

Zero-Reachability in Probabilistic Multi-Counter Automata*

Tomáš Brázdil

Faculty of Informatics, Masaryk
University, Brno, Czech Republic
brazdil@fi.muni.cz

Stefan Kiefer

Department of Computer Science,
University of Oxford, United Kingdom
stefan.kiefer@cs.ox.ac.uk

Antonín Kučera

Faculty of Informatics, Masaryk
University, Brno, Czech Republic
kucera@fi.muni.cz

Petr Novotný

Faculty of Informatics, Masaryk University,
Brno, Czech Republic
xnovot18@fi.muni.cz

Joost-Pieter Katoen

Department of Computer Science,
RWTH Aachen University, Germany
katoen@cs.rwth-aachen.de

Abstract

We study the qualitative and quantitative zero-reachability problem in probabilistic multi-counter systems. We identify the undecidable variants of the problems, and then we concentrate on the remaining two cases. In the first case, when we are interested in the probability of all runs that visit zero in *some* counter, we show that the qualitative zero-reachability is decidable in time which is polynomial in the size of a given pMC and doubly exponential in the number of counters. Further, we show that the probability of all zero-reaching runs can be effectively approximated up to an arbitrarily small given error $\varepsilon > 0$ in time which is polynomial in $\log(\varepsilon)$, exponential in the size of a given pMC, and doubly exponential in the number of counters. In the second case, we are interested in the probability of all runs that visit zero in some counter different from the last counter. Here we show that the qualitative zero-reachability is decidable and SQUAREROOTSUM-hard, and the probability of all zero-reaching runs can be effectively approximated up to an arbitrarily small given error $\varepsilon > 0$ (these results apply to pMC satisfying a suitable technical condition that can be verified in polynomial time). The proof techniques invented in the second case allow to construct counterexamples for some classical results about ergodicity in stochastic Petri nets.

Categories and Subject Descriptors G.3 [Probability and Statistics]: Markov processes

General Terms Theory

Keywords multi-counter automata, Markov chains, reachability

*Tomáš Brázdil is supported by the Czech Science Foundation, grant No P202/12/P612. S. Kiefer is supported by a Royal Society University Research Fellowship. J.-P. Katoen is supported by the Excellence Initiative of the federal German government.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CSL-LICS 2014, July 14–18, 2014, Vienna, Austria.
Copyright © 2014 ACM 978-1-4503-2886-9...\$15.00.
<http://dx.doi.org/10.1145/2603088.2603161>

1. Introduction

A *probabilistic multi-counter automaton (pMC)* \mathcal{A} of dimension $d \in \mathbb{N}$ is an abstract fully probabilistic computational device equipped with a finite-state control unit and d unbounded counters that can store non-negative integers. A *configuration* \mathbf{pv} of \mathcal{A} is given by the current control state p and the vector of current counter values \mathbf{v} . The dynamics of \mathcal{A} is defined by a finite set of *rules* of the form (p, α, c, q) where p is the current control state, q is the next control state, α is a d -dimensional vector of counter changes ranging over $\{-1, 0, 1\}^d$, and c is a subset of counters that are tested for zero. Moreover, each rule is assigned a positive integer *weight*. A rule (p, α, c, q) is *enabled* in a configuration \mathbf{pv} if the set of all counters with zero value in \mathbf{v} is precisely c and no component of $\mathbf{v} + \alpha$ is negative; such an enabled rule can be *fired* in \mathbf{pv} and generates a *probabilistic transition* $\mathbf{pv} \xrightarrow{x} q(\mathbf{v} + \alpha)$ where the probability x is equal to the weight of the rule divided by the total weight of all rules enabled in \mathbf{pv} . A special subclass of pMC are *probabilistic vector addition systems with states (pVASS)*, which are equivalent to (discrete-time) *stochastic Petri nets (SPN)*. Intuitively, a pVASS is a pMC where no subset of counters is tested for zero explicitly (see Section 2 for a precise definition).

The decidability and complexity of basic qualitative/quantitative problems for pMCs has so far been studied mainly in the one-dimensional case, and there are also some results about unbounded SPN (a more detailed overview of the existing results is given below). In this paper, we consider *multi-dimensional* pMC and the associated *zero-reachability* problem. That is, we are interested in the probability of all runs initiated in a given \mathbf{pv} that eventually visit a “zero configuration”. Since there are several counters, the notion of “zero configuration” can be formalized in various ways (for example, we might want to have zero in some counter, in all counters simultaneously, or in a given subset of counters). Therefore, we consider a general *stopping criterion* \mathcal{Z} which consists of *minimal* subsets of counters that are required to be simultaneously zero. For example, if $\mathcal{Z} = \mathcal{Z}_{all} = \{\{1\}, \dots, \{d\}\}$, then a run is stopped when reaching a configuration with zero in *some* counter; and if we put $\mathcal{Z} = \{\{1, 2\}\}$, then a run is stopped when reaching a configuration with zero in counters 1 and 2 (and possibly also in other counters). We use $\mathcal{P}(\text{Run}(\mathbf{pv}, \mathcal{Z}))$ to denote the probability of all runs initiated in \mathbf{pv} that reach a configuration satisfying the stopping criterion \mathcal{Z} . The main algorithmic problems considered in this paper are the following:

- *Qualitative \mathcal{Z} -reachability*: Is $\mathcal{P}(\text{Run}(\mathbf{pv}, \mathcal{Z})) = 1$?




Figure 1. Firing process may not be ergodic.

- *Approximation:* Can $\mathcal{P}(\text{Run}(p\mathbf{v}, \mathcal{Z}))$ be approximated up to a given absolute/relative error $\varepsilon > 0$?

We start by observing that the above problems are not effectively solvable in general, and we show that there are only two potentially decidable cases, where \mathcal{Z} is equal either to \mathcal{Z}_{all} (Case I) or to $\mathcal{Z}_{-i} = \mathcal{Z}_{\text{all}} \setminus \{\{i\}\}$ (Case II). Recall that if $\mathcal{Z} = \mathcal{Z}_{\text{all}}$, then a run is stopped when some counter reaches zero; and if $\mathcal{Z} = \mathcal{Z}_{-i}$, then a run is stopped when a counter different from i reaches zero. Cases I and II are analyzed independently and the following results are achieved:

Case I: We show that the qualitative \mathcal{Z}_{all} -reachability problem is decidable in time polynomial in the encoding size of \mathcal{A} (which we denote by $|\mathcal{A}|$) and doubly exponential in d . In particular, this means that the problem is decidable in *polynomial time for every fixed d* . Then, we show that $\mathcal{P}(\text{Run}(p\mathbf{v}, \mathcal{Z}_{\text{all}}))$ can be effectively approximated up to a given absolute/relative error $\varepsilon > 0$ in time which is polynomial in $\log(\varepsilon)$, exponential in $|\mathcal{A}|$, and doubly exponential in d (in the special case when $d = 1$, the problem is known to be solvable in time polynomial in $|\mathcal{A}|$ and $\log(\varepsilon)$, see [18]).

Case II: We analyze Case II only under a technical assumption that counter i is not critical; roughly speaking, this means that counter i has either a tendency to increase or a tendency to decrease when the other counters are positive. The problem whether counter i is critical or not is solvable in time polynomial in $|\mathcal{A}|$, so we can efficiently check whether a given pMC can be analyzed by our methods.

Under the mentioned assumption, we show how to construct a suitable martingale which captures the behaviour of certain runs in \mathcal{A} . Thus, we obtain a new and versatile tool for analyzing quantitative properties of runs in multi-dimensional pMC, which is more powerful than the martingale of [13] constructed for one-dimensional pMC. Using this martingale and the results of [7], we show that the qualitative \mathcal{Z}_{-i} -reachability problem is decidable. We also show that the problem is SQUARE-ROOM-SUM-hard, even for two-dimensional pMC satisfying the mentioned technical assumption. Further, we show that $\mathcal{P}(\text{Run}(p\mathbf{v}, \mathcal{Z}_{-i}))$ can be effectively approximated up to a given absolute error $\varepsilon > 0$. The main reason why we do not provide any upper complexity bounds in Case II is a missing upper bound for coverability in vector addition systems with one zero test (see [7]).

It is worth noting that the techniques developed in Case II reveal the existence of phenomena that should not exist according to the previous results about ergodicity in SPN. A classical paper in this area [22] has been written by Florin & Natkin in 80s. In the paper, it is claimed that if the state-space of a given SPN (with arbitrarily many unbounded places) is strongly connected, then the firing process is ergodic (see Section IV.B. in [22]). In the setting of discrete-time probabilistic Petri nets, this means that for almost all runs, the limit frequency of transitions performed along a run is defined and takes the same value. However, in Fig. 1 there is an example of a pVASS (depicted as SPN with weighted transitions) with two counters (places) and strongly connected state space

where the limit frequency of transitions may take two eligible values (each with probability 1/2). Intuitively, if both counters are positive, then both of them have a tendency to decrease (i.e., the trend of the only BSCC of $\mathcal{F}_{\mathcal{A}}$ is negative in both components, see Section 3.1). However, if we reach a configuration where the first counter is zero and the second counter is sufficiently large, then the second counter starts to *increase*, i.e., it never becomes zero again with some positive probability (cf. the *oc-trend* of the only BSCC D of \mathcal{B}_1 introduced in Section 3.2). The first counter stays zero for most of the time, because when it becomes positive, it is immediately emptied with a very large probability. This means that the frequency of firing t_2 will be much higher than the frequency of firing t_1 . When we reach a configuration where the first counter is large and the second counter is zero, the situation is symmetric, i.e., the frequency of firing t_1 becomes much higher than the frequency of firing t_2 . Further, almost every run eventually behaves according to one of the two scenarios, and therefore there are two eligible limit frequencies of transitions, each of which is taken with probability 1/2. So, we must unfortunately conclude that the results of [22] are not valid for general SPN.

Related Work. One-dimensional pMC and their extensions into decision processes and games were studied in [9–11, 13, 18–20]. In particular, in [18] it was shown that the termination probability (a “selective” variant of zero-reachability) in one-dimensional pMC can be approximated up to an arbitrarily small given error in polynomial time. In [13], it was shown how to construct a martingale for a given one-dimensional pMC which allows to derive tail bounds on termination time (we use this martingale in Section 3.1).

There is also a vast amount of literature about SPN (see, e.g., [5, 27]), and some of these works also consider algorithmic aspects of unbounded SPN (see, e.g., [1, 21, 22]).

A considerable amount of papers has been devoted to algorithmic analysis of so-called probabilistic lossy channel systems (PLCS) and their game extensions (see e.g. [2–4, 6, 23]). PLCS are a stochastic extension of lossy channel systems, i.e., an infinite-state model comprising several interconnected unbounded queues coupled with a finite-state control unit. The main ingredient, which makes results about PLCS incomparable with our results on pMCs, is that queues may lose messages with a fixed loss rate, which substantially simplifies the associated analysis.

2. Preliminaries

We use \mathbb{Z} , \mathbb{N} , \mathbb{N}^+ , \mathbb{Q} , and \mathbb{R} to denote the set of all integers, non-negative integers, positive integers, rational numbers, and real numbers, respectively. For a vector \mathbf{v} we denote by $\mathbf{v}[i]$ the i -component of \mathbf{v} .

