$\frac{\|\vec{V}^* - \vec{W}^*\|}{\|\vec{V}^*\|} \leq \delta \cdot 4uv.$$

PROOF. We apply Theorem 5.8 with

$$C := \begin{pmatrix} B & 0 \\ 0 & 1 \end{pmatrix} \qquad \text{and} \qquad \mathbf{c} := \begin{pmatrix} \mathbf{b} \\ 1 \end{pmatrix} \qquad \text{and} \qquad \Delta := \begin{pmatrix} \mathcal{E} & 0 \\ 0 & 0 \end{pmatrix};$$

that is, a single equation $x = 1$, for a new variable x is added to the system, without new errors. Notice that

$$C^{-1} = \begin{pmatrix} B^{-1} & 0 \\ 0 & 1 \end{pmatrix} \qquad \text{and} \qquad \vec{U}^* := \begin{pmatrix} \vec{V}^* \\ 1 \end{pmatrix}.$$

Further $\|C^{-1}\| = \max\{1, \|B^{-1}\|\}$. So we have $\|\Delta\| = \|\mathcal{E}\| < 1/v \leq 1/\max\{1, \|B^{-1}\|\} = 1/\|C^{-1}\|$. Thus, by Theorem 5.8, there is a unique solution of $(C + \Delta) \cdot \vec{U} = \mathbf{c}$, hence $\vec{W}^*$ is unique too. Moreover, we have

$$\frac{\|\Delta\|}{\|C\|} = \frac{\|\Delta\|}{\max\{1, \|B\|\}} \leq \|\Delta\| = \|\mathcal{E}\| \leq \delta \qquad \text{and}$$

$$4 \cdot \delta \cdot \kappa(C) = 4 \cdot \delta \cdot \max\{1, \|B\|\} \cdot \max\{1, \|B^{-1}\|\} \leq 4 \cdot \delta \cdot u \cdot v < 1,$$

so Theorem 5.8 implies

$$\frac{\|\vec{V}^* - \vec{W}^*\|}{\|\vec{V}^*\|} \leq 4 \cdot \delta \cdot \kappa(C) \leq \delta \cdot 4uv. \quad \square$$

[7]We use a slightly modified version of Theorem 3 presented in Chapter 2.1.2 of Isaacson and Keller [1966].

Now we have all tools needed to prove Proposition 3.6.

PROPOSITION 3.6. *Let* $b \geq \max\{E(p{\downarrow}q) \mid (p,q) \in T^{>0}_{<\infty}\}$. *For every* ε *such that* $0 < \varepsilon < 1$, *let* $\delta = \varepsilon/(12 \cdot b^2)$. *If* $\|G - H\| \leq \delta$, *then the perturbed system* $\vec{V} = G \cdot \vec{V} + \mathbf{1}$ *has a unique solution* $\vec{F}$ *such that* $|E(p{\downarrow}q) - \vec{F}_{pq}| \leq \varepsilon$ *for all* $(p,q) \in T^{>0}_{<\infty}$. *Here* $\vec{F}_{pq}$ *is the component of* $\vec{F}$ *corresponding to the variable* $V(p{\downarrow}q)$.

PROOF. Denote by $\vec{E}$ the vector of all finite $E(p{\downarrow}q)$, that is, $\vec{E} = (I-H)^{-1}\mathbf{1}$. We apply Proposition 5.9 using the following assignments: $B = I - H$, $B + \mathcal{E} = I - G$, $\mathbf{b} = \mathbf{1}$, $\vec{V}^* = \vec{E}$, and $\vec{W}^* = \vec{F}$. To find a suitable u, we need to find a bound on $\|I - H\|$. By comparing $\mathcal{L}$ with Equality (1), it follows that $\|H\mathbf{1}\| \leq 2$ and hence

