

Long-Run Average Behaviour of Probabilistic Vector Addition Systems

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
IST Austria
Klosterneuburg, Austria
petr.novotny@ist.ac.at

Abstract—We study the pattern frequency vector for runs in probabilistic Vector Addition Systems with States (pVASS). Intuitively, each configuration of a given pVASS is assigned one of finitely many *patterns*, and every run can thus be seen as an infinite sequence of these patterns. The pattern frequency vector assigns to each run the limit of pattern frequencies computed for longer and longer prefixes of the run. If the limit does not exist, then the vector is undefined. We show that for one-counter pVASS, the pattern frequency vector is defined and takes one of finitely many values for almost all runs. Further, these values and their associated probabilities can be approximated up to an arbitrarily small relative error in polynomial time. For stable two-counter pVASS, we show the same result, but we do not provide any upper complexity bound. As a byproduct of our study, we discover counterexamples falsifying some classical results about stochastic Petri nets published in the 80s.

I. INTRODUCTION

Stochastic extensions of Petri nets are intensively used in performance and dependability analysis as well as reliability engineering and bio-informatics. They have been developed in the early eighties [12], [2], and their token-game semantics yields a denumerable Markov chain. The analysis of stochastic Petri nets (SPNs) has primarily focused on long-run average behaviour. Whereas for safe nets long-run averages always exist and can be efficiently computed, the setting of infinite-state nets is much more challenging. This is a practically very relevant problem as, e.g., classical open queueing networks and biological processes typically yield nets with unbounded state space. The aim of this paper is to study the long-run average behaviour for infinite-state nets. We do so by considering probabilistic Vector Addition Systems with States (pVASS, for short), finite-state weighted automata equipped with a finite number of non-negative counters. A pVASS evolves by taking weighted rules along which any counter can be either incremented or decremented by one (or zero). The probability of performing a given enabled rule is given by its weight divided by the total weight of all enabled rules. This model is equivalent to discrete-time SPNs: a counter vector corresponds to the occupancy of the unbounded places in the net, and the bounded places are either encoded in the counters or in the control states. Producing a token yields an increment, whereas token consumption yields a decrement.

Discrete-time SPNs describe the probabilistic branching of the continuous-time Markov chains determined by SPNs, and many properties of continuous-time SPNs can be derived directly from the properties of their underlying discrete-time SPNs. In fact, discrete-time SPNs are a model of interest in itself, see e.g., [10].

Our study concentrates on long-run average *pattern frequencies* for pVASS. A *configuration* of a given pVASS \mathcal{A} is a pair $p\mathbf{v}$, where p is the current control state and $\mathbf{v} \in \mathbb{N}^d$ is the vector of current counter values. The *pattern* associated to $p\mathbf{v}$ is a pair $p\alpha$, where $\alpha \in \{0, *\}^d$, and α_i is either 0 or *, depending on whether v_i is zero or positive (for example, the pattern associated to $p(12, 0)$ is $p(*, 0)$). Every run in \mathcal{A} is an infinite sequence of configurations which determines a unique infinite sequence of the associated patterns. For every finite prefix of a run w , we can compute the frequency of each pattern in the prefix, and define the *pattern frequency vector* for w , denoted by $F_{\mathcal{A}}(w)$, as the limit of the sequence of frequencies computed for longer and longer prefixes of w . If the limit does not exist, we put $F_{\mathcal{A}}(w) = \perp$ and say that $F_{\mathcal{A}}$ is not well defined for w . Intuitively, a pattern represents the information sufficient to determine the set of enabled rules (recall that each rule can consume at most one token from each counter). Hence, if we know $F_{\mathcal{A}}(w)$, we can also determine the limit frequency of rules fired along w . However, we can also encode various predicates in the finite control of \mathcal{A} and determine the frequency of (or time proportion spent in) configurations in w satisfying the predicate. For example, we might wonder what is the proportion of time spent in configurations where the second counter is even, which can be encoded in the above indicated way.

The very basic questions about the pattern frequency vector include the following:

- Do we have $\mathcal{P}(F_{\mathcal{A}} = \perp) = 0$, i.e., is $F_{\mathcal{A}}$ well defined for almost all runs?
- Is $F_{\mathcal{A}}$ (seen as a random variable) discrete? If so, how many values can $F_{\mathcal{A}}$ take with positive probability?
- Can we somehow compute or approximate possible values of $F_{\mathcal{A}}$ and the probabilities of all runs that take these values?

These fundamental questions are rather difficult for general pVASS. In this paper, we concentrate on the subcase of pVASS with one or two counters, and we also observe that with three or more counters, there are some new unexpected phenomena that make the analysis even more challenging. Still, our results can be seen as a basis for designing algorithms that analyze the long-run average behaviour in certain subclasses of pVASS with arbitrarily many counters (see below). The main “algorithmic results” of our paper can be summarized as follows:

1. For a *one-counter* pVASS with n control states, we show that $F_{\mathcal{A}}$ is well defined and takes at most $\max\{2, 2n-1\}$ different values for almost all runs. These values and the associated probabilities may be irrational, but can be effectively approximated up to an arbitrarily small relative error $\varepsilon > 0$ in *polynomial time*.

2. For *two-counter* pVASS that are *stable*, we show that $F_{\mathcal{A}}$ is well defined and takes only finitely many values for almost all runs. Further, these values and the associated probabilities can be effectively approximated up to an arbitrarily small absolute/relative error $\varepsilon > 0$.

Intuitively, a two-counter pVASS \mathcal{A} is *unstable* if the changes of the counters are well-balanced so that certain infinite-state Markov chains used to analyze the behaviour of \mathcal{A} may become *null-recurrent*. Except for some degenerated cases, this null-recurrence is not preserved under small perturbations in transition probabilities. Hence, we can assume that pVASS models constructed by estimating some real-life probabilities are stable. Further, the analysis of null-recurrent Markov chains requires different methods and represents an almost independent task. Therefore, we decided to disregard unstable two-counter pVASS in this paper. Let us note that the problem whether a given two-counter pVASS \mathcal{A} is (un)stable is decidable in exponential time.

The above results for one-counter and stable two-counter pVASS are obtained by showing the following:

(a) There are finitely many sets of configurations called *regions*, such that almost every run eventually stays in some region, and almost all runs that stay in the same region share the same well-defined value of the pattern frequency vector.

(b) For every region R , the associated pattern frequency vector and the probability of reaching R can be computed/approximated effectively. For one-counter pVASS, we first identify families of regions (called *zones*) that share the same pattern frequency vector, and then consider these zones rather than individual regions.

For one-counter pVASS, we show that the total number of all regions (and hence also zones) cannot exceed $\max\{2, 2n-1\}$, where n is the number of control states. To compute/approximate the pattern frequency vector of a given zone Z and the probability of staying in Z , the tail bounds of [4] and the polynomial-time algorithm of [14] provide all the tools we need.

For two-counter pVASS, we do not give an explicit bound on the number of regions, but we show that all regions

are effectively semilinear (i.e., for each region there is a computable Presburger formula which represents the region). Here we repeatedly use the result of [11] which says that the reachability relation of a two-counter VASS is effectively semilinear. Technically, we show that every run eventually reaches a configuration where one or both counters become bounded or irrelevant (and we apply the results for one-counter pVASS), or a configuration of a special set C for which we show the existence and effective constructibility of a finite eager attractor¹. This is perhaps the most advanced part of our paper, where we need to establish new exponential tail bounds for certain random variables using an appropriately defined martingale. We believe that these tail bounds and the associated martingale are of broader interest, because they provide generic and powerful tools for quantitative analysis of two-counter pVASS. Hence, every run which visits C also visits its finite eager attractor, and the regions where the runs initiated in C eventually stay correspond to bottom strongly connected components of this attractor. For each of these bottom strongly connected components, we approximate the pattern frequency vector by employing the abstract algorithm of [1].

The overall complexity of our algorithm for stable two-counter pVASS could be estimated by developing lower/upper bounds on the parameters that are used in the lemmata of Section IV. Many of these parameters are “structural” (e.g., we consider the minimal length of a path from some configuration to some set of configurations). Here we miss a refinement of the results published in [11] which would provide explicit upper bounds. Another difficulty is that we do not have any lower bound on $|\tau_R|$ in the case when $\tau_R \neq 0$, where τ_R is the mean payoff defined in Section IV. Still, we conjecture that these “structural bounds” and hence also the complexity of our algorithm are not too high (perhaps, singly exponential in the size of \mathcal{A} and in $|\tau_R|$), but we leave this problem for future work.

The results summarized in (a) and (b) give a reasonably deep understanding of the long-run behaviour of a given one-counter or a stable two-counter pVASS, which can be used to develop algorithms for other interesting problems. For example, we can decide the existence of a finite attractor for the set of configuration reachable from a given initial configuration, we can provide a sufficient condition which guarantees that all pattern frequency vectors taken with positive probability are rational, etc. An obvious question is whether these results can be extended to pVASS with three or more counters. The answer is twofold.

I. The algorithm for stable two-dimensional pVASS presented in Section IV in fact “reduces” the analysis of a given two-counter pVASS \mathcal{A} to the analysis of several one-counter pVASS and the analysis of some “special” configurations of \mathcal{A} . It seems that this approach can be generalized to a recursive

¹A finite eager attractor [1] for a set of configurations C is a finite set of configurations $A \subseteq C$ such that the probability of reaching A from every configuration of $C \cup \text{post}^*(A)$ is equal to 1, and the probability of revisiting A in more than ℓ steps after leaving A decays (sub)exponentially in ℓ .

procedure which takes a pVASS \mathcal{A} with n counters, isolates certain subsets of runs whose properties can be deduced by analyzing pVASS with smaller number of counters, and checks that the remaining runs are sufficiently simple so that they can be analyzed directly. Thus, we would obtain a procedure for analyzing a subset of pVASS with n counters.