A *labelled transition system* is a tuple $\mathcal{V} = (V, L, \rightarrow)$, where V is a non-empty set of vertices, L a non-empty set of labels, and $\rightarrow \subseteq V \times L \times V$ a *total relation* (i.e., for every $v \in V$ there is at least one *outgoing* transition $(v, \ell, u) \in \rightarrow$). As usual, we write $v \xrightarrow{\ell} u$ instead of $(v, \ell, u) \in \rightarrow$, and $v \rightarrow u$ iff $v \xrightarrow{\ell} u$ for some $\ell \in L$. The reflexive and transitive closure of \rightarrow is denoted by \rightarrow^* . A *finite path* in \mathcal{V} of length $k \geq 0$ is a finite sequence of the form $v_0 \ell_0 v_1 \ell_1 \dots \ell_{k-1} v_k$, where $v_i \xrightarrow{\ell_i} v_{i+1}$ for all $0 \leq i < k$. The length of a finite path w is denoted by $\text{length}(w)$. A *run* in \mathcal{V} is an infinite alternating sequence w of vertices and labels $v_0 \ell_0 v_1 \ell_1 \dots$ such that every finite prefix of w ending in a vertex is a finite path in \mathcal{V} . For $w = v_0 \ell_0 v_1 \ell_1 \dots$ we denote by $w(i)$ the vertex v_i . The sets of all finite paths and all runs in \mathcal{V} are denoted by $FPath_{\mathcal{V}}$ and $Run_{\mathcal{V}}$, respectively. The sets of all finite paths and all runs in \mathcal{V} that start with a given finite path w are denoted by $FPath_{\mathcal{V}}(w)$ and $Run_{\mathcal{V}}(w)$, respectively. A *strongly connected component (SCC)* of \mathcal{V} is a maximal subset $C \subseteq V$ such that for all $v, u \in C$ we have

that $v \rightarrow^* u$. A SCC C of \mathcal{V} is a *bottom SCC (BSCC)* of \mathcal{V} if for all $v \in C$ and $u \in V$ such that $v \rightarrow u$ we have that $u \in C$.

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

Definition 1. A labeled Markov chain is a tuple $\mathcal{M} = (S, L, \rightarrow, \text{Prob})$ where $S \neq \emptyset$ is a finite or countably infinite set of states, $L \neq \emptyset$ is a finite or countably infinite set of labels, $\rightarrow \subseteq S \times L \times S$ 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 . We write $s \xrightarrow{\ell, x} t$ when $s \rightarrow t$ and x is the probability of (s, ℓ, t) .

All notions defined for labelled transition systems naturally extend to Markov chains. If $L = \{\ell\}$ is a singleton, we say that \mathcal{M} is *non-labeled*, and we omit both L and ℓ when specifying \mathcal{M} (in particular, we write $s \xrightarrow{x} t$ instead of $s \xrightarrow{\ell, x} t$). To every $s \in S$ we associate the standard probability space $(\text{Run}_{\mathcal{M}}(s), \mathcal{F}, \mathcal{P})$ of runs starting at s , where \mathcal{F} is the σ -field generated by all *basic cylinders* $\text{Run}_{\mathcal{M}}(w)$, where w is a finite path starting at s , and $\mathcal{P} : \mathcal{F} \rightarrow [0, 1]$ is the unique probability measure such that $\mathcal{P}(\text{Run}_{\mathcal{M}}(w)) = \prod_{i=1}^{\text{length}(w)} x_i$ where x_i is the probability of $w(i-1) \xrightarrow{\ell_i} w(i)$ for every $1 \leq i \leq \text{length}(w)$. If $\text{length}(w) = 0$, we put $\mathcal{P}(\text{Run}_{\mathcal{M}}(w)) = 1$.

Now we introduce probabilistic multi-counter automata (pMC). For technical convenience, we consider *labeled* rules, where the associated finite set of labels always contains a distinguished element τ . The role of the labels becomes clear in Section 3.2 where we abstract a (labeled) one-dimensional pMC from a given multi-dimensional one.

Definition 2. Let L be a finite set of labels such that $\tau \in L$, and let $d \in \mathbb{N}^+$. An L -labeled d -dimensional probabilistic multi-counter automaton (pMC) is a triple $\mathcal{A} = (Q, \gamma, W)$, where

- Q is a finite set of states,
- $\gamma \subseteq Q \times \{-1, 0, 1\}^d \times 2^{\{1, \dots, d\}} \times L \times Q$ is a set of rules such that for all $p \in Q$ and $c \subseteq \{1, \dots, d\}$ there is at least one outgoing rule of the form (p, α, c, ℓ, q) ,
- $W : \gamma \rightarrow \mathbb{N}^+$ is a weight assignment.

The encoding size of \mathcal{A} is denoted by $|\mathcal{A}|$, where the weights used in W and the counter indexes used in γ are encoded in binary.

A configuration of $\mathcal{A} = (Q, \gamma, W)$ is an element of $Q \times \mathbb{N}^d$, written as $p\mathbf{v}$. We use $Z(p\mathbf{v}) = \{i \mid 1 \leq i \leq d, \mathbf{v}[i] = 0\}$ to denote the set of all counters that are zero in $p\mathbf{v}$. A rule $(p, \alpha, c, \ell, q) \in \gamma$ is *enabled* in a configuration $p\mathbf{v}$ if $Z(p\mathbf{v}) = c$ and for all $1 \leq i \leq d$ where $\alpha[i] = -1$ we have that $\mathbf{v}[i] > 0$.

The semantics of a pMC \mathcal{A} is given by the associated L -labeled Markov chain $\mathcal{M}_{\mathcal{A}}$ whose states are the configurations of \mathcal{A} , and the outgoing transitions of a configuration $p\mathbf{v}$ are determined as follows:

- If no rule of γ is enabled in $p\mathbf{v}$, then $p\mathbf{v} \xrightarrow{1, \tau} p\mathbf{v}$ is the only outgoing transition of $p\mathbf{v}$;
- otherwise, for every rule $(p, \alpha, c, \ell, q) \in \gamma$ enabled in $p\mathbf{v}$ there is a transition $p\mathbf{v} \xrightarrow{x, \ell} q\mathbf{u}$ such that $\mathbf{u} = \mathbf{v} + \alpha$ and $x = W((p, \alpha, c, \ell, q))/T$, where T is the total weight of all rules enabled in $p\mathbf{v}$.

When $L = \{\tau\}$, we say that \mathcal{A} is *non-labeled*, and both L and τ are omitted when specifying \mathcal{A} . We say that \mathcal{A} is a *probabilistic vector addition system with states (pVASS)* if no subset of counters is tested for zero, i.e., for every $(p, \alpha, \ell, q) \in Q \times \{-1, 0, 1\}^d \times L \times Q$

we have that γ contains either all rules of the form (p, α, c, ℓ, q) (for all $c \subseteq \{1, \dots, d\}$) with the same weight, or no such rule. For every configuration $p\mathbf{v}$, we use $\text{state}(p\mathbf{v})$ and $\text{cval}(p\mathbf{v})$ to denote the control state p and the vector of counter values \mathbf{v} , respectively. We also use $\text{cval}_i(p\mathbf{v})$ to denote $\mathbf{v}[i]$.

Qualitative zero-reachability. A *stopping criterion* is a non-empty set $\mathcal{Z} \subseteq 2^{\{1, \dots, d\}}$ of pairwise incomparable non-empty subsets of counters. For every configuration $p\mathbf{v}$, let $\text{Run}(p\mathbf{v}, \mathcal{Z})$ be the set of all $w \in \text{Run}(p\mathbf{v})$ such that there exist $k \in \mathbb{N}$ and $\varrho \in \mathcal{Z}$ satisfying $\varrho \subseteq Z(w(k))$. Intuitively, \mathcal{Z} specifies the minimal subsets of counters that must be *simultaneously* zero to stop a run. The *qualitative \mathcal{Z} -reachability problem* is formulated as follows:

Instance: A d -dimensional pMC \mathcal{A} and a control state p of \mathcal{A} .
Question: Do we have $\mathcal{P}(\text{Run}(p\mathbf{1}, \mathcal{Z})) = 1$?

Here $\mathbf{1} = (1, \dots, 1)$ is a d -dimensional vector of 1's. We also use $\text{Run}(p\mathbf{v}, \neg\mathcal{Z})$ to denote $\text{Run}(p\mathbf{v}) \setminus \text{Run}(p\mathbf{v}, \mathcal{Z})$, and we say that $w \in \text{FPath}(p\mathbf{v})$ is *\mathcal{Z} -safe* if for all $w(i)$ where $0 \leq i < \text{length}(w)$ and all $\varrho \in \mathcal{Z}$ we have that $\varrho \not\subseteq Z(w(i))$.

3. The Results

We start by observing that the qualitative zero-reachability problem is undecidable in general, and we identify potentially decidable subclasses.

Observation 1. Let $\mathcal{Z} \subseteq 2^{\{1, \dots, d\}}$ be a stopping criterion satisfying one of the following conditions:

- there is $\varrho \in \mathcal{Z}$ with more than one element;
- there are $i, j \in \{1, \dots, d\}$ such that $i \neq j$ and for every $\varrho \in \mathcal{Z}$ we have that $\{i, j\} \cap \varrho = \emptyset$.

Then, the qualitative \mathcal{Z} -reachability problem is undecidable, even if the set of instances is restricted to pairs (\mathcal{A}, p) such that $\mathcal{P}(\text{Run}(p\mathbf{1}, \mathcal{Z}))$ is either 0 or 1 (hence, $\mathcal{P}(\text{Run}(p\mathbf{1}, \mathcal{Z}))$ cannot be effectively approximated up to an absolute error smaller than 0.5).

A proof of Observation 1 is immediate. For a given Minsky machine M (see [26]) with two counters initialized to one, we construct pMCs \mathcal{A}_a and \mathcal{A}_b of dimension 2 and 3, respectively, and a control state p such that

- if M halts, then we have $\mathcal{P}(\text{Run}_{\mathcal{M}_{\mathcal{A}_a}}(p\mathbf{1}, \{\{1, 2\}\})) = 1$ and $\mathcal{P}(\text{Run}_{\mathcal{M}_{\mathcal{A}_b}}(p\mathbf{1}, \{\{3\}\})) = 1$;
- if M does not halt, then $\mathcal{P}(\text{Run}_{\mathcal{M}_{\mathcal{A}_a}}(p\mathbf{1}, \{\{1, 2\}\})) = 0$ and $\mathcal{P}(\text{Run}_{\mathcal{M}_{\mathcal{A}_b}}(p\mathbf{1}, \{\{3\}\})) = 0$.

The construction of \mathcal{A}_a and \mathcal{A}_b is trivial (and hence omitted). Note that \mathcal{A}_b can faithfully simulate the instructions of M using the counters 1 and 2. The third counter is decreased to zero only when a control state corresponding to the halting instruction of M is reached. Similarly, \mathcal{A}_a simulates the instructions of M using its two counters, but here we need to ensure that a configuration where *both* counters are simultaneously zero is entered iff a control state corresponding to the halting instruction of M is reached. This is achieved by increasing both counters by 1 initially, and then decreasing/increasing counter i before/after simulating a given instruction of M operating on counter i .

Note that the construction of \mathcal{A}_a and \mathcal{A}_b can trivially be adapted to pMCs of higher dimensions satisfying the conditions (a) and (b) of Observation 1, respectively. However, there are two cases not covered by Observation 1:

- $\mathcal{Z}_{\text{all}} = \{\{1\}, \dots, \{d\}\}$, i.e., a run is stopped when *some* counter reaches zero.

II. $\mathcal{Z}_{-i} = \{\{1\}, \dots, \{d\}\} \setminus \{\{i\}\}$ where $i \in \{1, \dots, d\}$, i.e., a run is stopped when a counter different from i reaches zero. The counters different from i are called *stopping counters*.