$$\|I - H\| \leq 1 + \|H\| = 1 + \|H\mathbf{1}\| \leq 3\,. \tag{12}$$

Hence, we set $u := 3$. Further, we set $v := b$, so we need to show $\|(I-H)^{-1}\| \leq b$. By our assumption, $\|\vec{E}\| \leq b$. Recall that $\vec{E} = (I-H)^{-1}\mathbf{1}$, so if $(I-H)^{-1}$ is nonnegative, then $\|(I-H)^{-1}\| = \|(I-H)^{-1}\mathbf{1}\| = \|\vec{E}\| \leq b$, hence it remains to show that $(I-H)^{-1}$ is nonnegative. To see this, note that $\vec{E}$ is the (unique) fixed point of a linear function $\mathcal{F}$ which to every $\vec{V}$ assigns $H \cdot \vec{V} + \mathbf{1}$. This function is continuous and monotone, so by Kleene's theorem we get that $\vec{E} = \sup_{i \in \mathbb{N}} \mathcal{F}^i(\mathbf{0}) = \sum_{i=0}^{\infty} H^i \mathbf{1}$. Recall that $\vec{E}$ is finite, so the matrix series $H^* := \sum_{i=0}^{\infty} H^i$ converges and thus equals $(I-H)^{-1}$. Hence $(I-H)^{-1} = H^*$, which is nonnegative as H is nonnegative.

Now we are ready to apply Proposition 5.9. Since $\|G - H\| \leq \varepsilon/(12 \cdot b^2) < 1/v$, the perturbed system $\vec{V} = G \cdot \vec{V} + \mathbf{1}$ has a unique solution $\vec{F}$ as desired. By applying the second part of Proposition 5.9, we get

$$\frac{\|\vec{E} - \vec{F}\|}{\|\vec{E}\|} \leq \delta \cdot 12 \cdot b \qquad \text{for } \|G - H\| \leq \delta \leq 1/(12 \cdot b). \tag{13}$$

Hence,

$$\begin{aligned}
|E(p{\downarrow}q) - \vec{F}_{pq}| &\leq \|\vec{E} - \vec{F}\| && \text{(by the definition of the norm)}\\
&\leq b \cdot \frac{\|\vec{E} - \vec{F}\|}{\|\vec{E}\|} && \text{(by } \|\vec{E}\| \leq b\text{)}\\
&\leq b \cdot \delta \cdot 12 \cdot b && \text{(by (13))}\\
&= \varepsilon && \text{(by the definition of } \delta\text{)}. \quad \square
\end{aligned}$$

5.3. Proofs of Section 4

Recall that we assume a fixed 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 $\mathscr{A}$. Also recall the finite-state Markov chain $\mathcal{G}$ introduced in Section 4, and the system $(I - A)\vec{V} = \mathbf{b}$ of linear equations whose unique solution $\vec{V}^*$ in $\mathbb{R}^{|S_?|}$ is the vector of probabilities of reaching a good BSCC of $\mathcal{G}$ from the states of $S_?$. We show how to compute a relative ε-approximation of $\vec{V}^*$ in time polynomial in $|\mathscr{A}|$ and $\log(1/\varepsilon)$, assuming that Claim (b) of Section 4 has already been proven.

LEMMA 5.11. *Let* $c = 2|Q|$. *For every* $s \in S_?$, *let* R_s *be the probability of visiting a BSCC of* $\mathcal{G}$ *from* s *in at most* c *transitions, and let* $R = \min\{R_s \mid s \in S_?\}$. *Then* $R \geq x_{\min}^{6|Q|^4} \cdot t^3_{\min^+}/(7000 \cdot |Q|)$ *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* $(I-A')\vec{V} = \vec{b}'$ *has a unique solution* $\vec{U}^*$ *such that* $|\vec{V}_s^* - \vec{U}_s^*|/\vec{V}_s^* \leq \varepsilon$ *for every* $s \in S_?$.

Proof. The first step towards applying Theorem 5.8 is to estimate the condition number $\kappa = \|I-A\| \cdot \|(I-A)^{-1}\|$. Obviously, $\|I-A\| \leq 2$. Further, $\|(I-A)^{-1}\|$ is bounded by the expected number of steps needed to reach a BSCC of $\mathcal{G}$ from a state of $S_?$ (here we use a standard result about absorbing finite-state Markov chains). Since $S_?$ has at most c states, we have that $R_s > 0$, and hence $R \geq x_{\min}^{6|Q|^4} \cdot t_{\min^+}^3/(7000 \cdot |Q|)$ by applying the lower bounds of Claim (b). Obviously, the probability of *not visiting* a BSCC of $\mathcal{G}$ in at most i transitions from a state of $S_?$ is bounded by $(1-R)^{\lfloor i/c \rfloor}$. Hence, the probability of visiting a BSCC of $\mathcal{G}$ from a state of $S_?$ after *exactly* i transitions is bounded by $(1-R)^{\lfloor (i-1)/c \rfloor}$. Further, a simple calculation shows that