II. In Section V we give an example of a *three-counter* pVASS \mathcal{A} with strongly connected state-space whose long-run behaviour is *undefined* for almost all runs (i.e., $F_{\mathcal{A}}$ takes the \perp value), and this property is not sensitive to small perturbations in transition probabilities. Since we do not provide a rigorous mathematical analysis of \mathcal{A} in this paper, the above claims are formally just *conjectures* confirmed only by Monte Carlo simulations. Assuming that these conjectures are valid, the method used for two-counter pVASS is not sufficient for the analysis of general three-counter pVASS, i.e., there are new phenomena which cannot be identified by the methods used for two-counter pVASS.

Related work. The problem of studying pattern frequency vector is directly related to the study of ergodicity properties in stochastic Petri nets, particularly to the study of the so-called *firing process*. A classical paper in this area [8] has been written by Florin & Natkin in the 80s. In the paper, it is claimed that if the state-space of a given stochastic Petri net (with arbitrarily many unbounded places) is strongly connected, then the firing process is ergodic. In the setting of (discrete-time) probabilistic Petri nets, this implies that for almost all runs, the limit frequency of transitions performed along a run is defined and takes the same value. A simple counterexample to this claim is shown in Fig. 1. The net \mathcal{N} has two unbounded places and strongly connected state-space, but the limit frequency of transitions takes two values with positive probability (each with probability 1/2). Note that \mathcal{N} can be translated into an equivalent pVASS \mathcal{A} with two counters which is also shown in Fig. 1. Intuitively, if both places/counters are positive, then both of them have a tendency to decrease, i.e., the trend t_S of the only BSCC S of $\mathcal{C}_{\mathcal{A}}$ is negative in both components (see Section II). However, if we reach a configuration where the first place/counter is zero and the second place/counter is sufficiently large, then the second place/counter starts to *increase*, i.e., it never becomes zero again with some positive probability (i.e., the the mean payoff τ_{R_2} is positive, where R_2 is the only type II region of the one-counter pVASS \mathcal{A}_2 , see Section IV). The first place/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 place/counter is large and the second place/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 limit frequencies of transitions, each of which is taken with probability 1/2. This possibility




Fig. 1: A discrete-time SPN \mathcal{N} and an equivalent pVASS \mathcal{A} .

of reversing the “global” trend of the counters after hitting zero in some counter was not considered in [8]. Further, as we already mentioned, we conjecture the existence of a three-counter pVASS \mathcal{A} with strongly connected state-space (the one of Section V) where the limit frequency of transitions is undefined for almost all runs. So, we must unfortunately conclude that the results of [8] are invalid for fundamental reasons. On the other hand, the results achieved for one-counter pVASS are consistent with another paper by Florin & Natkin [7] devoted to stochastic Petri nets with only one unbounded place and strongly connected state-space, where the firing process is indeed ergodic (in our terms, the pattern frequency vector takes only one value with probability 1).

II. 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. The absolute value of a given $x \in \mathbb{R}$ is denoted by $|x|$. Let $\delta > 0$, $x \in \mathbb{Q}$, and $y \in \mathbb{R}$. We say that x approximates y up to a relative error δ , if either $y \neq 0$ and $|x - y|/|y| \leq \delta$, or $x = y = 0$. Further, we say that x approximates y up to an absolute error δ if $|x - y| \leq \delta$. We assume that rational numbers (including integers) are represented as fractions of binary numbers, and we use $\|x\|$ to denote the size (length) of this representation.

Let $\mathcal{V} = (V, \rightarrow)$, where V is a non-empty 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 $v \rightarrow u$). The reflexive and transitive closure of \rightarrow is denoted by \rightarrow^* , and the reflexive, symmetric and transitive closure of \rightarrow is denoted by \leftrightarrow^* . We say that \mathcal{V} is *weakly connected* if $s \leftrightarrow^* t$ for all $s, t \in V$. A *finite path* in \mathcal{V} of length $k \geq 0$ is a finite sequence of vertices v_0, \dots, 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 $\text{length}(w)$. 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 $w(0), w(1), \dots$. 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 $Run_{\mathcal{V}}(w)$ (or just by $FPath(w)$ and $Run(w)$ if \mathcal{V} is understood), respectively. For a given set $S \subseteq V$, we use $pre^*(S)$ and $post^*(S)$ to denote the set of all $v \in V$ such that $v \rightarrow^* s$ and $s \rightarrow^* v$ for some $s \in S$, respectively. Further, we say that a run w *stays* in S if there is a $k \in \mathbb{N}$ such that for

all $\ell \geq k$ we have that $w(\ell) \in S$. The set of all runs initiated in s that stay in S is denoted by $Run(s, S)$.

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*. Given events E, F , we say that E holds for almost all elements of F if $\mathcal{P}(E \cap F) = \mathcal{P}(F)$ (in particular, if $\mathcal{P}(F) = 0$, then any event holds for almost all elements of F). 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 Markov chain is a triple $\mathcal{M} = (S, \rightarrow, Prob)$ where S is a finite or countably infinite set of vertices, $\rightarrow \subseteq S \times S$ 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 $s \rightarrow t$.

To every $s \in S$ we associate the standard probability space $(Run_{\mathcal{M}}(s), \mathcal{F}, \mathcal{P})$ of runs starting at s , where \mathcal{F} is the σ -field generated by all *basic cylinders* $Run_{\mathcal{M}}(w)$, 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 $\mathcal{P}(Run_{\mathcal{M}}(w)) = 1$.

If $\mathcal{M} = (S, \rightarrow, Prob)$ is a strongly connected finite-state Markov chain, we use μ_S to denote the unique *invariant distribution* of \mathcal{M} . Recall that by the strong ergodic theorem, (see, e.g., [13], the limit frequency of visits to the states of S is defined for almost all $w \in Run(s)$ (where $s \in S$ is some initial state) and it is equal to μ_S .

Definition 2. A probabilistic vector addition system with states (pVASS) of dimension $d \geq 1$ is a triple $\mathcal{A} = (Q, \gamma, W)$, where Q is a finite set of control states, $\gamma \subseteq Q \times \{-1, 0, 1\}^d \times Q$ is a set of rules, and $W : \gamma \rightarrow \mathbb{N}^+$ is a weight assignment.

In the following, we often write $p \xrightarrow{\kappa} q$ to denote that $(p, \kappa, q) \in \gamma$, and $p \xrightarrow{\kappa, \ell} q$ to denote that $(p, \kappa, q) \in \gamma$ and $W((p, \kappa, q)) = \ell$. The encoding size of \mathcal{A} is denoted by $\|\mathcal{A}\|$, where the weights are encoded in binary.

Assumption 1. From now on (in the whole paper), we assume that (Q, \rightarrow) , where $p \rightarrow q$ iff $p \xrightarrow{\kappa} q$ for some q , is weakly connected. Further, we also assume that for every pair of control states p, q there is at most one rule of the form $p \xrightarrow{\kappa} q$.

The first condition of Assumption 1 is obviously safe (if (Q, \rightarrow) is not weakly connected, then \mathcal{A} is a “disjoint union” of several independent pVASS, and we can apply our results to each of them separately). The second condition is also safe because every pVASS \mathcal{A} can be easily transformed into another pVASS \mathcal{A}' satisfying this condition in the following way: for each control state s of \mathcal{A} and each rule of the form $r \xrightarrow{\kappa} s$

we add a fresh control state $s[r, \kappa]$ to \mathcal{A}' . Further, for every $s \xrightarrow{\kappa, \ell} t$ in \mathcal{A} we add $s[r, \kappa'] \xrightarrow{\kappa, \ell} t[s, \kappa]$ to \mathcal{A}' (for all states of the form $s[r, \kappa']$ in \mathcal{A}'). In other words, \mathcal{A}' is the same as \mathcal{A} , but it also “remembers” the rule that was used to enter a given control state.

A *configuration* of \mathcal{A} is an element of $conf(\mathcal{A}) = Q \times \mathbb{N}^d$, written as $p\mathbf{v}$. A rule $p \xrightarrow{\kappa} q$ is *enabled* in a configuration $p\mathbf{v}$ if $v_i > 0$ for all $1 \leq i \leq d$ with $\kappa_i = -1$. To \mathcal{A} we associate an infinite-state Markov chain $\mathcal{M}_{\mathcal{A}}$ whose vertices 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 $p\mathbf{v} \xrightarrow{1} p\mathbf{v}$ is the only outgoing transition of $p\mathbf{v}$;
- otherwise, for every rule $p \xrightarrow{\kappa, \ell} q$ enabled in $p\mathbf{v}$ there is a transition $p\mathbf{v} \xrightarrow{\ell/T} q(\mathbf{v} + \kappa)$ where T is the total weight of all rules enabled in $p\mathbf{v}$, and there are no other outgoing transitions of $p\mathbf{v}$.

In this paper, we also consider the underlying finite-state Markov chain of \mathcal{A} , denoted by $\mathcal{C}_{\mathcal{A}}$, whose vertices are the control states of \mathcal{A} , and $p \xrightarrow{x} q$ in $\mathcal{C}_{\mathcal{A}}$ iff $p \xrightarrow{\kappa, \ell} q$ in \mathcal{A} and $x = \ell/T_p > 0$, where T_p is the sum of the weights of all outgoing rules of p in \mathcal{A} . Note that every BSCC S of $\mathcal{C}_{\mathcal{A}}$ can be seen as a strongly connected finite-state Markov chain, and we use μ_S to denote the invariant distribution on the states of S . To each $s \in S$ we associate the vector

$$change(s) = \sum_{(s, \kappa, t) \in \gamma} \kappa \cdot \frac{W((s, \kappa, t))}{T_s}$$

of expected changes in counter values at s . Further, we define the *trend* of S , denoted by t_S , as the vector $t_S = \sum_{s \in S} \mu_S(s) \cdot change(s)$.