These cases are analyzed in the following subsections. The proofs omitted due to space constraints can be found in the full version of this paper [16].

3.1 Zero-Reachability, Case I

For the rest of this section, let us fix a (non-labeled) pMC $\mathcal{A} = (Q, \gamma, W)$ of dimension $d \in \mathbb{N}^+$ and a configuration $p\mathbf{v}$.

Our aim is to identify the conditions under which it holds $\mathcal{P}(\text{Run}(p\mathbf{v}, \neg\mathcal{Z}_{\text{all}})) > 0$. To achieve that, we first consider a (non-labeled) finite-state Markov chain $\mathcal{F}_{\mathcal{A}} = (Q, \hookrightarrow, \text{Prob})$ where $q \xrightarrow{x} r$ iff

$$x = \sum_{(q, \boldsymbol{\alpha}, \emptyset, r) \in \gamma} P_{\emptyset}(q, \boldsymbol{\alpha}, \emptyset, r) > 0.$$

Here $P_{\emptyset} : \gamma \rightarrow [0, 1]$ is the probability assignment for the rules defined as follows (we write $P_{\emptyset}(q, \boldsymbol{\alpha}, \emptyset, r)$ instead of $P_{\emptyset}((q, \boldsymbol{\alpha}, \emptyset, r))$):

- For every rule $(p, \boldsymbol{\alpha}, c, q)$ with $c \neq \emptyset$ we put $P_{\emptyset}(p, \boldsymbol{\alpha}, c, q) = 0$.
- $P_{\emptyset}(p, \boldsymbol{\alpha}, \emptyset, q) = W((p, \boldsymbol{\alpha}, \emptyset, q))/T$, where T is the total weight of all rules of the form $(p, \boldsymbol{\alpha}', \emptyset, q')$.

Intuitively, a state q of $\mathcal{F}_{\mathcal{A}}$ captures the behavior of configurations $q\mathbf{u}$ where all components of \mathbf{u} are positive.

Further, we partition the states of Q into SCCs C_1, \dots, C_m according to \hookrightarrow . Note that every run $w \in \text{Run}(p\mathbf{v})$ eventually stays in precisely one C_j , i.e., there is precisely one $1 \leq j \leq m$ such that for some $k \in \mathbb{N}$, the control state of every $w(k')$, where $k' \geq k$, belongs to C_j . We use $\text{Run}(p\mathbf{v}, C_j)$ to denote the set of all $w \in \text{Run}(p\mathbf{v}, \neg\mathcal{Z}_{\text{all}})$ that stay in C_j . Obviously,

$$\text{Run}(p\mathbf{v}, \neg\mathcal{Z}_{\text{all}}) = \text{Run}(p\mathbf{v}, C_1) \uplus \dots \uplus \text{Run}(p\mathbf{v}, C_m).$$

For any $n \in \mathbb{N}$ denote by P_n the probability that a run w initiated in $p\mathbf{v}$ satisfies the following for every $0 \leq i \leq n$: $\text{state}(w(i))$ does not belong to any BSCC of $\mathcal{F}_{\mathcal{A}}$ and $Z(w(i)) = \emptyset$. The following lemma shows that P_n decays exponentially fast.

Lemma 1. *For any $n \in \mathbb{N}$ we have*

$$P_n \leq (1 - p_{\min}^{|Q|})^{\lfloor \frac{n}{|\mathcal{Q}|} \rfloor},$$

where p_{\min} is the minimal positive transition probability in $\mathcal{M}_{\mathcal{A}}$. In particular, for any non-bottom SCC C of $\mathcal{F}_{\mathcal{A}}$ it holds that $\mathcal{P}(\text{Run}(p\mathbf{v}, C)) = 0$.

Proof. The lemma immediately follows from the fact that for every configuration $p\mathbf{v}$ there is a path (in \mathcal{A}) of length at most $|Q|$ to a configuration $q\mathbf{u}$ satisfying either $Z(q\mathbf{u}) \neq \emptyset$ or $q \in D$ for some BSCC D of $\mathcal{F}_{\mathcal{A}}$. \square

Now, let C be a BSCC of $\mathcal{F}_{\mathcal{A}}$. For every $q \in C$, let *change* ^{q} be a d -dimensional vector of *expected counter changes* given by

$$\text{change}^q[i] = \sum_{(q, \boldsymbol{\alpha}, \emptyset, r) \in \gamma} P_{\emptyset}(q, \boldsymbol{\alpha}, \emptyset, r) \cdot \boldsymbol{\alpha}[i].$$

Note that C can be seen as a finite-state irreducible Markov chain, and hence there exists the unique *invariant distribution* μ on the states of C (see, e.g., [24]) satisfying

$$\mu(q) = \sum_{r \xrightarrow{x} q} \mu(r) \cdot x.$$

The *trend* of C is a d -dimensional vector \mathbf{t} defined by

$$\mathbf{t}[i] = \sum_{q \in C} \mu(q) \cdot \text{change}^q[i].$$

Further, for every $i \in \{1, \dots, d\}$ and every $q \in C$, we denote by $\text{botfin}_i(q)$ the *least* $j \in \mathbb{N}$ such that for every configuration $q\mathbf{u}$ where $\mathbf{u}[i] = j$, there is *no* $w \in \text{FPath}_{\mathcal{M}_{\mathcal{A}}}(q\mathbf{u})$ where counter i is zero in the last configuration of w and all counters stay positive in every $w(k)$, where $0 \leq k < \text{length}(w)$. If there is no such j , we put $\text{botfin}_i(q) = \infty$. It is easy to show that if $\text{botfin}_i(q) < \infty$, then $\text{botfin}_i(q) \leq |C|$; and if $\text{botfin}_i(q) = \infty$, then $\text{botfin}_i(r) = \infty$ for all $r \in C$. Moreover, if $\text{botfin}_i(q) < \infty$, then there is a \mathcal{Z}_{-i} -safe finite path of length at most $|C| - 1$ from $q\mathbf{u}$ to a configuration with counter i equal to 0, where $\mathbf{u}[i] = \text{botfin}_i(q) - 1$ and $\mathbf{u}[\ell] = |C|$ for $\ell \neq i$. In particular, the number $\text{botfin}_i(q)$ is computable in time polynomial in $|C|$.

We say that counter i is *decreasing* in C if $\text{botfin}_i(q) = \infty$ for some (and hence all) $q \in C$.

Definition 3. *Let C be a BSCC of $\mathcal{F}_{\mathcal{A}}$ with trend \mathbf{t} , and let $i \in \{1, \dots, d\}$. We say that counter i is *diverging* in C if either $\mathbf{t}[i] > 0$, or $\mathbf{t}[i] = 0$ and the counter i is not decreasing in C .*

Intuitively, our aim is to prove that $\mathcal{P}(\text{Run}(p\mathbf{v}, C)) > 0$ iff all counters are diverging in C and $p\mathbf{v}$ can reach a configuration $q\mathbf{u}$ (via a \mathcal{Z}_{all} -safe finite path) where all components of \mathbf{u} are “sufficiently large”. To analyze the individual counters, for every $i \in \{1, \dots, d\}$ we introduce a (labeled) *one-dimensional* pMC which faithfully simulates the behavior of counter i and “updates” the other counters just symbolically in the labels.

Definition 4. *Let $L = \{-1, 0, 1\}^{d-1}$, and let $\mathcal{B}_i = (Q, \hat{\gamma}, \hat{W})$ be an L -labeled pMC of dimension one such that*

- $(q, j, \emptyset, \boldsymbol{\beta}, r) \in \hat{\gamma}$ iff $(q, \langle \boldsymbol{\beta}, j \rangle_i, \emptyset, r) \in \gamma$;
- $(q, j, \{1\}, \boldsymbol{\beta}, r) \in \hat{\gamma}$ iff $(q, \langle \boldsymbol{\beta}, j \rangle_i, \{i\}, r) \in \gamma$;
- $\hat{W}(q, j, \emptyset, \boldsymbol{\beta}, r) = W(q, \langle \boldsymbol{\beta}, j \rangle_i, \emptyset, r)$.
- $\hat{W}(q, j, \{1\}, \boldsymbol{\beta}, r) = W(q, \langle \boldsymbol{\beta}, j \rangle_i, \{i\}, r)$.

Here, $\langle (j_1, \dots, j_{d-1}), j \rangle_i = (j_1, \dots, j_{i-1}, j, j_i, \dots, j_{d-1})$.

Observe that the symbolic updates of the counters different from i “performed” in the labels of \mathcal{B}_i mimic the real updates performed by \mathcal{A} in configurations where all of these counters are positive.

Given a run $w \equiv p_0(v_0) \boldsymbol{\alpha}_0 p_1(v_1) \boldsymbol{\alpha}_1 p_2(v_2) \boldsymbol{\alpha}_2 \dots$ from $\text{Run}_{\mathcal{M}_{\mathcal{B}_i}}(p_0(v_0))$ and $k \in \mathbb{N}$, we denote by $\text{tot}(w; k)$ the vector $\sum_{n=0}^{k-1} \boldsymbol{\alpha}_n$, and given $j \in \{1, \dots, d\} \setminus \{i\}$, we denote by $\text{tot}_j(w; k)$ the number $\sum_{n=0}^{k-1} \boldsymbol{\alpha}_n[j]$ (i.e., the j -th component of $\text{tot}(w; k)$).

Let Υ_i be a function which for a given run $w \equiv p_0 \mathbf{v}_0 p_1 \mathbf{v}_1 \dots$ of $\text{Run}_{\mathcal{M}_{\mathcal{A}}}(p\mathbf{v}, \neg\mathcal{Z}_{-i})$ returns a run

$$\Upsilon_i(w) \equiv p_0(\mathbf{v}_0[i]) \boldsymbol{\alpha}_0 p_1(\mathbf{v}_1[i]) \boldsymbol{\alpha}_1 p_2(\mathbf{v}_2[i]) \boldsymbol{\alpha}_2 \dots$$

of $\text{Run}_{\mathcal{M}_{\mathcal{B}_i}}(p(\mathbf{v}[i]))$ where the label $\boldsymbol{\alpha}_j$ corresponds to the update in the abstracted counters performed in the transition $p_j \mathbf{v}_j \rightarrow p_{j+1} \mathbf{v}_{j+1}$, i.e., $\mathbf{v}_{j+1} - \mathbf{v}_j = \langle \boldsymbol{\alpha}_j, \mathbf{v}_{j+1}[i] - \mathbf{v}_j[i] \rangle_i$. The next lemma is immediate.

Lemma 2. *For all $w \in \text{Run}_{\mathcal{M}_{\mathcal{A}}}(p\mathbf{v}, \neg\mathcal{Z}_{-i})$ and $k \in \mathbb{N}$ we have*

- $\text{state}(w(k)) = \text{state}(\Upsilon_i(w)(k))$,
- $\text{eval}(w(k)) = \langle \text{tot}(\Upsilon_i(w); k), \text{eval}_1(\Upsilon_i(w)(k)) \rangle_i$.

Further, for every measurable set $R \subseteq \text{Run}_{\mathcal{M}_{\mathcal{A}}}(p\mathbf{v}, \neg\mathcal{Z}_{-i})$ we have that $\Upsilon_i(R)$ is measurable and

$$\mathcal{P}(R) = \mathcal{P}(\Upsilon_i(R)) \quad (1)$$

Now we examine the runs of $\text{Run}(p\mathbf{v}, C)$ where C is a BSCC of $\mathcal{F}_{\mathcal{A}}$ such that some counter is not diverging in C .