$$\begin{aligned}
\|(I-A)^{-1}\| &\leq \sum_{i=1}^{\infty} i \cdot (1-R)^{\lfloor (i-1)/c \rfloor} = \sum_{j=0}^{\infty} \sum_{i=jc+1}^{(j+1)c} i \cdot (1-R)^j \\
&= \sum_{j=0}^{\infty} (1-R)^j \sum_{i=jc+1}^{(j+1)c} i = \sum_{j=0}^{\infty} (1-R)^j \cdot \left(\frac{c(c+1)}{2} + jc^2 \right) \\
&= \frac{c(c+1)}{2R} + \frac{c^2(1-R)}{R^2} \leq \left(\frac{c+1}{R} \right)^2 .
\end{aligned}$$

Hence, $\kappa \leq 2(c+1)^2/R^2$. Let $\vec{V}^*$ be the unique solution of $(I-A)\vec{V} = \vec{b}$. Since $\|\vec{V}^*\| \leq 1$ and $\vec{V}_s^* \geq R$ for every $s \in S_?$, it suffices to compute an approximate solution $\vec{U}^*$ such that

$$\frac{\|\vec{V}^* - \vec{U}^*\|}{\|\vec{V}^*\|} \leq \varepsilon \cdot R.$$

By Theorem 5.8, we have that

$$\frac{\|\vec{V}^* - \vec{U}^*\|}{\|\vec{V}^*\|} \leq 4\tau\kappa \leq \frac{8\tau(c+1)^2}{R^2},$$

where τ is the relative error of A and $\vec{b}$. Hence, it suffices to choose τ so that

$$\tau \leq \frac{\varepsilon R^3}{8(c+1)^2}$$

and compute all transition probabilities in $\mathcal{G}$ up to the relative error τ. Note that the approximation A' of the matrix A which is obtained in this way is still nonsingular, because

$$\|A - A'\| \leq \tau \leq \frac{\varepsilon R^3}{8(c+1)^2} < \frac{R^2}{(c+1)^2} \leq \frac{1}{\|(I-A)^{-1}\|} . \quad \square$$

Now we prove the divergence gap theorem (i.e., Theorem 4.8). Let us fix a pOC $\hat{\mathscr{A}}$ with Q as the set of control states where the least positive probability used in the rules of $\hat{\mathscr{A}}$ is at least $x_{\min}$. Recall the underlying finite-state Markov chain $\hat{\mathcal{X}}$ introduced in Section 3.1. The technical core of our proof are some observations about the runs in BSCCs of $\hat{\mathcal{X}}$ with positive trend. For the rest of this section, we fix a BSCC $\mathscr{B}$ of $\hat{\mathcal{X}}$ *such that the trend* t *of* $\mathscr{B}$ *is positive*, and we use B to denote the transition matrix of $\mathscr{B}$.

LEMMA 5.12. *Assume* $[p\uparrow] < 1$ *for all* $p \in \mathscr{B}$*, and let* $q(k) \in \mathscr{B} \times \mathbb{N}$ *be a configuration such that the control state* q *satisfies* $\mathbf{v}_q = \mathbf{v}_{\max}$*, where* $\mathbf{v}$ *is the vector of Theorem* 3.4. *Let* $b \in \mathbb{N}$*, and let* $Run(q(k)\rightarrow b)$ *be the set of all* $w \in Run(q(k))$ *such that* w *visits a configuration with counter value at least* b *and the counter stays positive in all configurations preceding this visit. Then,* $\mathcal{P}(Run(q(k)\rightarrow b)) \geq 1/(b+1+\delta_{\mathbf{v}})$.

PROOF. If $k \geq b$, the lemma holds trivially. Now assume that $k < b$. We define a *stopping time* τ over $Run(q(k))$ (cf. Section 3.1) as follows:

$$\tau := \inf\{i \in \mathbb{N}_0 \mid m^{(i)} \leq \mathbf{v}_{\max} \ \vee \ \ m^{(i)} \geq b + \mathbf{v}_{\max}\}.$$