A *pattern* of \mathcal{A} is a pair $q\alpha \in Q \times \{0, *\}^d$, and the set of all patterns of \mathcal{A} is denoted by $Pat_{\mathcal{A}}$. A configuration $p\mathbf{v}$ matches a pattern $q\alpha \in Pat_{\mathcal{A}}$ if $p = q$ and for every $i \in \{1, \dots, d\}$ we have that $v_i = 0$ or $v_i > 0$, depending on whether $\alpha_i = 0$ or $\alpha_i = *$, respectively. Intuitively, a pattern represents exactly the information which determines the set of enabled rules. For all $w \in Run_{\mathcal{M}_{\mathcal{A}}}(p\mathbf{v})$, we define the *pattern frequency vector* $F_{\mathcal{A}}(w) : Pat_{\mathcal{A}} \rightarrow \mathbb{R}$ as follows:

$$F_{\mathcal{A}}(w)(q\alpha) = \lim_{k \rightarrow \infty} \frac{\#_{q\alpha}(w(0), \dots, w(k))}{k+1}$$

where $\#_{q\alpha}(w(0), \dots, w(k))$ denotes the total number of all indexes i such that $0 \leq i \leq k$ and $w(i)$ matches the pattern $q\alpha$. If the above limit does not exist for some $q\alpha \in Pat_{\mathcal{A}}$, we put $F_{\mathcal{A}}(w) = \perp$. We say that $F_{\mathcal{A}}$ is *well defined* for w if $F_{\mathcal{A}}(w) \neq \perp$. Note that if $F_{\mathcal{A}}$ is well defined for w , then the sum of all components of $F_{\mathcal{A}}(w)$ is equal to 1.

Let $R \subseteq Run(p\mathbf{v})$ be a measurable subset of runs, and let $\varepsilon > 0$. We say that a sequence $(H_1, P_1), \dots, (H_n, P_n)$, where $H_i : Pat_{\mathcal{A}} \rightarrow \mathbb{Q}$ and $P_i \in \mathbb{Q}$, approximates the pattern frequencies of R up to the absolute/relative error ε , if there are pairwise disjoint measurable subsets R_1, \dots, R_n of R and vectors F_1, \dots, F_n , where $F_i : Pat_{\mathcal{A}} \rightarrow \mathbb{R}$, such that

- $\sum_{i=1}^n \mathcal{P}(R_i) = \mathcal{P}(R)$;

- $F_{\mathcal{A}}(w) = F_i$ for almost all $w \in R_i$;
- $H_i(q\alpha)$ approximates $F_i(q\alpha)$ up to the absolute/relative error ε for every $q\alpha \in \text{Pat}_{\mathcal{A}}$;
- P_i approximates $\mathcal{P}(R_i)$ up to the absolute/relative error ε .

Note that if $(H_1, P_1), \dots, (H_n, P_n)$ approximates the pattern frequencies of R up to some absolute/relative error, then the pattern frequency vector is well defined for almost all $w \in R$ and takes only finitely many values with positive probability. Also note that neither F_1, \dots, F_n nor H_1, \dots, H_n are required to be pairwise different. Hence, it may happen that there exist $i \neq j$ such that $H_i \neq H_j$ and $F_i = F_j$ (or $H_i = H_j$ and $F_i \neq F_j$).

III. RESULTS FOR ONE-COUNTER pVASS

In this section we concentrate on analyzing the pattern frequency vector for one-dimensional pVASS. We show that $F_{\mathcal{A}}$ is well defined and takes at most $|Q| + b$ distinct values for almost all runs, where $|Q|$ is the number of control states of \mathcal{A} , and b is the number of BSCCs of $\mathcal{C}_{\mathcal{A}}$. Moreover, these values as well as the associated probabilities can be efficiently approximated up to an arbitrarily small positive relative error. More precisely, our aim is to prove the following:

Theorem 1. *Let $\mathcal{A} = (Q, \gamma, W)$ be a one-dimensional pVASS, and let b be the number of BSCCs of $\mathcal{C}_{\mathcal{A}}$. Then there is $n \leq |Q| + b$ computable in time polynomial in $\|\mathcal{A}\|$ such that for every $\varepsilon > 0$, there are $H_1, \dots, H_n : \text{Pat}_{\mathcal{A}} \rightarrow \mathbb{Q}$ computable in time polynomial in $\|\mathcal{A}\|$ and $\|\varepsilon\|$, such that for every initial configuration $p(k) \in \text{conf}(\mathcal{A})$ there are $P_1, \dots, P_n \in \mathbb{Q}$ computable in time polynomial in $\|\mathcal{A}\|$, $\|\varepsilon\|$, and k , such that the sequence $(P_1, H_1), \dots, (P_n, H_n)$ approximates the pattern frequencies of $\text{Run}(p(k))$ up to the relative error ε .*

Let us note that the “real” pattern frequency vectors F_i as well as the probabilities $\mathcal{P}(F_{\mathcal{A}}=F_i)$ may take irrational values, and they cannot be computed precisely in general.

Remark 1. *The $|Q| + b$ upper bound on n given in Theorem 1 is tight. To see this, realize that if $|Q| = 1$, then $b = 1$ and the trivial pVASS with the only rule $p \xrightarrow{0} p$ witnesses that the pattern frequency vector may take two different values. If $|Q| \geq 2$, we have that $b \leq |Q| - 1$. Consider a pVASS where $Q = \{p, q_1, \dots, q_k\}$ and γ contains the rules $p \xrightarrow{-1} p$, $p \xrightarrow{-1} q_i$, and $q_i \xrightarrow{0} q_i$ for all $1 \leq i \leq k$, where all of these rules have the same weight equal to 1. For $p(2)$ as the initial configuration, the vector $F_{\mathcal{A}}$ takes $2k + 1 = 2|Q| - 1$ pairwise different values with positive probability.*

For the rest of this section, we fix a one-dimensional pVASS $\mathcal{A} = (Q, \gamma, W)$. We start by identifying certain (possibly empty) subsets of configurations called *regions* that satisfy the following properties:

- there are at most $|Q| + b$ non-empty regions;
- almost every run eventually stays in precisely one region;
- almost all runs that stay in a given region have the same well defined pattern frequency vector.

In principle, we might proceed by considering each region R separately and computing/approximating the associated pattern

frequency vector and the probability of all runs that stay in R . However, this would lead to unnecessary technical complications. Instead, we identify situations when multiple regions share the *same* pattern frequency vector, consider unions of such regions (called *zones*), and then compute/approximate the pattern frequency vector and the probability of staying in Z for each zone Z . Thus, we obtain Theorem 1.

Technically, we distinguish among four *types* of regions determined either by a control state of \mathcal{A} or a BSCC of $\mathcal{C}_{\mathcal{A}}$.

- Let $p \in Q$. A *type I region determined by p* is either the set $\text{post}^*(p(0))$ or the empty set, depending on whether $\text{post}^*(p(0))$ is a finite set satisfying $\text{post}^*(p(0)) \subseteq \text{pre}^*(p(0))$ or not, respectively.
- Let $p \in S$, where S is a BSCC of $\mathcal{C}_{\mathcal{A}}$. A *type II region determined by p* is either the set $\text{post}^*(p(0))$ or the empty set, depending on whether $\text{post}^*(p(0))$ is an infinite set satisfying $\text{post}^*(p(0)) \subseteq \text{pre}^*(p(0))$ or not, respectively.
- Let S be a BSCC of $\mathcal{C}_{\mathcal{A}}$. A *type III region determined by S* consists of all $p(k) \in S \times \mathbb{N}^+$ that cannot reach a configuration with zero counter.
- Let S be a BSCC of $\mathcal{C}_{\mathcal{A}}$, and let $R_I(S)$ and $R_{II}(S)$ be the unions of all type I and all type II regions determined by the control states of S , respectively. Further, let $D(S)$ be the set

$$\left(S \times \mathbb{N} \cap \text{pre}^*(R_I(S)) \right) \setminus \left(R_I(S) \cup \text{pre}^*(R_{II}(S)) \right)$$

A *type IV region determined by S* is either the set $D(S)$ or the empty set, depending on whether $D(S)$ is infinite or finite, respectively.

Note that if R_1, R_2 are regions of \mathcal{A} such that $R_1 \cap R_2 \neq \emptyset$, then $R_1 = R_2$. Also observe that regions of type I, II, and III are closed under post^* , and each such region can thus be seen as a Markov chain. Finally, note that every configuration of a type IV region can reach a configuration of a type I region, and the size of every type I region is bounded by $|Q|^2$ (if $R = \text{post}^*(p(0))$ is a type I region and $p(0) \xrightarrow{*} q(j)$, then $j < |Q|$, because otherwise the counter could be pumped to an arbitrarily large value; hence, $|R| \leq |Q|^2$).

Let us note that all regions are *regular* in the following sense: We say that a set $C \subseteq \text{conf}(\mathcal{A})$ of configurations is *regular* if there is a non-deterministic finite automaton A over the alphabet $\{a\}$ such that the set of control states of A subsumes Q and for every configuration $p(k) \in \text{conf}(\mathcal{A})$ we have that $p(k) \in C$ iff the word a^k is accepted by A with p as the initial state. It follows, e.g., from the results of [6] that if $C \subseteq \text{conf}(\mathcal{A})$ is regular, then $\text{post}^*(C)$ and $\text{pre}^*(C)$ are also regular and the associated NFA are computable in time polynomial in $\|\mathcal{A}\|$, where A is the NFA representing C . Hence, all regions are effectively regular which becomes important in Section IV.

Let S be a SCC of $\mathcal{C}_{\mathcal{A}}$. If S is not a BSCC of $\mathcal{C}_{\mathcal{A}}$, then the control states of S may determine at most $|S|$ non-empty regions (of type I). If S is a BSCC of $\mathcal{C}_{\mathcal{A}}$, then the control states of S may determine at most $|S|$ non-empty regions of type I or II, and at most one additional non-empty region

which is either of type III or of type IV (clearly, it cannot happen that the type III and type IV regions determined by S are both non-empty). Hence, the total number of non-empty regions cannot exceed $|Q|+b$, where b is the number of BSCCs of \mathcal{C}_A (here we also use the assumption that \mathcal{C}_A is weakly connected).