Lemma 3. *Let C be a BSCC of $\mathcal{F}_{\mathcal{A}}$. If some counter is not diverging in C , then $\mathcal{P}(\text{Run}(p\mathbf{v}, C)) = 0$.*

It remains to consider the case when C is a BSCC of \mathcal{F}_A where all counters are diverging. Here we use the results of [13] which allow to derive a bound on divergence probability in one-dimensional pMC. These results are based on designing and analyzing a suitable martingale for one-dimensional pMC.

Lemma 4. *Let \mathcal{B} be a 1-dimensional pMC, let C be a BSCC of \mathcal{F}_B such that the trend t of the only counter in C is positive and let $\delta = 2|C|/x_{\min}^{|C|}$ where x_{\min} is the smallest non-zero transition probability in \mathcal{M}_B . Then for all $q \in C$ and $k > 2\delta/t$ we have that $\mathcal{P}(q(k), \neg \mathcal{Z}) \geq 1 - (a^k/(1-a))$, where $\mathcal{Z} = \{1\}$ and $a = \exp(-t^2/8(\delta+t+1)^2)$.*

Proof. Denote by $[q(k)\downarrow, \ell]$ the probability that a run initiated in $q(k)$ visits a configuration with zero counter value for the first time in exactly ℓ steps. By Proposition 7 of [13] we obtain for all $\ell \geq h = 2\delta/t^1$,

$$[q(k)\downarrow, \ell] \leq a^\ell$$

where $a = \exp(-t^2/8(\delta+t+1)^2)$ for $\delta \leq 2|C|/x_{\min}^{|C|}$.

Thus

$$\mathcal{P}(q(k), \neg \mathcal{Z}) \geq 1 - \sum_{\ell=k}^{\infty} [q(k)\downarrow, \ell] \geq 1 - \frac{a^k}{1-a}$$

□

Definition 5. *Let C be a BSCC of \mathcal{F}_A where all counters are diverging, and let $q \in C$. We say that a configuration \mathbf{qu} is above a given $n \in \mathbb{N}$ if $\mathbf{u}[i] \geq n$ for every i such that $\mathbf{t}[i] > 0$, and $\mathbf{u}[i] \geq \text{botfin}_i(q)$ for every i such that $\mathbf{t}[i] = 0$.*

Lemma 5. *Let C be a BSCC of \mathcal{F}_A where all counters are diverging. Then $\mathcal{P}(\text{Run}(\mathbf{pv}, C)) > 0$ iff there is a \mathcal{Z}_{all} -safe finite path of the form $\mathbf{pv} \rightarrow^* \mathbf{qu} \rightarrow^* \mathbf{qz}$ where $q \in C$, \mathbf{qu} is above 1, $\mathbf{z} - \mathbf{u} \geq \mathbf{0}$, and $(\mathbf{z} - \mathbf{u})[i] > 0$ for every i such that $\mathbf{t}[i] > 0$.*

Proof. We start with “ \Rightarrow ”. Let \mathbf{t} be the trend of C . We show that for almost all $w \in \text{Run}(\mathbf{pv}, C)$ and all $i \in \{1, \dots, d\}$, one of the following conditions holds:

- (A) $\mathbf{t}[i] > 0$ and $\liminf_{k \rightarrow \infty} \text{eval}_i(w(k)) = \infty$,
- (B) $\mathbf{t}[i] = 0$ and $\text{eval}_i(w(k)) \geq \text{botfin}_i(\text{state}(w(k)))$ for all k 's large enough.

First, recall that C is also a BSCC of \mathcal{F}_{B_i} , and realize that the trend of the (only) counter in the BSCC C of \mathcal{F}_{B_i} is $\mathbf{t}[i]$.

Concerning (A), it follows, e.g., from the results of [13], that almost all runs $w' \in \text{Run}_{\mathcal{M}_{B_i}}(p(1))$ that stay in C and do not visit a configuration with zero counter satisfy $\liminf_{k \rightarrow \infty} \text{eval}_1(w'(k)) = \infty$. In particular, this means that almost all $w' \in \Upsilon_i(\text{Run}(\mathbf{pv}, C))$ satisfy this property. Hence, by Lemma 2, for almost all $w \in \text{Run}(\mathbf{pv}, C)$ we have that $\liminf_{k \rightarrow \infty} \text{eval}_i(w(k)) = \infty$.

Concerning (B), note that almost all runs $w \in \text{Run}(\mathbf{pv}, C)$ satisfying $\text{eval}_i(w'(k)) < \text{botfin}_i(\text{state}(w(k)))$ for infinitely many k 's eventually visit zero in some counter (there is a path of length at most $|C|$ from each such $w(k)$ to a configuration with zero in counter i , or in one of the other counters).

The above claim immediately implies that for every $k \in \mathbb{N}$, almost every run of $\text{Run}(\mathbf{pv}, C)$ visits a configuration \mathbf{qu} above k . Hence, there must be a \mathcal{Z}_{all} -safe path of the form $\mathbf{pv} \rightarrow^* \mathbf{qu} \rightarrow^* \mathbf{qz}$ with the required properties.

“ \Leftarrow ”: If there is a \mathcal{Z}_{all} -safe path of the form $\mathbf{pv} \rightarrow^* \mathbf{qu} \rightarrow^* \mathbf{qz}$ where $q \in C$, \mathbf{qu} is above 1, $\mathbf{z} - \mathbf{u} \geq \mathbf{0}$, and $(\mathbf{z} - \mathbf{u})[i] > 0$ for

every i such that $\mathbf{t}[i] > 0$, then \mathbf{pv} can reach a configuration \mathbf{qy} above k for an arbitrarily large $k \in \mathbb{N}$ via a \mathcal{Z}_{all} -safe path.

By Lemma 4, there exists $k \in \mathbb{N}$ such that for every $i \in \{1, \dots, d\}$ where $\mathbf{t}[i] > 0$ and every $n \geq k$, the probability of all $w \in \text{Run}_{\mathcal{M}_{B_i}}(q(n))$ that visit a configuration with zero counter is strictly smaller than $1/d$. Let \mathbf{qy} be a configuration above k reachable from \mathbf{pv} via a \mathcal{Z}_{all} -safe path (as shown above, the existence of such a \mathbf{qy} follows from the existence of $\mathbf{pv} \rightarrow^* \mathbf{qu} \rightarrow^* \mathbf{qz}$). It suffices to show that $\mathcal{P}(\text{Run}(\mathbf{qy}, \mathcal{Z}_{\text{all}})) < 1$. For every $i \in \{1, \dots, d\}$ where $\mathbf{t}[i] > 0$, let R_i be the set of all $w \in \text{Run}(\mathbf{qy}, \mathcal{Z}_{\text{all}})$ such that $\text{eval}_i(w(k)) = 0$ for some $k \in \mathbb{N}$ and all counters stay positive in all $w(k')$ where $k' < k$. Clearly, $\text{Run}(\mathbf{qy}, \mathcal{Z}_{\text{all}}) = \bigcup_i R_i$, and thus we obtain

$$\mathcal{P}(\text{Run}(\mathbf{qy}, \mathcal{Z}_{\text{all}})) \leq \sum_i \mathcal{P}(R_i) = \sum_i \mathcal{P}(\Upsilon_i(R_i)) < d \cdot \frac{1}{d} = 1$$

□

The following lemma shows that it is possible to decide, whether for a given $n \in \mathbb{N}$ a configuration above n can be reached via a \mathcal{Z}_{all} -safe path. Its proof uses the results of [8] on the coverability problem in (non-stochastic) VASS.

Lemma 6. *Let C be a BSCC of \mathcal{F}_A where all counters are diverging and let $q \in C$. There is a \mathcal{Z}_{all} -safe finite path of the form $\mathbf{pv} \rightarrow^* \mathbf{qu}$ with \mathbf{qu} above some $n \in \mathbb{N}$ iff there is a \mathcal{Z}_{all} -safe finite path of length at most $(|Q| + 2 \cdot |\gamma|) \cdot (3+n)^{(3d)^{d+1}}$ of the form $\mathbf{pv} \rightarrow^* \mathbf{qu}'$ with \mathbf{qu}' above n . Moreover, the existence of such a path can be decided in time $(|A| \cdot n)^{c \cdot 2^{d \log(d)}}$ where c is a fixed constant independent of d and A .*

Proof. We employ a decision procedure of [8] for VASS coverability. Since we need to reach \mathbf{qu}' above n via a \mathcal{Z}_{all} -safe finite path, we transform \mathcal{A} into a (non-probabilistic) VASS \mathcal{A}' whose control states and rules are determined as follows: for every rule (p, α, θ, q) of \mathcal{A} , we add to \mathcal{A}' the control states p, q together with two auxiliary fresh control states q', q'' , and we also add the rules $(p, -\mathbf{1}, q')$, $(q', \mathbf{1}, q'')$, (q'', α, q) . Hence, \mathcal{A}' behaves like \mathcal{A} , but when some counter becomes zero, then \mathcal{A}' is stuck (i.e., no transition is enabled except for the self-loop). Now it is easy to check that \mathbf{pv} can reach a configuration \mathbf{qu} above n via a \mathcal{Z}_{all} -safe finite path in \mathcal{A} iff \mathbf{pv} can reach a configuration \mathbf{qu} above n via some finite path in \mathcal{A}' , which is exactly the coverability problem for VASS. Theorem 1 in [8] shows that such a configuration can be reached iff there is configuration \mathbf{qu}' above n reachable via some finite path of length at most $m = (|Q| + 2 \cdot |\gamma|) \cdot (3+n)^{(3d)^{d+1}}$. (The term $(|Q| + 2 \cdot |\gamma|)$ represents the number of control states of \mathcal{A}' .) This path induces, in a natural way, a \mathcal{Z}_{all} -safe path from \mathbf{pv} to \mathbf{qu}' in \mathcal{A} of length at most $m/2$. Moreover, Theorem 2 in [8] shows that the existence of such a path in \mathcal{A}' can be decided in time $(|Q| + 2 \cdot |\gamma|) \cdot (3+n)^{2^{O(d \log(d))}}$, which proves the lemma. □

Theorem 1. *The qualitative \mathcal{Z}_{all} -reachability problem for d -dimensional pMC is decidable in time $|A|^{\kappa \cdot 2^{d \log(d)}}$, where κ is a fixed constant independent of d and A .*

Proof. Note that the Markov chain \mathcal{F}_A is computable in time polynomial in $|A|$ and d , and we can efficiently identify all diverging BSCCs of \mathcal{F}_A . For each diverging BSCC C , we need to check the condition of Lemma 5. By applying Lemma 2.3. of [29], we obtain that if there exist some \mathbf{qu} above 1 and a \mathcal{Z}_{all} -safe finite path of the form $\mathbf{qu} \rightarrow^* \mathbf{qz}$ such that $\mathbf{z} - \mathbf{u} \geq \mathbf{0}$ and $(\mathbf{z} - \mathbf{u})[i] > 0$ for every i where $\mathbf{t}[i] > 0$, then such a path exists for every \mathbf{qu} above $|A|^{c \cdot d}$ and its length is bounded by $|A|^{c \cdot d}$. Here c is a fixed constant independent of $|A|$ and d (let us note that Lemma 2.3. of

¹ The precise bound on h is given in the proof of Proposition 7 [14].

² The bound on δ is given in Proposition 6 [14].