Here $m^{(0)}, m^{(1)}, \ldots$ is the martingale of Theorem 3.4. Note that $1 + \mathbf{v}_{\max} \leq m^{(0)} < b + \mathbf{v}_{\max}$, that is, $\tau \geq 1$. Let E be the set of all $w \in Run(q(k))$ where $\tau(w) < \infty$ and $m^{(\tau)}(w) \geq b + \mathbf{v}_{\max}$, that is, E is the event that the martingale $m^{(i)}$ reaches a value of $b + \mathbf{v}_{\max}$ or higher without previously reaching a value of $\mathbf{v}_{\max}$ or lower. Similarly, let D be the set of all $w \in Run(q(k))$ such that the counter reaches a value of b or higher without previously hitting 0. To prove the lemma, we need to show $\mathcal{P}(D) \geq 1/(b+1+\delta_{\mathbf{v}})$. We will do that by showing that $E \subseteq D$ and $\mathcal{P}(E) \geq 1/(b+1+\delta_{\mathbf{v}})$.

First, we show $E \subseteq D$. Consider any run $w \in E$, that is, $m^{(\tau)}(w) \geq b + \mathbf{v}_{\max}$ and $m^{(i)}(w) > \mathbf{v}_{\max}$ for all $i \leq \tau$. So, for all $i \leq \tau$, we have

$$m^{(i)}(w) = c^{(i)}(w) + \mathbf{v}_{p^{(i)}(w)} - i \cdot t > \mathbf{v}_{\max},$$

implying $c^{(i)}(w) > 0$. Similarly, $m^{(\tau)}(w) = c^{(\tau)}(w) + \mathbf{v}_{p^{(\tau)}(w)} - \tau(w) \cdot t \geq b + \mathbf{v}_{\max}$, which means $c^{(\tau)}(w) \geq b$ and hence $w \in D$. It remains to show $\mathcal{P}(E) \geq 1/(b+1+\delta_{\mathbf{v}})$.

Next we argue that $\mathbb{E}[\tau]$ is finite, that is, τ is indeed a stopping time. Since $[p\uparrow] < 1$ for all $p \in \mathscr{B}$, there are constants $\ell \in \mathbb{N}$ and $x \in (0, 1]$ such that, given any initial configuration $p(c) \in \mathscr{B} \times \mathbb{N}$, the probability of decreasing the counter by 1 in at most ℓ steps is at least x. Since $\mathscr{B}$ is strongly connected, it follows that there are constants $\ell' \in \mathbb{N}$ and $x' \in (0, 1]$ such that, given any configuration $p(c) \in \mathscr{B} \times \mathbb{N}$, the probability of reaching a configuration with zero counter or a configuration $p(c-b)$ in at most ℓ' steps is at least x'. It follows that whenever $m^{(i)} < b + \mathbf{v}_{\max}$, the probability that there is $j \leq \ell'$ with $m^{(i+j)} \leq \mathbf{v}_{\max}$ is at least x'. Hence, we have

$$\mathbb{E}[\tau] \ = \ \sum_{j=0}^{\infty} \mathcal{P}(\tau > j) \ \leq \ \ell' \sum_{j=0}^{\infty} \mathcal{P}(\tau > \ell' \cdot j) \ \leq \ \ell' \sum_{j=0}^{\infty} (1-x')^j \ = \frac{\ell'}{x'} \ < \ \infty.$$

Consequently, the *optional stopping theorem* (cf., Section 3.1) is applicable and asserts

$$\mathbb{E}\big[m^{(\tau)}\big] = \mathbb{E}\big[m^{(0)}\big] = m^{(0)} \geq 1 + \mathbf{v}_{\max}\,. \tag{14}$$

For all runs in E, we have $m^{(\tau-1)} < b + \mathbf{v}_{\max}$. Since the value of $m^{(i)}$ can increase by at most $1 + \delta_{\mathbf{v}}$ in a single step, we have $m^{(\tau)} \leq b + \mathbf{v}_{\max} + 1 + \delta_{\mathbf{v}}$ for all runs in E. It follows that

$$\begin{aligned}\mathbb{E}\left[m^{(\tau)}\right] \ &\leq \ \mathcal{P}(E) \cdot (b + \mathbf{v}_{\max} + 1 + \delta_{\mathbf{v}}) + (1 - \mathcal{P}(E)) \cdot \mathbf{v}_{\max} \\ &= \ \mathbf{v}_{\max} + \mathcal{P}(E) \cdot (b + 1 + \delta_{\mathbf{v}})\,.\end{aligned}$$

Combining this inequality with (14) yields $\mathcal{P}(E) \geq 1/(b+1+\delta_{\mathbf{v}})$. This completes the proof. □

Recall that for every configuration $q(k) \in \mathscr{B} \times \mathbb{N}$, we use $Run(q(k){\Downarrow}0)$ to denote the set of all runs initiated in $q(k)$ that visit a configuration with zero counter. Further, $R_{q(k)\Downarrow 0}$ is a random variable which for every $w \in Run(q(k))$ returns either the least i such that the counter value in $w(i)$ is zero, or $\perp$ if there is no such i. The following lemma gives an upper bound on $\mathcal{P}(Run(q(k){\Downarrow}0))$.