Now we prove that every configuration can reach some region in a bounded number of steps. This fact is particularly important for the analysis of two-counter pVASS in Section IV.

Lemma 1. *Every configuration of \mathcal{A} can reach a configuration of some region in at most $11|Q|^4$ transitions.*

By Lemma 1, the probability of reaching (some) region from an arbitrary initial configuration is at least $x_{\min}^{11|Q|^4}$, where x_{\min} is the least positive transition probability of \mathcal{M}_A . This implies that almost every $w \in \text{Run}(p(k))$ visits some region R . If R is of type I, II, or III, then w inevitably stays in R because these regions are closed under post^* . If R is a type IV region, then w either stays in R , or later visits a configuration of a type I region where it stays. Thus, we obtain the following:

Lemma 2. *Let $p(k)$ be a configuration of \mathcal{A} . Then almost every run initiated in $p(k)$ eventually stays in precisely one region.*

As we already mentioned, computing the pattern frequency vector and the probability of staying in R for each region R separately is technically complicated. Therefore, we also introduce *zones*, which are unions of regions that are guaranteed to share the same pattern frequency vector. Formally, a *zone* of \mathcal{A} is a set $Z \subseteq \text{conf}(\mathcal{A})$ satisfying one of the following conditions (recall that t_S denotes the trend of a BSCC S):

- $Z = R$, where R is a region of type I.
- $Z = R$, where R is a type III region determined by a BSCC S of \mathcal{C}_A such that $t_S \leq 0$.
- $Z = R$, where R is a type II region determined by $p \in S$ where S is a BSCC of \mathcal{C}_A satisfying $t_S < 0$.
- $Z = R_{II}(S)$, where S is a BSCC of \mathcal{C}_A such that $t_S = 0$ and $R_{II}(S)$ is the union of all type II regions determined by the control states of S .
- $Z = R_{II}(S) \cup R_{III}(S) \cup R_{IV}(S)$, where S is a BSCC of \mathcal{C}_A such that $t_S > 0$, $R_{II}(S)$ is the union of all type II regions determined by the control states of S , and $R_{III}(S)$ and $R_{IV}(S)$ are the type III and the type IV regions determined by S , respectively.

The next two lemmata are nontrivial and represent the technical core of this section (proofs can be found in [5]). They crucially depend on the results presented recently in [4] and [14]. In the proof of Lemma 3, we also characterize situations when some elements of pattern frequency vectors take irrational values.

Lemma 3. *Let $p(k)$ be a configuration of \mathcal{A} and Z a zone of \mathcal{A} . Then F_A is well defined for almost all $w \in \text{Run}(p(k), Z)$, and there exists $F : \text{Pat}_A \rightarrow \mathbb{R}$ such that $F_A(w) = F$ for almost all $w \in \text{Run}(p(k), Z)$. Further, for*

every rational $\varepsilon > 0$, there is a vector $H : \text{Pat}_A \rightarrow \mathbb{Q}$ computable in time polynomial in $\|\mathcal{A}\|$ and $\|\varepsilon\|$ such that $H(q\alpha)$ approximates $F(q\alpha)$ up to the relative error ε for every $q\alpha \in \text{Pat}_A$.

Lemma 4. *Let $p(k)$ be a configuration of \mathcal{A} . Then almost every run initiated in $p(k)$ eventually stays in precisely one zone of \mathcal{A} . Further, for every zone Z and every rational $\varepsilon > 0$, there is a $P \in \mathbb{Q}$ computable in time polynomial in $\|\mathcal{A}\|$, $\|\varepsilon\|$, and k such that P approximates $\mathcal{P}(\text{Run}(p(k), Z))$ up to the relative error ε .*

IV. RESULTS FOR TWO-COUNTER pVASS

In this section we analyze the long-run average behavior of two-counter pVASS. We show that if a given two-counter pVASS is *stable* (see Definition 5 below), then the pattern frequency vector is well defined takes one of finitely many values for almost all runs. Further, these values and the associated probabilities can be effectively approximated up to an arbitrarily small positive absolute/relative error.

Let \mathcal{A} be a two-counter pVASS. When we say that some object (e.g., a number or a vector) is *computable* for every $\sigma \in \Sigma$, where Σ is some set of parameters, we mean that there exists an algorithm which inputs the encodings of \mathcal{A} and σ , and outputs the object. Typically, the parameter σ is some rational $\varepsilon > 0$, of a pair $(\varepsilon, p\mathbf{v})$ where $p\mathbf{v}$ is a configuration. The parameter can also be void, which means that the algorithm inputs just the encoding of \mathcal{A} .

A *semilinear constraint* φ is a function $\varphi : Q \rightarrow \Phi$, where Φ is the set of all formulae of Presburger arithmetic with two free variables x, y . Each φ determines a semilinear set $\llbracket \varphi \rrbracket \subseteq \text{conf}(\mathcal{A})$ consisting of all $p(v_1, v_2)$ such that $\varphi(p)[x/v_1, y/v_2]$ is a valid formula. Since the reachability relation \rightarrow^* of \mathcal{A} is effectively semilinear [11] and semilinear sets are closed under complement and union, all of the sets of configurations we work with (such as $C[R_1, R_2]$ defined below) are effectively semilinear, i.e., the associated semilinear constraint is computable. In particular, the membership problem for these sets is decidable.

Given $p\mathbf{v} \in \text{conf}(\mathcal{A})$ and $D \subseteq \text{conf}(\mathcal{A})$, we use $\text{Run}(p\mathbf{v} \rightarrow^* D)$ to denote the set of all $w \in \text{Run}(p\mathbf{v})$ that visit a configuration of D , and $\text{Run}(p\mathbf{v} \not\rightarrow^* D)$ to denote the set $\text{Run}(p\mathbf{v}) \setminus \text{Run}(p\mathbf{v} \rightarrow^* D)$. Note that if $D = \emptyset$, then $\text{Run}(p\mathbf{v} \not\rightarrow^* D) = \text{Run}(p\mathbf{v})$.

Intuitively, our aim is to prove that the set $C = \text{conf}(\mathcal{A})$ is “good” in the sense that there is a computable $n \in \mathbb{N}$ such that for every rational $\varepsilon > 0$, there exists a computable sequence of rational vectors H_1, \dots, H_n such that for every $p\mathbf{v} \in C$, there are computable rational P_1, \dots, P_n such that the sequence $(P_1, H_1), \dots, (P_n, H_n)$ that approximates the pattern frequencies of $\text{Run}(p\mathbf{v})$ up to the absolute/relative error ε . This is achieved by first showing that certain simple subsets of configurations are good, and then (repeatedly) demonstrating that more complicated subsets are also good because they can be “reduced” to simpler subsets that are already known to be good. Thus, we eventually prove that the whole set $\text{conf}(\mathcal{A})$ is good.

For our purposes, it is convenient to parameterize the notion of a “good” subset C by another subset of “dangerous” configurations D so that the above conditions are required to hold only for those runs that do not visit D . Further, we require that every configuration of C can avoid visiting D with some positive probability which is bounded away from zero.

Definition 3. Let $\mathcal{A} = (Q, \gamma, W)$ be a pVASS of dimension two, and let $C, D \subseteq \text{conf}(\mathcal{A})$. We say that C is good for D if the following conditions are satisfied:

- There is $\delta > 0$ such that $\mathcal{P}(\text{Run}(p\mathbf{v} \rightarrow^* D)) \leq 1 - \delta$ for every $p\mathbf{v} \in C$.
- There is a computable $n \in \mathbb{N}$ such that for every $\varepsilon > 0$, there are computable $H_1, \dots, H_n : \text{Pat}_{\mathcal{A}} \rightarrow \mathbb{Q}$ such that for every $p\mathbf{v} \in C$ there are computable $P_{p\mathbf{v},1}, \dots, P_{p\mathbf{v},n} \in \mathbb{Q}$ such that $(P_{p\mathbf{v},1}, H_1), \dots, (P_{p\mathbf{v},n}, H_n)$ approximate the pattern frequencies of $\text{Run}(p\mathbf{v} \not\rightarrow^* D)$ up to the absolute error ε .

Note that in Definition 3, we require that $(P_{p\mathbf{v},1}, H_1), \dots, (P_{p\mathbf{v},n}, H_n)$ approximate the pattern frequencies of $\text{Run}(p\mathbf{v} \not\rightarrow^* D)$ up to the absolute error ε . As we shall see, we can always compute a lower bound for each positive $P_{p\mathbf{v},i}$ and H_i , which implies that if $P_{p\mathbf{v},i}$ and H_i can be effectively approximated up to an arbitrarily small absolute error $\varepsilon > 0$, they can also be effectively approximated up to an arbitrarily small relative error $\varepsilon > 0$.

The next definition and lemma explain what we mean by reducing the analysis of runs initiated in configurations of C to the analysis of runs initiated in “simpler” configurations of C_1, \dots, C_k .

Definition 4. Let \mathcal{A} be a pVASS of dimension two, $C \subseteq \text{conf}(\mathcal{A})$, and $\mathcal{E} = \{C_1, \dots, C_k\}$ a set of pairwise disjoint subsets of $\text{conf}(\mathcal{A})$. We say that C is reducible to \mathcal{E} if, for every $\varepsilon > 0$, there are computable semilinear constraints $\varphi_1, \dots, \varphi_k$ such that

- $\llbracket \varphi_i \rrbracket \subseteq C_i$ for every $1 \leq i \leq k$;
- for all $1 \leq i \leq k$ and $p\mathbf{v} \in \llbracket \varphi_i \rrbracket$, we have that $\mathcal{P}(\text{Run}(p\mathbf{v} \rightarrow^* D_i)) \leq \varepsilon$, where $D_i = \bigcup_{j \neq i} C_j$.
- for every $p\mathbf{v} \in C$ and every $\delta > 0$, there is a computable $\ell \in \mathbb{N}$ such that the probability of reaching a configuration of $\llbracket \varphi_1 \rrbracket \cup \dots \cup \llbracket \varphi_k \rrbracket$ in at most ℓ transitions is at least $1 - \delta$.