[29] is formulated for vector addition systems without states and a non-strict increase in every counter, but the corresponding result for VASS is easy to derive; see also Lemma 15 in [12]). Hence, the existence of such a path for a given $q \in C$ can be decided in $|\mathcal{A}|^{d^{\mathcal{O}(1)}}$ time, e.g., by simple inductive enumeration of all configurations that can be reached via a path of length at most $|\mathcal{A}|^{c \cdot d}$. It remains to check whether $p\mathbf{v}$ can reach a configuration $q\mathbf{u}$ above $|\mathcal{A}|^{c \cdot d}$ via a \mathcal{Z}_{all} -safe finite path. By Lemma 6 this can be done in time $(|\mathcal{A}| \cdot |\mathcal{A}|^{c \cdot d})^{c' \cdot 2^{d \log(d)}}$ for another constant c' . This gives us the desired complexity bound. \square

Note that for every fixed dimension d , the qualitative \mathcal{Z}_{all} -reachability problem is solvable in polynomial time.

Now we show that $\mathcal{P}(\text{Run}(p\mathbf{v}, \mathcal{Z}_{all}))$ can be effectively approximated up to an arbitrarily small absolute/relative error $\varepsilon > 0$. A full proof of Theorem 2 can be found in [16].

Theorem 2. *For a given d -dimensional pMC \mathcal{A} and its initial configuration $p\mathbf{v}$, the probability $\mathcal{P}(\text{Run}(p\mathbf{v}, \mathcal{Z}_{all}))$ can be approximated up to a given absolute error $\varepsilon > 0$ in time $(\exp(|\mathcal{A}|) \cdot \log(1/\varepsilon))^{\mathcal{O}(d \cdot d!)}$.*

Proof sketch. First we check whether $\mathcal{P}(\text{Run}(p\mathbf{v}, \mathcal{Z}_{all})) = 1$ (using the algorithm of Theorem 1) and return 1 if it is the case. Otherwise, we first show how to approximate $\mathcal{P}(\text{Run}(p\mathbf{v}, \mathcal{Z}_{all}))$ under the assumption that p is in some diverging BSCC of $\mathcal{F}_{\mathcal{A}}$, and then we show how to drop this assumption.

So, let C be a diverging BSCC of $\mathcal{F}_{\mathcal{A}}$ with $\mathcal{P}(\text{Run}(p\mathbf{v}, C)) < 1$, and let us assume that $p \in C$. We show how to compute a number $\nu > 0$ such that $|\mathcal{P}(\text{Run}(p\mathbf{v}, \mathcal{Z}_{all})) - \nu| \leq d \cdot \varepsilon$ in time $(\exp(|\mathcal{A}|) \cdot \log(1/\varepsilon))^{\mathcal{O}(d!)}$. We proceed by induction on d . The key idea of the inductive step is to find a sufficiently large constant K such that if some counter reaches K , it can be safely “forgotten”, i.e., replaced by ∞ , without influencing the probability of reaching zero in some counter by more than ε . Hence, whenever we visit a configuration $q\mathbf{u}$ where some counter value in \mathbf{u} reaches K , we can apply induction hypothesis and approximate the probability of reaching zero in some counter from $q\mathbf{u}$ by “forgetting” the large counter a thus reducing the dimension. Obviously, there are only finitely many configurations where all counters are below K , and here we employ the standard methods for finite-state Markov chains. The number K is computed by using the bounds of Lemma 4.

Let us note that the base (when $d = 1$) is handled by relying only on Lemma 4. Alternatively, we could employ the results of [18]. This would improve the complexity for $d = 1$, but not for higher dimensions.

Finally, we show how to approximate $\mathcal{P}(\text{Run}(p\mathbf{v}, \mathcal{Z}_{all}))$ when the control state p does not belong to a BSCC of $\mathcal{F}_{\mathcal{A}}$. Here we use the bound of Lemma 1. \square

Note that if $\mathcal{P}(\text{Run}(p\mathbf{v}, \mathcal{Z}_{all})) > 0$, then this probability is at least $p_{min}^{m \cdot |\mathcal{Q}|}$ where p_{min} is the least positive transition probability in $\mathcal{M}_{\mathcal{A}}$ and m is the maximal component of \mathbf{v} . Hence, Theorem 2 can also be used to approximate $\mathcal{P}(\text{Run}(p\mathbf{v}, \mathcal{Z}_{all}))$ up to a given relative error $\varepsilon > 0$.

3.2 Zero-Reachability, Case II

Let us fix a (non-labeled) pMC $\mathcal{A} = (Q, \gamma, W)$ of dimension $d \in \mathbb{N}^+$ and $i \in \{1, \dots, d\}$. As in the previous section, our aim is to identify the conditions under which $\text{Run}(p\mathbf{1}, \neg\mathcal{Z}_{-i}) > 0$. Without restrictions, we assume that $i = d$, i.e., we consider $\mathcal{Z}_{-d} = \{\{1\}, \dots, \{d-1\}\}$. Also, for technical reasons, we assume that $\mathcal{P}(\text{Run}(p\mathbf{1}, \neg\mathcal{Z}_{-d})) = \mathcal{P}(\text{Run}(p\mathbf{u}^{in}, \neg\mathcal{Z}_{-d}))$ where $\mathbf{u}^{in}[i] = 1$ for all $i \in \{1, \dots, d-1\}$ but $\mathbf{u}^{in}[d] = 0$. (Note that

every pMC can be easily modified in polynomial time so that this condition is satisfied.)

To analyze the runs of $\text{Run}(p\mathbf{u}^{in}, \neg\mathcal{Z}_{-d})$, we re-use the finite-state Markov chain $\mathcal{F}_{\mathcal{A}}$ from Section 3.1. Intuitively, the chain $\mathcal{F}_{\mathcal{A}}$ is useful for analyzing those runs of $\text{Run}(p\mathbf{u}^{in}, \neg\mathcal{Z}_{-d})$ where all counters stay positive. Since the structure of $\text{Run}(p\mathbf{u}^{in}, \neg\mathcal{Z}_{-d})$ is more complex than in Section 3.1, we also need some new analytic tools.

We also re-use the L -labeled 1-dimensional pMC \mathcal{B}_d to deal with runs that visit zero in counter d infinitely many times. To simplify notation, we use \mathcal{B} to denote \mathcal{B}_d . The behaviour of \mathcal{B} is analyzed using the finite-state Markov chain \mathcal{X} (see Definition 6 below) that has been employed already in [13] to design a model-checking algorithm for linear-time properties in one-dimensional pMC.

Let us denote by $[q \downarrow r]$ the probability that a run of $\mathcal{M}_{\mathcal{B}}$ initiated in $q(0)$ visits the configurations $r(0)$ without visiting any configuration of the form $r'(0)$ (where $r' \neq r$) in between. Given $q \in Q$, we denote by $[q \uparrow]$ the probability $1 - \sum_{r \in Q} [q \downarrow r]$ that a run initiated in $q(0)$ never visits a configuration with zero counter value (except for the initial one).

Definition 6. *Let $\mathcal{X}_{\mathcal{B}} = (X, \rightarrow, \text{Prob})$ be a non-labelled finite-state Markov chain where $X = Q \cup \{q \uparrow \mid q \in Q\}$ and the transitions are defined as follows:*

- $q \xrightarrow{x} r$ iff $0 < x = [q \downarrow r]$;
- $q \xrightarrow{x} q \uparrow$ iff $0 < x = [q \uparrow]$;
- for every $q \in Q$ we have $q \uparrow \xrightarrow{1} q \uparrow$;
- there are no other transitions.

The correspondence between the runs of $\text{Run}_{\mathcal{M}_{\mathcal{B}}}(p(0))$ and $\text{Run}_{\mathcal{X}_{\mathcal{B}}}(p)$ is captured by a function $\Phi : \text{Run}_{\mathcal{M}_{\mathcal{B}}}(p(0)) \rightarrow \text{Run}_{\mathcal{X}_{\mathcal{B}}}(p) \cup \{\perp\}$, where $\Phi(w)$ is obtained from a given $w \in \text{Run}_{\mathcal{M}_{\mathcal{B}}}(p(0))$ as follows:

- First, each maximal subpath in w of the form $q(0), \dots, r(0)$ such that the counter stays positive in all of the intermediate configurations is replaced with a single transition $q \rightarrow r$.
- Note that if w contained infinitely many configurations with zero counter, then the resulting sequence is a run of $\text{Run}_{\mathcal{X}_{\mathcal{B}}}(p)$, and thus we obtain our $\Phi(w)$. Otherwise, the resulting sequence takes the form $v \hat{w}$, where $v \in \text{FPath}_{\mathcal{X}_{\mathcal{B}}}(p)$ and \hat{w} is a suffix of w initiated in a configuration $r(1)$. Let q be the last state of v . Then, $\Phi(w)$ is either $v(q \uparrow)^\omega$ or \perp , depending on whether $[q \uparrow] > 0$ or not, respectively (here, $(q \uparrow)^\omega$ is a infinite sequence of $q \uparrow$).

Lemma 7. *For every measurable subset $R \subseteq \text{Run}_{\mathcal{X}_{\mathcal{B}}}(p)$ we have that $\Phi^{-1}(R)$ is measurable and $\mathcal{P}(R) = \mathcal{P}(\Phi^{-1}(R))$.*

A proof of Lemma 7 is straightforward (it suffices to check that the lemma holds for all basic cylinders $\text{Run}_{\mathcal{X}_{\mathcal{B}}}(w)$ where $w \in \text{FPath}_{\mathcal{X}_{\mathcal{B}}}(p)$). Note that Lemma 7 implies $\mathcal{P}(\Phi = \perp) = 0$.

Let D_1, \dots, D_k be all BSCCs of $\mathcal{X}_{\mathcal{B}}$ reachable from p . Further, for every D_j , we use $\text{Run}(p\mathbf{u}^{in}, D_j)$ to denote the set of all $w \in \text{Run}_{\mathcal{M}_{\mathcal{A}}}(p\mathbf{u}^{in}, \neg\mathcal{Z}_{-d})$ such that $\Phi(\Upsilon_d(w)) \neq \perp$ and $\Phi(\Upsilon_d(w))$ visits D_j . Observe that

$$\mathcal{P}(\text{Run}_{\mathcal{M}_{\mathcal{A}}}(p\mathbf{u}^{in}, \neg\mathcal{Z}_{-d})) = \sum_{j=1}^k \mathcal{P}(\text{Run}(p\mathbf{u}^{in}, D_j)) \quad (2)$$

Indeed, note that almost all runs w of $\text{Run}_{\mathcal{X}_{\mathcal{B}}}(p)$ visit some D_j , and hence by Lemma 7, we obtain that $\Phi(w)$ visits some D_j for almost all $w \in \text{Run}_{\mathcal{M}_{\mathcal{B}}}(p(1))$. In particular, for almost all w of $\Upsilon_d(\text{Run}_{\mathcal{M}_{\mathcal{A}}}(p\mathbf{u}^{in}, \neg\mathcal{Z}_{-d}))$ we have that $\Phi(w)$ visits some D_j .

By Lemma 2, for almost all $w \in \text{Run}_{\mathcal{M}_A}(\mathbf{p}\mathbf{u}^{in}, \neg\mathcal{Z}_{-d})$, the run $\Phi(\Upsilon_d(w))$ visits some D_j , which proves Equation (2).

Now we examine the runs of $\text{Run}(\mathbf{p}\mathbf{u}^{in}, D_j)$ in greater detail and characterize the conditions under which $\mathcal{P}(\text{Run}(\mathbf{p}\mathbf{u}^{in}, D_j)) > 0$. Note that for every BSCC D in \mathcal{X}_B we have that either $D = \{q\uparrow\}$ for some $q \in Q$, or $D \subseteq Q$. We treat these two types of BSCCs separately, starting with the former.