LEMMA 5.13. *Let $q(k) \in \mathscr{B} \times \mathbb{N}$ such that $k \geq \delta_{\mathbf{v}}$. Then, $\mathcal{P}(Run(q(k){\Downarrow}0)) \leq a^k/(1-a)$, where*

$$a := \exp\left(-\frac{t^2}{2(\delta_{\mathbf{v}}+t+1)^2}\right).$$

Note that $0 < a < 1$. Further, if $k \geq 6(\delta_{\mathbf{v}}+t+1)^3/t^3$, then $\mathcal{P}(Run(q(k){\Downarrow}0)) \leq 1/2$.

PROOF. Clearly, $\mathcal{P}(Run(q(k){\Downarrow}0)) = \sum_{i=k}^{\infty} \mathcal{P}(R_{q(k)\Downarrow 0} = i)$. For all runs $w \in Run(q(k))$ such that $R_{q(k)\Downarrow 0}(w) = i$, we have $m^{(i)} = \mathbf{v}_{p^{(i)}} - i \cdot t$ and so

$$m^{(0)} - m^{(i)} = k + \mathbf{v}_q - \mathbf{v}_{p^{(i)}} + i \cdot t.$$

It follows that

$$\begin{aligned}
\mathcal{P}\big(R_{q(k)\Downarrow 0} = i\big) &= \mathcal{P}\big(R_{q(k)\Downarrow 0} = i \ \wedge \ m^{(0)} - m^{(i)} = k + \mathbf{v}_q - \mathbf{v}_{p^{(i)}} + i \cdot t\big) \\
&\leq \mathcal{P}\big(m^{(0)} - m^{(i)} = k + \mathbf{v}_q - \mathbf{v}_{p^{(i)}} + i \cdot t\big) \\
&\leq \mathcal{P}\big(m^{(0)} - m^{(i)} \geq k - \delta_{\mathbf{v}} + i \cdot t\big) \\
&\leq \mathcal{P}\big(m^{(0)} - m^{(i)} \geq i \cdot t\big) \qquad \text{(as } k \geq \delta_{\mathbf{v}}\text{)}.
\end{aligned}$$

In each step, the martingale value changes by at most $\delta_{\mathbf{v}} + t + 1$. Hence, Azuma's inequality (cf., Section 3.1) asserts

$$\mathcal{P}\big(R_{q(k)\Downarrow 0} = i\big) \leq \exp\left(-\frac{i \cdot t^2}{2(\delta_{\mathbf{v}}+t+1)^2}\right) = a^i.$$

It follows that

$$\mathcal{P}(Run(q(k){\Downarrow}0)) = \sum_{i=k}^{\infty} \mathcal{P}\big(R_{q(k)\Downarrow 0} = i\big) \leq \sum_{i=k}^{\infty} a^i = a^k/(1-a).$$

This proves the first statement. For the second statement, we need to find a condition on k such that $\mathcal{P}(Run(q(k){\Downarrow}0)) \leq 1/2$. The condition $a^k/(1-a) \leq 1/2$ is equivalent to

$$k \geq \frac{\ln(1-a) - \ln 2}{\ln a}.$$

Define $d := \frac{t^2}{2(\delta_{\mathbf{v}}+t+1)^2}$. Note that $a = \exp(-d)$ and $0 < d < 1$. It is straightforward to verify that

$$\frac{\ln(1-\exp(-d)) - \ln 2}{-d} \leq \frac{2}{d^{3/2}} \quad \text{for all } 0 < d < 1.$$