Lemma 5. If C is reducible to $\mathcal{E} = \{C_1, \dots, C_k\}$ and every C_i is good for $D_i = \bigcup_{j \neq i} C_j$, then C is good for \emptyset .

Proof: For every $1 \leq i \leq k$, let n_i be the computable constant for C_i which exists by Definition 3. The constant n for C is defined as $n = \sum_{i=1}^k n_i$. Now let us fix some $\varepsilon > 0$. Since C is reducible to $\{C_1, \dots, C_k\}$, there are computable constraints $\varphi_1, \dots, \varphi_k$ such that, for every $1 \leq i \leq k$, we have that $\llbracket \varphi_i \rrbracket \subseteq C_i$ and $\mathcal{P}(\text{Run}(p_i\mathbf{v}_i \rightarrow^* D_i)) \leq \varepsilon/4$ for every $p_i\mathbf{v}_i \in \llbracket \varphi_i \rrbracket$. Further, there are computable $H_{i,1}, \dots, H_{i,n_i} : \text{Pat}_{\mathcal{A}} \rightarrow \mathbb{Q}$ such that for every $p_i\mathbf{v}_i \in \llbracket \varphi_i \rrbracket$, there are computable $P_{p_i\mathbf{v}_i,1}, \dots, P_{p_i\mathbf{v}_i,n_i} \in \mathbb{Q}$ such that

$(P_{p_i\mathbf{v}_i,1}, H_{i,1}), \dots, (P_{p_i\mathbf{v}_i,n_i}, H_{i,n_i})$ approximate the pattern frequencies of $\text{Run}(p_i\mathbf{v}_i \not\rightarrow^* D_i)$ up to the absolute error $\varepsilon/4$. Now let $p\mathbf{v} \in C$. Then there is a computable $\ell \in \mathbb{N}$ such that the probability of reaching a configuration of $\llbracket \varphi_1 \rrbracket \cup \dots \cup \llbracket \varphi_k \rrbracket$ in at most ℓ transitions is at least $1 - \varepsilon/4$. Hence, we can effectively construct a finite tree T rooted by $p\mathbf{v}$ which represents the (unfolding of) the part of $\mathcal{M}_{\mathcal{A}}$ reachable from $p\mathbf{v}$. A branch in this tree is terminated when a configuration of $\llbracket \varphi_1 \rrbracket \cup \dots \cup \llbracket \varphi_k \rrbracket$ is visited, or when the length of the branch reaches ℓ . For every $1 \leq i \leq k$, let L_i be the set of all leafs α of T labeled by configurations of $\llbracket \varphi_i \rrbracket$. We use P_α to denote the (rational and computable) probability of reaching α from the root of T , and $\text{label}(\alpha)$ to denote the configuration which is the label of α . For every $1 \leq i \leq k$ and every $1 \leq j \leq n_i$, we put $P_{p\mathbf{v},i,j} = \sum_{\alpha \in L_i} P_\alpha \cdot P_{\text{label}(\alpha),j}$. It is straightforward to verify that the sequence

$$\begin{aligned} & (P_{p\mathbf{v},1,1}, H_{1,1}), \dots, (P_{p\mathbf{v},1,n_1}, H_{1,n_1}), \\ & (P_{p\mathbf{v},2,1}, H_{2,1}), \dots, (P_{p\mathbf{v},2,n_2}, H_{2,n_2}), \\ & \vdots \\ & (P_{p\mathbf{v},k,1}, H_{k,1}), \dots, (P_{p\mathbf{v},k,n_k}, H_{k,n_k}) \end{aligned}$$

approximates the pattern frequencies of $\text{Run}(p\mathbf{v})$ up to the absolute error ε . In particular, realize that almost every $w \in \text{Run}(p\mathbf{v})$ eventually “decides” for some C_i , i.e., there is $m \in \mathbb{N}$ such that $w(m) \in C_i$ and for all $m' > m$ we have $w(m') \notin D_i$ (this is where we use the first condition of Definition 3). Hence, the pattern frequency vector is well defined and approximated up to the absolute error $\varepsilon/4$ by some of the above $H_{i,j}$ for almost all $w \in \text{Run}(p\mathbf{v})$. ■

For the rest of this section, we fix a two-counter pVASS $\mathcal{A} = (Q, \gamma, W)$ (recall that \mathcal{A} satisfies Assumption 1). For $i \in \{1, 2\}$, we define a one-counter pVASS $\mathcal{A}_i = (Q, \gamma_i, W_i)$ and a labeling $L_i : \gamma_i \rightarrow \{-1, 0, 1\}$ as follows: $s \xrightarrow{\kappa(i)} t$ in \mathcal{A}_i and $L_i((s, \kappa(i), t)) = \kappa(3-i)$ iff $s \xrightarrow{\kappa(i)} t$ in \mathcal{A} . Note that \mathcal{A}_i is obtained by “preserving” the i -th counter; the change of the other counter is encoded in L_i . Also observe that $\mathcal{C}_{\mathcal{A}}, \mathcal{C}_{\mathcal{A}_1}$, and $\mathcal{C}_{\mathcal{A}_2}$ are the same Markov chains.

The results of Section III are applicable to \mathcal{A}_1 and \mathcal{A}_2 . Let R be a type II or a type IV region of \mathcal{A}_i . We claim that there is a unique $\tau_R \in \mathbb{R}$ such that for almost all runs $w \in \text{Run}(p(k), R)$, where $p(k) \in R$, we have that the limit

$$\lim_{n \rightarrow \infty} \frac{\sum_{j=0}^{n-1} L_i(\text{rule}(w, j))}{n}$$

exists and it is equal to τ_R (here, $\text{rule}(w, j)$ is the unique rule of γ_i which determines the transition $w(j) \rightarrow w(j+1)$; cf. Assumption 1). In other words, τ_R is the unique mean payoff determined by the labeling L_i associated to R . To see this, consider the trend t_S of the associated BSCC S of $\mathcal{C}_{\mathcal{A}}$. If R is a type IV region, then $\tau_R = t_S(3-i)$ for almost all $w \in \text{Run}(p(k), R)$ (in particular, note that if $t_S(i) \leq 0$ then $\mathcal{P}(\text{Run}(p(k), R)) = 0$; see Section III). If $t_S(i) \geq 0$ and R is a type II region, then $\tau_R = t_S(3-i)$, because the frequency of visits to configurations with zero counter is zero for almost

all $w \in \text{Run}(p(k))$, where $p(k) \in R$ (see [4]). Finally, if $t_S(i) < 0$ and R is a type II region, then R is ergodic because the mean recurrence time in every configuration of R is finite [4], and hence τ_R takes the same value for almost all $w \in \text{Run}(p(k), R)$, where $p(k) \in R$.

Although the value of τ_R may be irrational when R is of type II and $t_S(i) < 0$, there exists a formula $\Phi(x)$ of Tarski algebra with a fixed alternation depth of quantifiers computable in polynomial time such that $\Phi[x/c]$ is valid iff $c = \tau_R$. Hence, the problem whether τ_R is zero (or positive, or negative) is decidable in exponential time [9]; and if $\tau_R < 0$ (or $\tau_R > 0$), there is a computable $x \in \mathbb{Q}$ such that $x < 0$ (or $x > 0$) and $|x| \leq |\tau_R|$.

Definition 5. Let $\mathcal{A} = (Q, \gamma, W)$ be a pVASS of dimension two. We say that \mathcal{A} is stable if the following conditions are satisfied:

- Let S be a BSCC of $\mathcal{C}_{\mathcal{A}}$ such that the type IV region determined by S is non-empty in \mathcal{A}_1 or \mathcal{A}_2 , or there is $p \in S$ such that the type II region determined by $p(0)$ is non-empty in \mathcal{A}_1 or \mathcal{A}_2 . Then the trend t_S is non-zero in both components.
- Let R be a type II region in \mathcal{A}_i such that $t_S(i) < 0$, where S is the BSCC of $\mathcal{C}_{\mathcal{A}}$ associated to R . Then $\tau_R \neq 0$.

Note that the problem whether a given two-counter pVASS \mathcal{A} is stable is decidable in exponential time. Our aim is to prove the following theorem:

Theorem 2. Let $\mathcal{A} = (Q, \gamma, W)$ be a stable pVASS of dimension two. Then the set $\text{conf}(\mathcal{A})$ is good for \emptyset .

For the rest of this section, we fix a pVASS \mathcal{A} of dimension two a present a sequence of observations that imply Theorem 2. Note that \mathcal{A} is not necessarily stable, i.e., the presented observations are valid for *general* two-dimensional pVASS. The stability condition is used to rule out some problematic subcases that are not covered by these observations.