Lemma 8. $\mathcal{P}(\bigcup_{q \in Q} \text{Run}(\mathbf{p}\mathbf{u}^{in}, \{q\uparrow\})) > 0$ iff there exists a BSCC C of \mathcal{F}_A with all counters diverging and a \mathcal{Z}_{-d} -safe finite path of the form $\mathbf{p}\mathbf{u}^{in} \rightarrow^* \mathbf{q}\mathbf{u} \rightarrow^* \mathbf{q}\mathbf{z}$ where the subpath $\mathbf{q}\mathbf{u} \rightarrow^* \mathbf{q}\mathbf{z}$ is \mathcal{Z}_{all} -safe, $q \in C$, $\mathbf{q}\mathbf{u}$ is above 1, $\mathbf{z} - \mathbf{u} \geq \mathbf{0}$, and $(\mathbf{z} - \mathbf{u})[i] > 0$ for every i such that $\mathbf{t}[i] > 0$.

Now let D be a BSCC of \mathcal{X}_B reachable from p such that $D \subseteq Q$ (i.e., $D \neq \{q\uparrow\}$ for any $q \in Q$). Let $\mathbf{e} \in [1, \infty]^D$ where $\mathbf{e}[q]$ is the expected number of transitions needed to revisit a configuration with zero counter from $q(0)$ in \mathcal{M}_B .

Proposition 1 ([13], Corollary 6). *The problem whether $\mathbf{e}[q] < \infty$ is decidable in polynomial time.*

From now on, we assume that $\mathbf{e}[q] < \infty$ for all $q \in D$.

In Section 3.1, we used the trend $\mathbf{t} \in \mathbb{R}^d$ to determine tendency of counters either to diverge, or to reach zero. As defined, each $\mathbf{t}[i]$ corresponds to the long-run average change per transition of counter i as long as all counters stay positive. Allowing zero value in counter d , the trend $\mathbf{t}[i]$ is no longer equal to the long-run average change per transition of counter i and hence it does not correctly characterize its behavior. Therefore, we need to redefine the notion of trend in this case.

Recall that \mathcal{B} is $L = \{-1, 0, 1\}^{d-1}$ -labeled pMC. Given $i \in \{1, \dots, d-1\}$, we denote by $\boldsymbol{\delta}_i \in \mathbb{R}^Q$ the vector where $\boldsymbol{\delta}_i[q]$ is the i -th component of the expected total reward accumulated along a run from $q(0)$ before revisiting another configuration with zero counter. Formally, $\boldsymbol{\delta}_i[q] = \mathbb{E}T_i$ where T_i is a random variable which to every $w \in \text{Run}_{\mathcal{M}_B}(q(0))$ assigns $\text{tot}_i(w; \ell)$ such that $\ell > 0$ is the least number satisfying $w(\ell) = r(0)$ for some $r \in D$.

Let $\boldsymbol{\mu}_{oc} \in [0, 1]^D$ be the invariant distribution of the BSCC D of \mathcal{X}_B , i.e., $\boldsymbol{\mu}_{oc}$ is the unique solution of

$$\boldsymbol{\mu}_{oc}[q] = \sum_{r \in D, r \overset{x}{\rightarrow} q} \boldsymbol{\mu}_{oc}[r] \cdot x$$

The *oc-trend* of D is a $(d-1)$ -dimensional vector $\mathbf{t}_{oc} \in [-1, 1]^{d-1}$ defined by

$$\mathbf{t}_{oc}[i] = \left(\boldsymbol{\mu}_{oc}^T \cdot \boldsymbol{\delta}_i \right) / \left(\boldsymbol{\mu}_{oc}^T \cdot \mathbf{e} \right)$$

The following lemma follows from the standard results about irreducible Markov chains (see, e.g., [28]).

Lemma 9. *For almost all $w \in \text{Run}_{\mathcal{M}_B}(q(0))$ we have that*

$$\mathbf{t}_{oc}[i] = \lim_{k \rightarrow \infty} \frac{\text{tot}_i(w; k)}{k}$$

That is, $\mathbf{t}_{oc}[i]$ is the i -th component of the expected long-run average reward per transition in a run of $\text{Run}_{\mathcal{M}_B}(q(0))$, and as such, determines the long-run average change per transition of counter i as long as all counters of $\{1, \dots, d-1\}$ remain positive.

Further, for every $i \in \{1, \dots, d-1\}$ and every $q \in D$, we denote by $\text{botinf}_i(q)$ the least $j \in \mathbb{N}$ such that every $w \in \text{FPath}_{\mathcal{M}_B}(q(0))$ ending in $q(0)$ satisfies $\text{tot}_i(w; k) \geq -j$ for all $0 \leq k \leq \text{length}(w)$. If there is no such j , we put $\text{botinf}_i(q) = \infty$. It is easy to show that if $\text{botinf}_i(q) = \infty$, then $\text{botinf}_i(r) = \infty$ for all $r \in D$.

Lemma 10. *If $\text{botinf}_i(q) < \infty$, then $\text{botinf}_i(q) \leq (|Q| + 1)^3$ and the exact value of $\text{botinf}_i(q)$ is computable in time polynomial in $|A|$.*

We say that counter i is *oc-decreasing* in D if $\text{botinf}_i(q) = \infty$ for some (and hence all) $q \in D$.

Definition 7. *For a given $i \in \{1, \dots, d-1\}$, we say that the i -th reward is *oc-diverging* in D if either $\mathbf{t}_{oc}[i] > 0$, or $\mathbf{t}_{oc}[i] = 0$ and counter i is not *oc-decreasing* in D .*

Lemma 11. *If some reward is not *oc-diverging* in D , then we have $\mathcal{P}(\text{Run}(\mathbf{p}\mathbf{u}^{in}, D)) = 0$.*

It remains to analyze the case when all rewards are *oc-diverging* in D . Similarly to Case I, we need to obtain a bound on probability of divergence of an arbitrary counter $i \in \{1, \dots, d-1\}$ with $\mathbf{t}_{oc}[i] > 0$. The following lemma (an analogue of Lemma 4) is crucial in the process.

Lemma 12. *Let \mathcal{D} be a $\{-1, 0, 1\}$ -labeled one-dimensional pMC with a set of states Q , and let $D \subseteq Q$ be a BSCC of \mathcal{X}_D such that the *oc-trend* \mathbf{t}_{oc} of the only reward in D is positive. Then for all $q \in D$, there exist computable constants h' and A_0 where $0 < A_0 < 1$, such that for all $h \geq h'$ we have that the probability that a run $w \in \text{Run}_{\mathcal{M}_D}(q(0))$ satisfies*

$$\inf_{k \in \mathbb{N}} \text{tot}_1(w; k) \geq -h$$

is at least $1 - A_0^h$.

A proof of Lemma 12 is the most involved part of this paper, where we need to construct new analytic tools. A sketch of the proof is included at the end of this section.

Definition 8. *Let D be a BSCC of \mathcal{X}_B where all rewards are *oc-diverging*, and let $q \in D$. We say that a configuration $\mathbf{q}\mathbf{u}$ is *oc-above* a given $n \in \mathbb{N}$ if $\mathbf{u}[i] \geq n$ for every $i \in \{1, \dots, d-1\}$ such that $\mathbf{t}_{oc}[i] > 0$, and $\mathbf{u}[i] \geq \text{botinf}_i(q)$ for every $i \in \{1, \dots, d-1\}$ such that $\mathbf{t}_{oc}[i] = 0$.*

The next lemma is an analogue of Lemma 5 and it is proven using the same technique, using Lemma 12 instead of Lemma 4.

Lemma 13. *Let D be a BSCC of \mathcal{X}_B where all rewards are *diverging*. Then there exists a computable constant $n \in \mathbb{N}$ such that $\mathcal{P}(\text{Run}(\mathbf{p}\mathbf{u}^{in}, D)) > 0$ iff there is a \mathcal{Z}_{-d} -safe finite path of the form $\mathbf{p}\mathbf{u}^{in} \rightarrow^* \mathbf{q}\mathbf{u}$ where $q \in D$, \mathbf{u} is *oc-above* n , and $\mathbf{u}[d] = 0$.*

A direct consequence of Lemma 13 and the results of [7] is the following:

Theorem 3. *The qualitative \mathcal{Z}_{-d} -reachability problem for d -dimensional pMC is decidable (assuming $\mathbf{e}[q] < \infty$ for all $q \in D$ in every BSCC of \mathcal{X}_B).*

A proof of Theorem 3 is straightforward, since we can effectively compute the structure of \mathcal{X}_B (in time polynomial in $|A|$), express its transition probabilities and *oc-trends* in BSCCs of \mathcal{X}_B in the existential fragment of Tarski algebra, and thus effectively identify all BSCCs of \mathcal{X}_B where all rewards are *oc-diverging*. To check the condition of Lemma 13, we use the algorithm of [7] for constructing finite representation of filtered covers in VAS with one zero test. This is the only part where we miss an upper complexity bound, and therefore we cannot provide any bound in Theorem 3. It is worth noting that the qualitative \mathcal{Z}_{-d} -reachability problem is SQUARE-ROOT-SUM-hard (see below), and hence it cannot be solved efficiently without a breakthrough results in the complexity of exact algorithms. For more comments and a proof of the next Proposition, see [16].

Proposition 2. *The qualitative \mathcal{Z}_{-d} -reachability problem is SQUARE-ROOT-SUM-hard, even for two-dimensional pMC where $\mathbf{e}[q] < \infty$ for all $q \in D$ in every BSCC of $\mathcal{X}_{\mathcal{B}}$.*

Using Lemma 13, we can also approximate $\mathcal{P}(\text{Run}(\mathbf{p}\mathbf{v}, \mathcal{Z}_{-d}))$ up to an arbitrarily small absolute error $\varepsilon > 0$ (due to the problems mentined above, we do not provide any complexity bounds). The procedure mimics the one of Theorem 2. The difference is that now we eventually use methods for one-dimensional pMC instead of the methods for finite-state Markov chains.

Theorem 4. *For a given d -dimensional pMC \mathcal{A} and its initial configuration $\mathbf{p}\mathbf{v}$, the probability $\mathcal{P}(\text{Run}(\mathbf{p}\mathbf{v}, \mathcal{Z}_{-d}))$ can be effectively approximated up to a given absolute error $\varepsilon > 0$.*

A Proof of Lemma 12. The lemma differs from Lemma 4 in that it effectively bounds the probability of not reaching zero in one of the counters of a two-dimensional pMC. Hence, the results on one-dimensional pMCs are not sufficient here. Below, we sketch a stronger method that allows us to prove the lemma. The method is again based on analysing a suitable martingale; however, the construction and structure of the martingale is much more complex than in the one-dimensional case.

Before we show how to construct the desired martingale, let us mention the following useful lemma:

Lemma 14. *Let $r \in D$. Given a run $w \in \text{Run}_{\mathcal{M}_{\mathcal{B}}}(r(0))$, we denote by $E(w) = \inf\{\ell > 0 \mid \text{cval}_1(w(\ell)) = 0\}$, i.e., the time it takes w to re-visit zero counter value. Then there are constants $c' \in \mathbb{N}$ and $a \in (0, 1)$ computable in polynomial space such that for all $k \geq c'$ we have*

$$\mathcal{P}(E \geq k) \leq a^k$$

Proof. This follows immediately from Proposition 6 and Theorem 7 in [15]. \square

Let us fix an 1-dimensional pMC \mathcal{D} with the set of states Q and let $D \subseteq Q$ be a BSCC of $\mathcal{X}_{\mathcal{D}}$ in which the oc-trend of the only reward is positive. Let us summarize the notation used throughout the proof.