Since

$$\frac{2}{d^{3/2}} = \frac{2 \cdot 2^{3/2} \cdot (\delta_{\mathbf{v}}+t+1)^3}{t^3} \leq \frac{6(\delta_{\mathbf{v}}+t+1)^3}{t^3},$$

the second statement follows. □

LEMMA 5.14. *Assume $[p{\uparrow}] < 1$ for all $p \in \mathscr{B}$. Let $q \in \mathscr{B}$ with $\mathbf{v}_q = \mathbf{v}_{\max}$, where $\mathbf{v}$ is the vector of Theorem 3.4. Then*

$$[q{\uparrow}] \geq \frac{t^3}{12(2\delta_{\mathbf{v}}+4)^3}.$$

PROOF. Define b as the smallest integer such that $b \geq 6(\delta_{\mathbf{v}}+t+1)^3/t^3$. By Lemma 5.12, we have $\mathcal{P}(Run(q(1)\rightarrow b)) \geq 1/(b+1+\delta_{\mathbf{v}})$. Since $0 < t \leq 1$, we have

$$b+1+\delta_{\mathbf{v}} \leq 6(\delta_{\mathbf{v}}+t+2)^3/t^3+1+\delta_{\mathbf{v}} \leq 6(2\delta_{\mathbf{v}}+4)^3/t^3$$

and so

$$\mathcal{P}(Run(q(1)\rightarrow b)) \geq \frac{t^3}{6(2\delta_{\mathbf{v}}+4)^3}.$$

Using Lemma 5.13, we obtain

$$[q\uparrow] \geq \frac{t^3}{12(2\delta_{\mathbf{v}}+4)^3}. \quad \Box$$

Now we can prove Theorem 4.8.

THEOREM 4.8. *Let $\hat{\mathscr{A}}$ be a pOC with Q as the set of control states where the least positive probability used in the rules of $\hat{\mathscr{A}}$ is at least $x_{\min}$. For all $p, q \in Q$, let $[p,q]$ be the probability of all $w \in Run_{\hat{\mathscr{A}}}(p(1))$ such that w starts with a zero-safe finite path from $p(1)$ to $q(k)$, where $k \geq 1$. Let $p \in Q$ such that the probability $[p\uparrow]$ (considered in $\hat{\mathscr{A}}$) is positive. Then there are two possibilities.*

(1) *There is $q \in Q$ such that $[p,q] > 0$ and $[q\uparrow] = 1$. Hence, $[p\uparrow] \geq [p,q] \geq x_{\min}^{|Q|^2}$.*
(2) *There exists a BSCC $\mathscr{B}$ of the underlying finite-state Markov chain $\hat{\mathcal{X}}$ of $\hat{\mathscr{A}}$ and a state q of $\mathscr{B}$ such that the trend t of $\mathscr{B}$ is positive, $[p,q] > 0$, and $\mathbf{v}_q = \mathbf{v}_{\max}$. Here, $\mathbf{v}$ is the vector of Theorem 3.4, and $\mathbf{v}_{\max}$ is the maximal component of $\mathbf{v}$; all of these are considered in $\mathscr{B}$. Further,*

$$[p\uparrow] \;\geq\; \frac{x_{\min}^{4|Q|^2} \cdot t^3}{7000 \cdot |Q|^3}.$$

PROOF. For a given BSCC $\hat{\mathscr{B}}$ of $\hat{\mathcal{X}}$, let $Run(p\uparrow, \hat{\mathscr{B}})$ be the set of all $w \in Run(p\uparrow)$ that visit $\hat{\mathscr{B}}$. Since almost all runs of $Run(p\uparrow)$ visit some BSCC of $\hat{\mathcal{X}}$, there is a BSCC $\mathscr{B}$ of $\hat{\mathcal{X}}$ such that $\mathcal{P}(Run(p\uparrow, \mathscr{B})) > 0$. According to the results of Brázdil et al. [2010b, Section 3], the trend of $\mathscr{B}$ must be nonnegative; and if the trend of $\mathscr{B}$ is equal to zero, there is a configuration $q(k) \in \mathscr{B} \times \mathbb{N}$ such that $[q\uparrow] = 1$ and $p(1)$ can reach $q(k)$ via a zero-safe finite path. Hence, we can distinguish two possibilities.