In our constructions, we need to consider the following subsets of configurations:

- $C[R_1, R_2]$, where $R_1 \in \text{Reg}(\mathcal{A}_1)$ and $R_2 \in \text{Reg}(\mathcal{A}_2)$, is the set of all $p(m_1, m_1) \in \text{conf}(\mathcal{A})$ such that $p(m_1) \in R_1$ and $p(m_2) \in R_2$;
- $B[b]$, where $b \in \mathbb{N}$, consists of all $p\mathbf{v} \in \text{conf}(\mathcal{A})$ such that for every $q\mathbf{u} \in \text{post}^*(p\mathbf{v})$ we have that $\mathbf{u}(1) \leq b$ or $\mathbf{u}(2) \leq b$;
- $C_S[c_1 \sim b_1 \wedge c_2 \approx b_2]$, where $S \subseteq Q$, $b_1, b_2 \in \mathbb{N}$, and \sim, \approx are numerical comparisons (such as $=$ or \leq) consists of all $p(m_1, m_2) \in \text{conf}(\mathcal{A})$ such that $p \in S$, $m_1 \sim b_1$, and $m_2 \approx b_2$. Trivial constraints of the form $c_i \geq 0$ can be omitted. For example, $C_Q[c_1 = 0 \wedge c_2 \geq 6]$ is the set of all $q(0, m) \in \text{conf}(\mathcal{A})$ where $m \geq 6$, and $C_S[c_1 \leq 2]$ is the set of all $q(n, m) \in \text{conf}(\mathcal{A})$ where $q \in S$ and $n \leq 2$.
- Z_S , where $S \subseteq Q$, consists of all $p(m_1, m_2)$ such that $p \in S$ and some counter is zero (i.e., $m_1 = 0$ or $m_2 = 0$).
- $E_S[b_1, b_2]$, where $S \subseteq Q$ and $b_1, b_2 \in \mathbb{N}$, consists of all $p(m_1, m_2)$ such that $p \in S$, some counter is zero,

and every $q(n_1, n_2) \in \text{post}^*(p(m_1, m_2))$ satisfies the following:

- if $n_1 = 0$, then $n_2 \leq b_2$;
- if $n_2 = 0$, then $n_1 \leq b_1$.

Note that all of these sets are semilinear and the associated semilinear constraints are computable.

A direct consequence of Lemma 1 is the following:

Lemma 6. Let $b = 11|Q|^4$, and let \mathcal{E} be a set consisting of $B[b]$ and all $C[R_1, R_2]$ where $R_1 \in \text{Reg}(\mathcal{A}_1)$, $R_2 \in \text{Reg}(\mathcal{A}_2)$. Then $\text{conf}(\mathcal{A})$ is reducible to \mathcal{E} .

To prove Lemma 6, it suffices to realize that there is a computable $k \in \mathbb{N}$ such that every $p\mathbf{v} \in \text{conf}(\mathcal{A})$ can reach a configuration of some $C[R_1, R_2]$ or $B[b]$ in at most k transitions.

Hence, it suffices to prove that $B[b]$ and all $C[R_1, R_2]$ are good for \emptyset . All cases except for those where R_1 and R_2 are of type II or type IV follow almost immediately. To handle the remaining cases, we need to develop new tools, which we present now. We start by introducing some notation.

Given a finite path or a run w in $\mathcal{M}_{\mathcal{A}}$ and $\ell \in \mathbb{N}$, where $\ell \leq \text{length}(w)$, we denote by $x_1^{(\ell)}(w)$, $x_2^{(\ell)}(w)$, and $p^{(\ell)}(w)$ the value of the first counter, the value of the second counter, and the control state of the configuration $w(\ell)$, respectively. Further, $T(w)$ denotes either the least ℓ such that $x_1^{(\ell)}(w) = 0$, or ∞ if there is no such ℓ . For every $i \in \mathbb{N}$, $[p\mathbf{v} \rightarrow^* q\mathbf{u}, i]$ denotes the probability of all $w \in \text{Run}(p\mathbf{v})$ such that $T(w) \geq i$, $w(i) = q\mathbf{u}$, and $w(j) \neq q\mathbf{u}$ for all $0 \leq j < i$. By $[p\mathbf{v} \rightarrow^* q\mathbf{u}] = \sum_{i=0}^{\infty} [p\mathbf{v} \rightarrow^* q\mathbf{u}, i]$ we denote the probability of reaching $q\mathbf{u}$ from $p\mathbf{v}$ before time T . We also put

$$[p\mathbf{v} \rightarrow^* q(0, *), i] = \sum_{k=0}^{\infty} [p\mathbf{v} \rightarrow^* q(0, k), i]$$

and

$$[p\mathbf{v} \rightarrow^* q(0, *)] = \sum_{k=0}^{\infty} [p\mathbf{v} \rightarrow^* q(0, k)].$$

For a measurable function X over the runs of $\mathcal{M}_{\mathcal{A}}$, we use $\mathbb{E}_{p\mathbf{v}}[X]$ to denote the expected value of X over $\text{Run}(p\mathbf{v})$.

The following theorems are at the very core of our analysis, and represent new non-trivial quantitative bounds obtained by designing and analyzing a suitable martingale. Proofs can be found in [5].

Theorem 3. Let S be a BSCC of $\mathcal{C}_{\mathcal{A}}$ such that $t_S(2) < 0$, and let R be a type II region of \mathcal{A}_2 determined by some state of S . Then there are rational $a_1, b_1 > 0$ and $0 < z_1 < 1$ computable in polynomial space such that the following holds for all $p(0) \in R$, $n \in \mathbb{N}$, and $i \in \mathbb{N}^+$:

$$\mathcal{P}_{p(n,0)}(T < \infty \wedge x_2^{(T)} \geq i) \leq a_1 \cdot z_1^{b_1 \cdot i}.$$

Moreover, if $\mathcal{P}_{p(n,0)}(T < \infty) = 1$, then

$$\mathbb{E}_{p(n,0)} \left[x_2^{(T)} \right] \leq \frac{a_1 \cdot z_1^{b_1}}{1 - z_1^{b_1}}.$$

In particular, none of the bounds depends on n .

Theorem 4. Let S be a BSCC of \mathcal{C}_A such that $t_S(2) < 0$, and let R be a type II region determined by some state of S such that $\tau_R > 0$. Then there are rational $a_2, b_2 > 0$ and $0 < z_2 < 1$ computable in polynomial space such that for all configurations $p(n, 0)$, where $p(0) \in R$, and all $q \in Q$, the following holds:

$$[p(n, 0) \rightarrow^* q(0, *)] \leq n \cdot a_2 \cdot z_2^{n \cdot b_2}$$

Theorem 5. Let R be a type II region of \mathcal{A}_2 such that $\tau_R < 0$. Then there are rational $a_3, b_3, d_3 > 0$ and $0 < z_3 < 1$ computable in polynomial space such that for all configurations $p(n, 0)$, where $p(0) \in R$, and all $q \in Q$, the following holds for all $i \geq \frac{H \cdot n}{-\tau_R}$, where H is computable in polynomial space:

$$[p(n, 0) \rightarrow^* q(0, *), i] \leq i \cdot a_3 \cdot z_3^{\sqrt{n \cdot \tau_R \cdot b_3 + i \cdot d_3}}.$$

The above theorems are used to prove that certain configurations are *eagerly attracted* by certain sets of configurations in the following sense:

Definition 6. Let $C, D \subseteq \text{conf}(\mathcal{A})$. We say that $p\mathbf{v} \in C$ is *eagerly attracted* by D if $\mathcal{P}(\text{Run}(p\mathbf{v} \rightarrow^* D)) = 1$ and there are computable constants $a, z \in \mathbb{Q}$, $\ell \in \mathbb{N}$, and $k \in \mathbb{N}^+$ (possibly dependent on $p\mathbf{v}$), where $a > 0$ and $0 < z < 1$, such that for every $\ell' \geq \ell$, the probability of visiting D from $p\mathbf{v}$ in at most ℓ' transitions is at least $1 - a \cdot z^{\frac{k}{\ell'}}$. Further, we say that C is *eagerly attracted* by D if all configurations of C are *eagerly attracted* by D , and D is a finite eager attractor if D is finite and $\text{post}^*(D)$ is *eagerly attracted* by D .

Markov chains with finite eager attractors were studied in [1]. The only subtle difference is that in [1], the probability of revisiting the attractor in at most ℓ transitions is at least $1 - z^\ell$. However, all arguments of [1] are valid also for the sub-exponential bound $1 - a \cdot z^{\frac{k}{\ell}}$ adopted in Definition 6 (note that some quantitative bounds given in [1], such as the bound on K in Lemma 5.1 of [1], need to be slightly adjusted to accommodate the sub-exponential bound). In [1], it was shown that various limit properties of Markov chains with finite eager attractors can be effectively approximated up to an arbitrarily small absolute error $\varepsilon > 0$. A direct consequence of these results is the following:

Proposition 1. Let $D \subseteq \text{conf}(\mathcal{A})$ be a finite eager attractor. Then D is good for \emptyset .

Let us also formulate one simple consequence of Theorem 4.

Corollary 1. For every BSCC S of \mathcal{C}_A we have the following:

- If t_S is negative in some component, then every configuration $p\mathbf{v}$ where $p \in S$ is *eagerly attracted* by Z_S .
- If both components of t_S are positive, then for every $\varepsilon > 0$ there is a computable b_ε such that for every configuration $p\mathbf{v}$ where $p \in S$ and $\mathbf{v} \geq (b_\varepsilon, b_\varepsilon)$ we have that $\mathcal{P}(\text{Run}(p\mathbf{v} \rightarrow^* Z_S)) < \varepsilon$.

The following theorem follows from the results about one-counter pVASS presented in [4].

Theorem 6. For every BSCC S of \mathcal{C}_A we have the following:

- If t_S is negative in some component, then every configuration $p\mathbf{v}$ where $p \in S$ is *eagerly attracted* by Z_S .
- If both components of t_S are positive, then for every $\varepsilon > 0$ there is a computable b_ε such that for every configuration $p\mathbf{v}$ where $p \in S$ and $\mathbf{v} \geq (b_\varepsilon, b_\varepsilon)$ we have that $\mathcal{P}(\text{Run}(p\mathbf{v} \rightarrow^* Z_S)) \leq \varepsilon$.

In the next lemmata, we reduce the study of pattern frequencies for certain runs in \mathcal{M}_A to the study of pattern frequencies for runs in one-counter pVASS (i.e., to the results of Section III). This is possible because in each of these cases, one of the counters is either bounded or irrelevant. Proofs of the following lemmata are straightforward.