- Recall that $\mathbf{e} \in [1, \infty]^D$ is the vector such that $\mathbf{e}[q]$ is the expected total time of a nonempty run from $q(0)$ to the first visit of $r(0)$ for some $r \in Q$. By our assumptions, \mathbf{e}_{\downarrow} is finite.
- Let $\text{Post}(D)$ be the set of all states $q \in Q$ such that there is an $r \in D$ and $\ell \in \mathbb{N}$ such that $q(\ell)$ is reachable from $r(0)$ in D .
- Let $\mathbf{e}_{\downarrow} \in [1, \infty]^{\text{Post}(D)}$ be the vector such that $\mathbf{e}_{\downarrow}[q]$ is the expected total time of a run from $q(1)$ to the first visit of $r(0)$ for some $r \in Q$. Since \mathbf{e} is finite, also \mathbf{e}_{\downarrow} is finite.
- Recall that $\boldsymbol{\delta}_1 \in \mathbb{R}^D$ is the vector such that $\boldsymbol{\delta}_1[q]$ is the expected total reward accumulated during a nonempty run from $q(0)$ to the first visit of $r(0)$ for some $r \in Q$. Since $|\boldsymbol{\delta}_1[q]| \leq |\mathbf{e}[q]|$ holds for all $q \in Q$, the vector $\boldsymbol{\delta}_1$ is finite.
- Let $\boldsymbol{\delta}_{\downarrow} \in \mathbb{R}^{\text{Post}(D)}$ be the vector such that $\boldsymbol{\delta}_{\downarrow}[q]$ is the expected total reward accumulated during a run from $q(1)$ to the first visit of $r(0)$ for some $r \in Q$. Similarly as before, $\boldsymbol{\delta}_{\downarrow}$ is finite.
- Let us denote by $A \in \mathbb{R}^{D \times D}$ transition matrix of the irreducible Markov chain induced by BSCC D of $\mathcal{X}_{\mathcal{D}}$, i.e., $A[q, r]$ is the probability that starting from $q(0)$ the configuration $r(0)$ is visited before visiting any configuration $r'(0)$ for any $r' \neq r$. Clearly the matrix A is stochastic and irreducible.
- Let $G \in \mathbb{R}^{\text{Post}(D) \times \text{Post}(D)}$ denote the matrix such that $G[q, r]$ is the probability that starting from $q(1)$ the configuration $r(0)$ is visited before visiting any configuration $r'(0)$ for any $r' \neq r$. By our assumptions the matrix G is stochastic, i.e., $G\mathbf{1} = \mathbf{1}$.

- Recall that $\boldsymbol{\mu}_{oc}^T = \boldsymbol{\mu}_{oc}^T A \in [0, 1]^D$ denotes the invariant distribution of the finite Markov chain $\mathcal{X}_{\mathcal{D}}$ induced by A .
- Recall that $t = (\boldsymbol{\mu}_{oc}^T \boldsymbol{\delta}_1) / (\boldsymbol{\mu}_{oc}^T \mathbf{e}) \in [-1, +1]$ is the oc-trend of D , so intuitively t is the expected average reward per step accumulated during a run started from $q(0)$ for some $q \in D$.
- Let $\mathbf{r}_{\downarrow} := \boldsymbol{\delta}_{\downarrow} - t\mathbf{e}_{\downarrow} \in \mathbb{R}^{\text{Post}(D)}$ and let $\mathbf{r}_0 := \boldsymbol{\delta}_1 - t\mathbf{e} \in \mathbb{R}^D$.

Lemma 15. *There exists a vector $\mathbf{g}(0) \in \mathbb{R}^{\text{Post}(D)}$ such that*

$$\mathbf{g}(0)[D] = \mathbf{r}_0 + A\mathbf{g}(0)[D], \quad (3)$$

where $\mathbf{g}(0)[D]$ denotes the vector obtained from $\mathbf{g}(0)$ by deleting the non- D -components.

Proof sketch. The proof is based on the notion of *group inverses* for matrices [17]. Close connections of this concept to (finite) Markov chains are discussed in [25]. In [16] we show that for any non-negative irreducible matrix P with spectral radius³ equal to 1 there is a matrix, denoted by $(I - P)^{\#}$, such that $(I - P) \cdot (I - P)^{\#} = I - W$, where W is a matrix whose rows are scalar multiples of a dominant left eigenvector of P (i.e., a left eigenvector corresponding to the eigenvalue of maximal absolute value). Then, we prove Lemma 15 as follows:

Recall that the matrix A is stochastic and irreducible. Also, note that the invariant distribution $\boldsymbol{\mu}_{oc}^T$ satisfies $\boldsymbol{\mu}_{oc}^T A = \boldsymbol{\mu}_{oc}^T$, i.e., it is a left eigenvector of A with the corresponding eigenvalue equal to 1. By the Perron-Frobenius theorem, an eigenvector with strictly positive components (such as $\boldsymbol{\mu}_{oc}^T$) of an irreducible non-negative matrix must be a dominant eigenvector of this matrix. It follows that the spectral radius of A is 1, so $(I - A)^{\#}$ exists.

Now define $\mathbf{g}(0)[D] := (I - A)^{\#} \mathbf{r}_0$. The non- D -components can be set arbitrarily, for instance, they can be set to 0. So we have $\mathbf{g}(0)[D] = \mathbf{r}_0 + A\mathbf{g}(0)[D] - W\mathbf{r}_0$, where the rows of W are multiples of $\boldsymbol{\mu}_{oc}^T$. We have:

$$\begin{aligned} \boldsymbol{\mu}_{oc}^T \mathbf{r}_0 &= \boldsymbol{\mu}_{oc}^T \left(\boldsymbol{\delta}_1 - \frac{\boldsymbol{\mu}_{oc}^T \boldsymbol{\delta}_1}{\boldsymbol{\mu}_{oc}^T \mathbf{e}} \mathbf{e} \right) \quad \text{by the definitions of } \mathbf{r}_0 \text{ and } t \\ &= 0. \end{aligned}$$

So (3) follows. \square

Now take an arbitrary vector $\mathbf{g}(0)$ satisfying (3) and extend it to a function $\mathbf{g} : \mathbb{N} \rightarrow \mathbb{R}^{\text{Post}(D)}$ inductively by putting

$$\mathbf{g}(n+1) = \mathbf{r}_{\downarrow} + G\mathbf{g}(n) \quad \text{for all } n \in \mathbb{N}. \quad (4)$$

Let us fix any $q \in D$ and any $h \in \mathbb{N}$. For a run $w \in \text{Run}_{\mathcal{M}_{\mathcal{D}}}(q(0))$ and all $\ell \in \mathbb{N}$ let $p^{(\ell)} \in Q$ and $x_1^{(\ell)}, x_2^{(\ell)} \in \mathbb{N}$ be such that $p^{(\ell)} = \text{state}(w(\ell))$, $x_2^{(\ell)} = \text{cval}(w(\ell))$ and $x_1^{(\ell)} = h + \text{tot}(w; \ell)$.

Finally, let us define

$$m^{(\ell)} := x_1^{(\ell)} - t\ell + \mathbf{g}(x_2^{(\ell)})[p^{(\ell)}] \quad \text{for all } \ell \in \mathbb{N}. \quad (5)$$

Then we have:

Proposition 3. *Write \mathcal{E} for the expectation with respect to \mathcal{P} . We have for all $\ell \in \mathbb{N}$:*

$$\mathcal{E} \left(m^{(\ell+1)} \mid w(\ell) \right) = m^{(\ell)}.$$

In other words, the stochastic process $\{m^{(\ell)}\}_{\ell=0}^{\infty}$ is a martingale, and this holds for any choice of vector $\mathbf{g}(0)$ satisfying (3) and any choice of $x^{(0)} = h \in \mathbb{N}$. For our purposes we need to pick $\mathbf{g}(0)$ and h in a rather specific way.

³The spectral radius of P is the maximal absolute value of an eigenvalue of P .

Lemma 16. *There is $\mathbf{g}(0)$ satisfying (3) such that for the function $\mathbf{g} : \mathbb{N} \rightarrow \mathbb{R}^{\text{Post}(D)}$ defined by (4) the following holds: There exists a constant c effectively computable in polynomial space such that for every $q \in D$ and $n \geq 1$ we have $|\mathbf{g}(0)[q]| \leq c$ and $|\mathbf{g}(n)[q]| \leq c \cdot n$.*

Proof sketch. First we show that there is a vector $\mathbf{g}^* \in \mathbb{R}^D$ with $\mathbf{g}^* = \mathbf{r}_0 + A\mathbf{g}^*$ and

$$0 \leq \mathbf{g}^*[q] \leq \frac{e_{\max}|D|}{y_{\min}^{|D|}} \quad \text{for all } q \in D, \quad (6)$$

where $e_{\max} := 1 + \max_{q \in \text{Post}(D)} \mathbf{e}_\downarrow[q]$ and y_{\min} denotes the smallest nonzero entry in the matrix A . To see this, pick an arbitrary vector $\mathbf{g}'(0)[D]$ satisfying (3) (whose existence is shown by Lemma 15) and observe that since $A\mathbf{1} = \mathbf{1}$, for every $\kappa \in \mathbb{R}$ it holds that $\mathbf{g}'(0)[D] + \kappa\mathbf{1} = \mathbf{r}_0 + A(\mathbf{g}'(0)[D] + \kappa\mathbf{1})$, i.e., $\mathbf{g}'(0)[D] + \kappa\mathbf{1}$ also satisfies (3). So let κ be such that for $\mathbf{g}^* := \mathbf{g}'(0)[D] + \kappa\mathbf{1}$ it holds $\max_{q \in D} \mathbf{g}^*[q] \leq e_{\max}|D|/y_{\min}^{|D|}$ (such κ clearly exists). In [16] we show that then the vector \mathbf{g}^* is non-negative, which shows that (6) holds.

Now put $\mathbf{g}(0)[D] = \mathbf{g}^*$ and define the non- D components of $\mathbf{g}(0)$ arbitrarily (e.g., make them zero). A straightforward induction, which we leave to the full version [16], shows that for every $n \in \mathbb{N}$ it holds $|\mathbf{g}(n)| \leq |\mathbf{g}(0)| + n|\mathbf{r}_\downarrow|$. Using (6) and the fact that $|\mathbf{r}_\downarrow| \leq |\mathbf{e}_\downarrow| \leq e_{\max}$, we get $|\mathbf{g}(n)| \leq e_{\max}|D|/y_{\min}^{|D|} + ne_{\max}$, from which the lemma follows. \square

Now let $\mathbf{g}(0)$ be the vector from Lemma 16 and $h \in \mathbb{N}$ be such that $(t \cdot \sqrt[4]{h})/c \geq c'$, where c is from Lemma 16 and c' from Lemma 14. As shown by Proposition 3, the stochastic process $\{m^{(\ell)}\}_{\ell=0}^\infty$ defined by (5) (where $\mathbf{g}(0)$ is extended to \mathbb{N} using (4)) is a martingale. Unfortunately, this martingale may still have unbounded differences, i.e. $|m_i^{(\ell+1)} - m_i^{(\ell)}|$ may become arbitrarily large with increasing ℓ , which prohibits us from applying standard tools of martingale theory (such as Azuma's inequality) directly on $\{m^{(\ell)}\}_{\ell=0}^\infty$. We now show how to overcome this difficulty.