—The trend of $\mathscr{B}$ is equal to zero and there is a configuration $q(k) \in \mathscr{B} \times \mathbb{N}$ such that $[q\uparrow] = 1$ and $p(1)$ can reach $q(k)$ via a zero-safe finite path. Hence, $\mathcal{P}(Run(p\uparrow, \mathscr{B})) \geq [p,q] > 0$. Now it suffices to show that if $[p,q] > 0$, there is a zero-safe finite path from $p(1)$ to $q(k)$ (for some $k \in \mathbb{N}$) of length at most $|Q|^2$; this implies $[p,q] \geq x_{\min}^{|Q|^2}$. If $[p,q] > 0$, there is some zero-safe finite path w from $p(1)$ to $q(k)$, where $k \in \mathbb{N}$, such that every configuration is visited at most once along w. If the counter value is strictly bounded by $|Q|$ along w, then w visits at most $|Q| \cdot (|Q|-1)$ distinct configurations and we are done. Otherwise, let ℓ by the least number such that the counter value in $w(\ell)$ is $|Q|$. The number of configurations visited along $w(0), \ldots, w(\ell-1)$ is at most $|Q| \cdot (|Q|-1)$, hence $\ell \leq |Q| \cdot (|Q|-1)$. Obviously, $w(\ell)$ can reach $q(k')$ (for some $k' \in \mathbb{N}$) in at most $|Q|$ transitions, and hence there is a zero-safe finite path from $p(1)$ to $q(k')$ of length at most $|Q| \cdot (|Q|-1) + |Q| = |Q|^2$.
—The trend of $\mathscr{B}$ is positive. Let $q \in \mathscr{B}$ be a control state such that $\mathbf{v}_q = \mathbf{v}_{\max}$. Observe that almost all runs of $Run(p\uparrow, \mathscr{B})$ visit a configuration with control state q infinitely

often. Hence, $[p, q] \geq \mathcal{P}(Run(p\uparrow, \mathscr{B})) > 0$. By applying Lemma 5.14, we obtain

$$\mathcal{P}(Run(p\uparrow, \mathscr{B})) \geq \frac{[p, q]t^3}{12(2\delta_{\mathbf{v}} + 4)^3} \geq \frac{x_{\min}^{4|Q|^2} \cdot t^3}{7000 \cdot |Q|^3}.$$

In the last step, we use the bound $\delta_{\mathbf{v}} \leq 2|\mathscr{B}|/x_{\min}^{|\mathscr{B}|}$ given in Theorem 3.4. □

6. CONCLUSIONS

We believe that the methods developed in this article can also be used to approximate other interesting quantities and numerical characteristics of pOC, related to both finite paths and infinite runs. 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.

REFERENCES

E. Allender, P. Bürgisser, J. Kjeldgaard-Pedersen, and P. B. Miltersen. 2008. On the complexity of numerical analysis. *SIAM J. Comput.* 38, 1987–2006.

P. Billingsley. 1995. *Probability and Measure*. Wiley.

T. Brázdil, V. Brožek, J. Holeček, and A. Kučera. 2008. Discounted properties of probabilistic pushdown automata. In *Proceedings of LPAR'08.* Lecture Notes in Computer Science Series, Vol. 5330, Springer, 230–242.

T. Brázdil, V. Brožek, and K. Etessami. 2010a. One-counter stochastic games In *Proceedings of FST&TCS'10*. Leibniz International Proceedings in Informatics, Vol. 8, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 108–119.

T. Brázdil, V. Brožek, K. Etessami, and A. Kučera. 2011a. Approximating the termination value of one-counter MDPs and stochastic games. In *Proceedings of ICALP'11*, Part II. Lecture Notes in Computer Science, vol. 6756, Springer, 332–343.

T. Brázdil, V. Brožek, K. Etessami, A. Kučera, and D. Wojtczak. 2010b. One-counter Markov decision processes. In *Proceedings of SODA'10*. SIAM, 863–874.

T. Brázdil, J. Esparza, S. Kiefer, and A. Kučera. 2013. Analyzing probabilistic pushdown automata. *Formal Methods Syst. Design* 43, 2, 124–163. DOI 10.1007/s10703-012-0166-0.

T. Brázdil, J. Esparza, and A. Kučera. 2005a. Analysis and prediction of the long-run behavior of probabilistic sequential programs with recursion. In *Proceedings of FOCS'05*. IEEE Computer Society Press, 521–530.

T. Brázdil, S. Kiefer, and A. Kučera. 2011b. Efficient analysis of probabilistic programs with an unbounded counter. In *Proceedings of CAV'11.* Lecture Notes in Computer Science, vol. 6806, Springer, 208–224.

T. Brázdil, A. Kučera, P. Novotný, and D. Wojtczak. 2012. Minimizing expected termination time in one-counter Markov decision processes. In *Proceedings of ICALP'12,* Part II. Lecture Notes in Computer Science, vol. 7392, Springer, 141–152.