Lemma 7. For every $b \in \mathbb{N}$, the set $B[b]$ is good for \emptyset .

Lemma 8. The set $C[R_1, R_2]$, where R_1 or R_2 is a type I or a type III region, is good for \emptyset .

So, it remains to consider sets of the form $C[R_1, R_2]$, where the regions R_1, R_2 are of type II or type IV. We start with the simple case when the trend t_S of the associated BSCC is positive in both components.

Lemma 9. Let $C[R_1, R_2]$ be a set such that R_1, R_2 are regions of type II or type IV, and the trend t_S of the associated BSCC S of \mathcal{C}_A is positive in both components. Then $C[R_1, R_2]$ is good for \emptyset .

Proof: Let $b \in \mathbb{N}$ be a bound such that for every $p\mathbf{v} \in \text{conf}(\mathcal{A})$ where $p \in S$ and $\mathbf{v} \geq (b, b)$ we have that there exists a “pumpable path” of the form $p\mathbf{v} \rightarrow^* p(\mathbf{v} + \mathbf{u})$ where \mathbf{u} is positive in both components. Note that such a b exists and it is computable (in fact, one can give an explicit upper bound on b in the size of S ; see, e.g., [3]).

By Lemma 7, $B[b]$ is good for \emptyset . We show that $C_S[c_1 \geq b \wedge c_2 \geq b]$ is good for $B[b]$. By our choice of b and Theorem 6, there is $\delta > 0$ such that $\mathcal{P}(\text{Run}(p\mathbf{v} \not\rightarrow^* B[b])) \geq \delta$ for every $p\mathbf{v} \in C_S[c_1 \geq b \wedge c_2 \geq b]$. Further, almost all runs of $\text{Run}(p\mathbf{v} \not\rightarrow^* B[b])$ have the same pattern frequency vector F_S where $F_S(q(*, *)) = \mu_S(q)$ for all $q \in S$, and $F_S(\alpha) = 0$ for the other patterns.

Now we prove that $C[R_1, R_2]$ is reducible to $\{B[b], C_S[c_1 \geq b \wedge c_2 \geq b]\}$. By Theorem 6, we obtain that for every $\varepsilon > 0$ there is a computable b_ε such that for every configuration of $q\mathbf{u}$ where $\mathbf{u} \geq (b_\varepsilon, b_\varepsilon)$ we have that $\mathcal{P}(\text{Run}(q\mathbf{u} \rightarrow^* Z_S)) \leq \varepsilon$. Let φ be a semilinear constraint where $\varphi(s) = x \geq b + b_\varepsilon \wedge y \geq b + b_\varepsilon$ for all $s \in S$, and $\varphi(s) = \text{false}$ for all $s \in Q \setminus S$. Then $\llbracket \varphi \rrbracket \subseteq C_S[c_1 \geq b \wedge c_2 \geq b]$ and for every $q\mathbf{u} \in \llbracket \varphi \rrbracket$ we have that $\mathcal{P}(\text{Run}(q\mathbf{u} \rightarrow^* B[b])) \leq \varepsilon$. Further, there exists a computable $k \in \mathbb{N}$ such that every configuration of $C[R_1, R_2]$ can reach a configuration of $B[b] \cup \llbracket \varphi \rrbracket$ in at most k transitions. This implies that for every $\delta > 0$, there is a computable $\ell \in \mathbb{N}$ such that every configuration of $C[R_1, R_2]$ reaches a configuration of $B[b] \cup \llbracket \varphi \rrbracket$ in at most ℓ steps with probability at least $1 - \delta$. ■

To prove Theorem 2, it suffices to show that the following sets of configurations are good for \emptyset , where we disregard the subcases ruled out by the stability condition. In particular, due to Lemma 9 we can safely assume that some component of t_S is negative.

- (a) $C[R_1, R_2]$, where both R_1 and R_2 are of type II.
- (b) $C[R_1, R_2]$, where R_1 is of type IV and R_2 is of type II, or R_1 is of type II and R_2 is of type IV.
- (c) $C[R_1, R_2]$, where both R_1 and R_2 are of type IV.

The most interesting (and technically demanding) is the following subcase of Case (a). Here we only sketch the main ideas, a full proof can be found in [5].

Lemma 10. *Let $C[R_1, R_2]$ be a set of configurations where R_1 and R_2 are of type II, $t_S(2) < 0$, $\tau_{R_1} < 0$, and $\tau_{R_2} < 0$. Then $C[R_1, R_2]$ is good for \emptyset .*

Proof Sketch: Let C be the set of all configurations of the form $q(0, m) \in C[R_1, R_2]$ satisfying $m \leq (a_1 \cdot z_1^{b_1}) / (1 - z_1^{b_1})$, where a_1, b_1, z_1 are the computable constants of Theorem 3. We prove that $C[R_1, R_2]$ is eagerly attracted by C . This immediately implies that C is a finite eager attractor, hence C is good for \emptyset by Proposition 1. We also immediately obtain that $C[R_1, R_2]$ is reducible to $\{C\}$, which means that $C[R_1, R_2]$ is good for \emptyset by Lemma 5.

Let $p\mathbf{v} \in C[R_1, R_2]$. Since $t_S(2) < 0$ and $\tau_{R_1} < 1$, almost every run $w \in \text{Run}(p\mathbf{v})$ eventually visits a configuration of $C_S[c_2 = 0]$, and, from that moment on, visits configurations of both $C_S[c_2 = 0]$ and $C_S[c_1 = 0]$ infinitely often.

Denote by $\Theta_0(w)$ the least ℓ such that $w(\ell) \in C_S[c_2 = 0]$. Given $k \geq 1$, denote by $\Theta_k(w)$ the least $\ell \geq \Theta_{k-1}(w)$ such that the following holds:

- If k is odd, then $w(\ell) \in C_S[c_1 = 0]$.
- If k is even, then $w(\ell) \in C_S[c_2 = 0]$.

We use Theorems 6, 3, and 5 to show that there are computable constants $\hat{a} > 0$ and $0 < \hat{z} < 1$ such that for all $k \geq 0$ and all $\ell \in \mathbb{N}$ we have that

$$\mathcal{P}_{p\mathbf{v}}(\Theta_k - \Theta_{k-1} \geq \ell) \leq \hat{a} \cdot (\hat{z})^{\sqrt{\ell}}$$

Here $\Theta_{-1} = 0$. Observe that Θ_0 is the sum of the number of transitions needed to visit Z_S for the first time (the first phase) and the number of transitions need to reach $C_S[c_2 = 0]$ subsequently (the second phase). Due to Theorem 6, the probability that the first phase takes more than ℓ transitions is bounded by $a \cdot z^\ell$ for some computable $a > 0$ and $0 < z < 1$. Note that the length of the second phase depends on the value of c_2 after the first phase. However, the probability that this value will be larger than ℓ can be bounded by $a \cdot z^\ell$ as well. Finally, assuming that the first phase ends in a configuration $q(0, m)$, Theorem 5 gives a bound $a' \cdot (z')^{\sqrt{\ell-m}}$ on the probability of reaching $C_S[c_2 = 0]$ in at least ℓ transitions. By combining these bounds appropriately, we obtain the above bound on Θ_0 .

Now let us consider $\Theta_k - \Theta_{k-1}$ for $k > 0$. Let us assume that k is even (the other case follows similarly). The only difference from the previous consideration (for Θ_0) is that

now the first phase consists of the part of the run up to the Θ_{k-1} -th configuration, and the second phase from there up to the Θ_k -th configuration. Using Theorem 3 and induction hypothesis, we derive a bound $a \cdot z^\ell$ on the probability that the height of the second counter in the Θ_{k-1} -th configuration will be at least ℓ . Then, as above, we combine this bound with the bound on the probability of reaching $C_S[c_2 = 0]$ in ℓ steps from a fixed configuration of $C_S[c_1 = 0]$.

In order to finish the proof, we observe that the probability of reaching a configuration of C between the Θ_{k-1} -th and Θ_k -th configuration is bounded away from zero by a computable constant. This follows immediately from Theorem 3 which basically bounds the expected value of c_2 in the Θ_k -th configuration. Denoting by $\text{Rounds}(w)$ the least number k such that $w(\Theta_k(w)) \in C$, we may easily show that $\mathcal{P}_{p\mathbf{v}}(\text{Rounds} \geq \ell) \leq \bar{c}^\ell$ for a computable constant $0 \leq \bar{c} < 1$.

Finally, we combine the bound on the number of rounds (i.e., the bound on $\mathcal{P}_{p\mathbf{v}}(\text{Rounds} \geq \ell)$) with the bound on the length of each round (i.e., the bound on $\mathcal{P}_{p\mathbf{v}}(\Theta_k - \Theta_{k-1} \geq \ell)$), and thus obtain the desired bound on the number of steps to visit C . ■

For the other cases (incl. Cases (b) and (c)), we show that the set of configurations C we aim to analyze is eagerly attracted by computable semilinear sets of configurations C_1, \dots, C_k , where each C_i is either good for \emptyset or good for $\bigcup_{i \neq j} C_j$. In all these cases, it is easy to see that the configurations of C reach a configuration of $\bigcup_{i=1}^k C_i$ with probability one, and the argument that C is *eagerly* attracted $\bigcup_{i=1}^k C_i$ is a simplified version of the proof of Lemma 10 (in some cases, the proof is substantially simpler than the one of Lemma 10). Therefore, in these cases we just list the sets C_1, \dots, C_k and add some intuitive comments which explain possible behaviour of the runs initiated in configurations of C .