Let us fix any $i \in \mathbb{N}$ such that $i \geq h$ and denote $K = (t \cdot \sqrt[4]{i})/c$. We define a new stochastic process $\{m_i^{(\ell)}\}_{\ell=0}^\infty$ as follows:

$$m_i^{(\ell)} := \begin{cases} m^{(\ell)} & \text{if } x_2^{(\ell)} \leq K \text{ for all } \ell' \leq \ell \\ m_i^{(\ell-1)} & \text{otherwise.} \end{cases} \quad (7)$$

Observe that $\{m_i^{(\ell)}\}_{\ell=0}^\infty$ is also a martingale. Moreover, using the bound of Lemma 16 we have for every $\ell \in \mathbb{N}$ that $|m_i^{(\ell+1)} - m_i^{(\ell)}| \leq 1 + t + 2cK \leq 4t\sqrt[4]{i}$, i.e., $\{m_i^{(\ell)}\}_{\ell=0}^\infty$ is a bounded-difference martingale.

Now let H_i be the set of all runs w that satisfy $x_1^{(i)} = 0$ and $x_1^{(\ell)} > 0$ for all $0 \leq \ell < i$. Moreover, denote by Over_i the set of all runs w such that $x_2^{(\ell)} \geq K$ for some $0 \leq \ell \leq i$, and by $\neg\text{Over}_i$ the complement of Over_i .

Note that every run can perform at most i -revisits of zero counter value during the first i steps. By Lemma 14 the probability that counter value at least K is reached between two visits of zero counter is at most a^K . It follows that $\mathcal{P}(\text{Over}_i) \leq i \cdot a^{(t \cdot \sqrt[4]{i})/c}$.

Next, for every run $w \in \neg\text{Over}_i \cap H_i$ it holds

$$\begin{aligned} (m_i^{(i)} - m_i^{(0)})(w) &= (m^{(i)} - m^{(0)})(w) \\ &= -it + \mathbf{g}(x_2^{(i)})[p^{(i)}] - h - \mathbf{g}(0)[p^{(0)}] \\ &\leq -it + 2cK = -it + 2 \cdot t \cdot \sqrt[4]{i} \leq -i \frac{t}{2}, \end{aligned}$$

where the first inequality follows from the bound on $\mathbf{g}(n)$ in Lemma 16 and the last inequality holds since $\sqrt[4]{i} \leq i/2$ for all $i \geq 3$.

Using the Azuma's inequality, we get

$$\begin{aligned} \mathcal{P}(\text{Over}_i \cap H_i) &\leq \mathcal{P}(m_i^{(i)} - m_i^{(0)} \leq -it/2) \\ &\leq \exp\left(-\frac{i^2 \cdot t^2}{8i(4t\sqrt[4]{i})^2}\right) = \exp\left(-\frac{\sqrt{i}}{128}\right). \end{aligned}$$

Altogether, we have

$$\begin{aligned} \mathcal{P}(H_i) &= \mathcal{P}(H_i \cap \text{Over}_i) + \mathcal{P}(H_i \cap \neg\text{Over}_i) \\ &\leq i \cdot a^{(t \cdot \sqrt[4]{i})/c} + e^{-\sqrt{i}/128} \leq i \cdot A^{\sqrt[4]{i}}, \end{aligned}$$

where $A = \max\{a^{t/c}, 2^{-1/128}\}$. Note that A is also computable in polynomial space.

We now have all the tools needed to prove Lemma 12. We have

$$\mathcal{P}(\inf_{k \in \mathbb{N}} \text{tot}_1(w; k) \leq -h) = \sum_{i \geq h} \mathcal{P}(H_i) \leq \sum_{i \geq h} i \cdot A^{\sqrt[4]{i}}.$$

Note that $\sum_{\ell=h}^\infty \ell \cdot A^{\sqrt[4]{\ell}} = \sum_{j=\lfloor \sqrt[4]{h} \rfloor}^\infty \sum_{\ell=j^4}^{(j+1)^4-1} \ell \cdot A^{\sqrt[4]{\ell}} \leq \sum_{j=\lfloor \sqrt[4]{h} \rfloor}^\infty \sum_{\ell=j^4}^{(j+1)^4-1} (j+1)^4 A^j \leq \sum_{j=\lfloor \sqrt[4]{h} \rfloor}^\infty 8(j+1)^7 A^j$. Using standard methods of calculus we can bound the last sum by $(c'' \cdot h^7 \cdot A^h)/(1-A)^8$ for some known constant c'' independent of \mathcal{B} . Thus, from the knowledge of A and c'' we can easily compute, again in polynomial space, numbers $h_0 \in \mathbb{N}$, $A_0 \in (0, 1)$ such that for all $h \geq h_0$ it holds

$$\mathcal{P}(\inf_{k \in \mathbb{N}} \text{tot}_1(w; k) \geq h) \geq 1 - A_0^h.$$

4. Conclusions

We have shown that the qualitative zero-reachability problem is decidable in Case I and II, and the probability of all zero-reaching runs can be effectively approximated. Let us note when the technical condition adopted in Case II is not satisfied, then the oc-trends may be undefined and the problem requires a completely different approach. An important technical contribution of this paper is the new martingale defined in Section 3.2, which provides a versatile tool for attacking other problems of pMC analysis (model-checking, expected termination time, constructing (sub)optimal strategies in multi-counter decision processes, etc.) similarly as the martingale of [13] for one-dimensional pMC.

References

- [1] P. A. Abdulla, N. Henda, and R. Mayr. Decisive Markov chains. *LMCS*, 3, 2007.
- [2] P. A. Abdulla, N. Henda, R. Mayr, and S. Sandberg. Limiting behavior of Markov chains with eager attractors. In *Proceedings of 3rd Int. Conf. on Quantitative Evaluation of Systems (QEST'06)*, pages 253–264. IEEE, 2006.
- [3] P. A. Abdulla, N. Bertrand, A. M. Rabinovich, and P. Schnoebelen. Verification of probabilistic systems with faulty communication. *Inf. Comput.*, 202(2):141–165, 2005.
- [4] P. A. Abdulla, L. Clemente, R. Mayr, and S. Sandberg. Stochastic parity games on lossy channel systems. In *Proceedings of 10th Int. Conf. on Quantitative Evaluation of Systems (QEST'13)*, volume 8054 of *LNCS*, pages 338–354. Springer, 2013.
- [5] M. Ajmone Marsan, G. Conte, and G. Balbo. A class of generalized stochastic Petri nets for the performance evaluation of multiprocessor systems. *ACM Trans. Comput. Syst.*, 2(2):93–122, 1984.
- [6] C. Baier and B. Engelen. Establishing qualitative properties for probabilistic lossy channel systems: an algorithmic approach. In *Proceedings of 5th International AMAST Workshop on Real-Time and Probabilistic Systems (ARTS'99)*, volume 1601 of *LNCS*, pages 34–52. Springer, 1999.
- [7] R. Bonnet, A. Finkel, J. Leroux, and M. Zeitoun. Model checking vector addition systems with one zero test. *LMCS*, 8(2), 2012.

- [8] L. Bozzelli and P. Ganty. Complexity analysis of the backward coverability algorithm for VASS. In *Reachability Problems*, volume 6945 of *LNCS*, pages 96–109. Springer, 2011.
- [9] T. Brázdil, V. Brožek, and K. Etessami. One-counter stochastic games. In *Proceedings of FST&TCS 2010*, volume 8 of *LIPICs*, pages 108–119. Schloss Dagstuhl, 2010.
- [10] T. Brázdil, V. Brožek, K. Etessami, and A. Kučera. Approximating the termination value of one-counter MDPs and stochastic games. In *Proceedings of ICALP 2011, Part II*, volume 6756 of *LNCS*, pages 332–343. Springer, 2011.
- [11] T. Brázdil, V. Brožek, K. Etessami, A. Kučera, and D. Wojtczak. One-counter Markov decision processes. In *Proceedings of SODA 2010*, pages 863–874. SIAM, 2010.
- [12] T. Brázdil, P. Jančar, and A. Kučera. Reachability games on extended vector addition systems with states. Technical report, 2010. Available at <http://arxiv.org/abs/1002.2557>.
- [13] T. Brázdil, S. Kiefer, and A. Kučera. Efficient analysis of probabilistic programs with an unbounded counter. In *Proceedings of CAV 2011*, volume 6806 of *LNCS*, pages 208–224. Springer, 2011.
- [14] T. Brázdil, S. Kiefer, and A. Kučera. Efficient analysis of probabilistic programs with an unbounded counter. Technical report, 2011. Available at <http://arxiv.org/abs/1102.2529>.
- [15] T. Brázdil, S. Kiefer, A. Kučera, and I. Hutařová Vařeková. Runtime analysis of probabilistic programs with unbounded recursion. In *Proceedings of ICALP 2011, Part II*, volume 6756 of *LNCS*, pages 319–331. Springer, 2011.
- [16] T. Brázdil, S. Kiefer, A. Kučera, P. Novotný, and J.-P. Katoen. Zero-Reachability in Probabilistic Multi-Counter Automata. Technical report, 2014. Available at <http://arxiv.org/abs/1401.6840>.
- [17] I. Erdelyi. On the matrix equation $Ax = \lambda Bx$. *Journal of Mathematical Analysis and Applications*, 17(1):119–132, 1967.
- [18] K. Etessami, A. Stewart, and M. Yannakakis. Polynomial time algorithms for multi-type branching processes and stochastic context-free grammars. In *Proceedings of STOC 2012*, pages 579–588. ACM Press, 2012.
- [19] K. Etessami, D. Wojtczak, and M. Yannakakis. Quasi-birth-death processes, tree-like QBDs, probabilistic 1-counter automata, and push-down systems. In *Proceedings of 5th Int. Conf. on Quantitative Evaluation of Systems (QEST'08)*. IEEE, 2008.
- [20] K. Etessami, D. Wojtczak, and M. Yannakakis. Quasi-birth-death processes, tree-like QBDs, probabilistic 1-counter automata, and push-down systems. *Performance Evaluation*, 67(9):837–857, 2010.
- [21] G. Florin and S. Natkin. One-place unbounded stochastic Petri nets: Ergodic criteria and steady-state solutions. *Journal of Systems and Software*, 6(1-2):103–115, 1986.
- [22] G. Florin and S. Natkin. Necessary and sufficient ergodicity condition for open synchronized queueing networks. *IEEE Trans. Software Eng.*, 15(4):367–380, 1989.
- [23] S. Iyer and M. Narasimha. Probabilistic lossy channel systems. In *Proceedings of TAPSOFT'97*, volume 1214 of *LNCS*, pages 667–681. Springer, 1997.
- [24] J. Kemeny and J. Snell. *Finite Markov chains*. D. Van Nostrand Company, 1960.
- [25] C. Meyer. The role of the group generalized inverse in the theory of finite Markov chains. *SIAM Review*, 17(3):443–464, 1975.
- [26] M. Minsky. *Computation: Finite and Infinite Machines*. Prentice-Hall, 1967.
- [27] M. K. Molloy. Performance analysis using stochastic Petri nets. *IEEE Trans. Computers*, 31(9):913–917, 1982.
- [28] J. Norris. *Markov Chains*. Cambridge University Press, 1998.
- [29] L. Rosier and H.-C. Yen. A multiparameter analysis of the boundedness problem for vector addition systems. *JCSS*, 32:105–135, 1986.