T. Brázdil, A. Kučera, and O. Stražovský. 2005. On the decidability of temporal properties of probabilistic pushdown automata, *Proceedings of STACS'05*. Lecture Notes in Computer Science, Vol. 3404, Springer. 145–157.

J. Canny. 1988. Some algebraic and geometric computations in PSPACE. In *Proceedings of STOC'88*. ACM Press, 460–467.

K. Chatterjee and L. Doyen. 2010. Energy parity games. In *Proceedings of ICALP'10*, Part II. Lecture Notes in Computer Science, vol. 6199, Springer, 599–610.

K. Chatterjee, L. Doyen, T. Henzinger, and J.-F. Raskin. 2010. Generalized mean-payoff and energy games *Proceedings of FST&TCS'10*. Leibniz International Proceedings in Informatics Series, vol. 8, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 505–516.

J. Esparza, D. Hansel, P. Rossmanith, and S. Schwoon. 2000. Efficient algorithms for model checking pushdown systems. In *Proceedings of CAV'00.* Lecture Notes in Computer Science, vol. 1855, Springer, 232–247.

J. Esparza, A. Kučera, and R. Mayr. 2004. Model-checking probabilistic pushdown automata. In *Proceedings of LICS'04* (13). IEEE Computer Society Press, 12–21.

J. Esparza, A. Kučera, and R. Mayr. 2005. Quantitative analysis of probabilistic pushdown automata: Expectations and variances. In *Proceedings of LICS'05*. IEEE Computer Society Press, 117–126.

K. Etessami, D. Wojtczak, and M. Yannakakis. 2008. Quasi-birth-death processes, tree-like QBDs, probabilistic 1-counter automata, and pushdown systems. In *Proceedings of 5th International Conference on Quantitative Evaluation of Systems (QEST'08)*. IEEE Computer Society Press.

K. Etessami, D. Wojtczak, and M. Yannakakis. 2010. Quasi-birth-death processes, tree-like QBDs, probabilistic 1-counter automata, and pushdown systems. *Performance Eval.* 67, 9, 837–857.

K. Etessami and M. Yannakakis. 2005a. Algorithmic verification of recursive probabilistic systems. In *Proceedings of TACAS'05.* Lecture Notes in Computer Science, vol. 3440, Springer, 253–270.

K. Etessami and M. Yannakakis. 2005b. Checking LTL properties of recursive Markov chains. In *Proceedings of 2nd International Conference on Quantitative Evaluation of Systems (QEST'05)*. IEEE Computer Society Press, 155–165.

K. Etessami and M. Yannakakis. 2005c. Recursive Markov chains, stochastic grammars, and monotone systems of non-linear equations. In *Proceedings of STACS'05*. Lecture Notes in Computer Science Series, vol. 3404, Springery, 340–352.

J. Hopcroft and J. Ullman. 1979. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley.

E. Isaacson and H. B. Keller. 1966. *Analysis of Numerical Methods*. Wiley.

J. Kemeny and J. Snell. 1960. *Finite Markov chains*. D. Van Nostrand Company.

S. Kiefer, M. Luttenberger, and J. Esparza. 2007. On the convergence of Newton's method for monotone systems of polynomial equations. In *Proceedings of STOC'07*. ACM, 217–226.

J. Křetínský and R. Ledesma-Garza. 2013. Rabinizer 2: Small deterministic automata for LTL$\setminus$GU. In *Proceedings of ATVA'13.* Lecture Notes in Computer Science, vol. 8172. Springer, 446–450.

M. Neuts. 1981. *Matrix-geometric Solutions in Stochastic Models: An Algorithmic Approach*. Courier Dover Publications.

J . Rosenthal. 2006. *A First Look at Rigorous Probability Theory*. World Scientific Publishing.

A. Stewart, K. Etessami, and M. Yannakakis. 2013. Upper bounds for Newton's method on monotone polynomial systems, and P-time model checking of probabilistic one-counter automata. In *Proceedings of CAV'13.* Lecture Notes in Computer Science, vol. 8044, Springer, 495–510.

W. Thomas. 1991. Automata on infinite objects. *Handbook of Theoretical Computer Science B*, 135–192.

D. Williams. 1991. *Probability with Martingales*. Cambridge University Press.

Received April 2012; revised September 2013; accepted April 2014