When defining the aforementioned sets C_1, \dots, C_k , we use the following computable constants $B_{II}, B_{IV}, D_{II} \in \mathbb{N}$, which are numbers (not necessarily the least ones) satisfying the following conditions:

- if $p(0) \in R$, where R is a type IV region of \mathcal{A}_i for some $i \in \{1, 2\}$, then $p(0)$ can reach a configuration a type I region in at most B_{IV} transitions.
- if $p(0) \in R$, where R is a type II region of \mathcal{A}_i such that $t_S(i) < 0$ and $\tau_R > 0$, then there is a finite path w from $p(0)$ to $p(0)$ of length smaller than B_{II} such that the total L_i -reward of all transitions executed in w is positive.
- for every $p\mathbf{v} \in \text{conf}(\mathcal{A})$ and every $i \in \{1, 2\}$, if $\mathbf{v}(i) = 0$, $\mathbf{v}(3-i) \geq D_{II}$, and $p(0) \in R$ for some type II region of \mathcal{A}_i such that either $t_S(i) > 0$ and $t_S(3-i) < 0$, or $t_S(i) < 0$ and $\tau_R < 0$, then there exists $qu \in \text{post}^*(p\mathbf{v})$ such that $\mathbf{u}(i) \geq \max\{B_{II}, B_{IV}\}$ and $\mathbf{u}(3-i) = 0$.

The existence and computability of B_{II} , B_{IV} , and D_{II} follows from simple observations about the transition structure of $\mathcal{M}_{\mathcal{A}}$ (these constants are in fact small and their size can be explicitly bounded in $\|\mathcal{A}\|$).

Lemma 11. *For all $m, n \in \mathbb{N}$ and a BSCC S of $\mathcal{C}_{\mathcal{A}}$ such that*

t_S is negative in some component, the set $E_S[m, n]$ is good for \emptyset .

Proof: From the definition of $E_S[m, n]$ and Theorem 6, we immediately obtain that $E_S[m, n]$ is a finite eager attractor (even if $E_S[m, n] = \emptyset$). Hence, the claim follows from Proposition 1. ■

Now we consider the remaining subcases of Case (a).

Lemma 12. *Let $C[R_1, R_2]$ be a set of configurations where R_1 and R_2 are of type II, $t_S(2) < 0$, $t_S(1) > 0$, and $\tau_{R_2} > 0$. Then $C[R_1, R_2]$ is good for \emptyset .*

Proof: Let $\mathcal{E} = \{E[B_{II}, D_{II}], C_S[c_2=0 \wedge c_1 \geq B_{II}]\}$. Observe that $E[B_{II}, D_{II}]$ is good for \emptyset by Lemma 11. We show that $C_S[c_2=0 \wedge c_1 \geq B_{II}]$ is good for $E[B_{II}, D_{II}]$ and that $C[R_1, R_2]$ reducible to \mathcal{E} . Hence, $C[R_1, R_2]$ is good for \emptyset by Lemma 5.

To see that $C_S[c_2=0 \wedge c_1 \geq B_{II}]$ is good for $E[B_{II}, D_{II}]$, realize that for every $p\mathbf{v} \in C_S[c_2=0 \wedge c_1 \geq B_{II}]$ we have that almost all runs of $Run(p\mathbf{v})$ that do not visit a configuration of $E[B_{II}, D_{II}]$ eventually behave as if the first counter did not exist, which means that the long-run behaviour of almost all of these runs is the same as the behavior of the runs of \mathcal{A}_2 initiated in $p(0)$ (here we also use the defining property of D_{II}). Further, it follows from the definition of B_{II} and Theorem 4 that there exists a $\delta > 0$ such that $\mathcal{P}(Run(p\mathbf{v}) \not\rightarrow^* E[B_{II}, D_{II}]) > \delta$ for every $p\mathbf{v} \in C_S[c_2=0 \wedge c_1 \geq B_{II}]$.

By Theorem 4, for every $\varepsilon > 0$ there exists a computable semilinear constraint φ such that $\llbracket \varphi \rrbracket \subseteq C_S[c_2=0 \wedge c_1 \geq B_{II}]$ and for every $q\mathbf{u} \in \llbracket \varphi \rrbracket$ we have that the probability of visiting $C[c_1=B_{II}]$ (and hence also $E[B_{II}, D_{II}]$) is bounded by ε .

Now let $p\mathbf{v} \in C[R_1, R_2]$ and $\delta > 0$. We need to show that there is a computable $\ell \in \mathbb{N}$ such that the probability of reaching a configuration of $E[B_{II}, D_{II}] \cup \llbracket \varphi \rrbracket$ in at most ℓ transitions is at least $1 - \delta$. Since $t_S(2) < 0$, every $p\mathbf{v} \in C[R_1, R_2]$ is eagerly attracted by Z_S . Similarly as in the proof of Lemma 10, we show that almost every run visits $C_S[c_2=0]$ infinitely many times, and that the probability that the length between two consecutive visits to $C_S[c_2=0]$ exceeds ℓ decays sub-exponentially in ℓ . Further, the probability of visiting a configuration of $E[B_{II}, D_{II}] \cup \llbracket \varphi \rrbracket$ from a configuration of $C_S[c_2=0]$ is bounded away from zero by a fixed constant. Hence, we can argue as in the proof of Lemma 10. ■

Lemma 13. *Let $C[R_1, R_2]$ be a set of configurations where R_1 and R_2 are of type II, $t_S(2) < 0$, $t_S(1) < 0$, $\tau_{R_1} > 0$, and $\tau_{R_2} > 0$. Then $C[R_1, R_2]$ is good for \emptyset .*

Proof: Let \mathcal{E} be the set consisting of $E[B_{II}, B_{II}]$, $C_S[c_2=0 \wedge c_1 \geq B_{II}]$, and $C_S[c_1=0 \wedge c_2 \geq B_{II}]$. Clearly, each $C \in \mathcal{E}$ is either good for \emptyset or good for the union of all sets in $\mathcal{E} \setminus \{C\}$ (see Lemma 11 and the proof of Lemma 12). For every $\varepsilon > 0$, there are computable semilinear constraint φ_1, φ_2 such that $\llbracket \varphi_1 \rrbracket \subseteq C_S[c_2=0 \wedge c_1 \geq B_{II}]$, $\llbracket \varphi_2 \rrbracket \subseteq C_S[c_1=0 \wedge c_2 \geq B_{II}]$ satisfying the requirements of Definition 4. Note that there is a $k \in \mathbb{N}$ such that for every

configuration of Z_S there is a finite path of length at most k to a configuration of $E[B_{II}, B_{II}] \cup \llbracket \varphi_1 \rrbracket \cup \llbracket \varphi_2 \rrbracket$. The rest of the argument is even simpler than in Lemma 12. ■

Lemma 14. *Let $C[R_1, R_2]$ be a set of configurations where R_1 and R_2 are of type II, $t_S(2) < 0$, $t_S(1) < 0$, $\tau_{R_1} < 0$, and $\tau_{R_2} > 0$. Then $C[R_1, R_2]$ is good for \emptyset .*

Proof: Let $\mathcal{E} = \{E[B_{II}, D_{II}], C_S[c_2=0 \wedge c_1 \geq B_{II}]\}$. We show that $C[R_1, R_2]$ reducible to \mathcal{E} similarly as in Lemma 12. ■

The case when R_1 and R_2 are of type II, $t_S(2) < 0$, $t_S(1) < 0$, $\tau_{R_1} > 0$, and $\tau_{R_2} < 0$ is symmetric to the case considered in Lemma 14.

Now we continue with Case (b)

Lemma 15. *Let $C[R_1, R_2]$ be a set of configurations where R_1 is of type IV and R_2 is of type II such that $t_S(2) < 0$ and $\tau_{R_2} > 0$. Then $C[R_1, R_2]$ is good for \emptyset .*

Proof: Let \mathcal{E} be the set consisting of $E[B_{II}, B_{IV}]$, $C_S[c_2=0 \wedge c_1 \geq B_{II}]$, and all $C[R'_1, R_2]$, where R'_1 is a type I region reachable from R_1 in \mathcal{A}_1 . We show that $C[R_1, R_2]$ reducible to \mathcal{E} similarly as in previous lemmata. ■

Lemma 16. *Let $C[R_1, R_2]$ be a set of configurations where R_1 is of type IV and R_2 is of type II such that $t_S(2) < 0$ and $\tau_{R_2} < 0$. Then $C[R_1, R_2]$ is good for \emptyset .*

Proof: Let \mathcal{E} be the set consisting of $E[D_{II}, B_{IV}]$ and all $C[R'_1, R_2]$, where R'_1 is a type I region reachable from R_1 in \mathcal{A}_1 . Then $C[R_1, R_2]$ reducible to \mathcal{E} and each $C \in \mathcal{E}$ is good for \emptyset . ■

Lemma 17. *Let $C[R_1, R_2]$ be a set of configurations where R_1 is of type II and R_2 is of type IV such that $t_S(2) < 0$ and $t_S(1) > 0$. Then $C[R_1, R_2]$ is good for \emptyset .*

Proof: Let \mathcal{E} be the set consisting of $E[B_{IV}, D_{II}]$ and all $C[R_1, R'_2]$, where R'_2 is a type I region reachable from R_2 in \mathcal{A}_2 . Then $C[R_1, R_2]$ reducible to \mathcal{E} . Further, all elements of \mathcal{E} are good for \emptyset . ■

Note that the case when R_1 is of type II and R_2 is of type IV such that $t_S(2) < 0$ and $t_S(1) < 0$ is symmetric to the cases covered in Lemma 15 and Lemma 16.

Finally, in the next lemma we consider Case (c).

Lemma 18. *Let $C[R_1, R_2]$ be a set of configurations where both R_1 and R_2 are type IV regions, and the trend t_S of the associated BSCC S of \mathcal{C}_A is negative in some component. Then $C[R_1, R_2]$ is good for \emptyset .*

Proof: Let \mathcal{E} be the set consisting of $E[B_{IV}, B_{IV}]$ and all $C[R_1, R'_2]$, $C[R'_1, R_2]$, $C[R'_1, R'_2]$, where R'_i is a type I region reachable from R_i in \mathcal{A}_i (for $i \in \{1, 2\}$). We show that $C[R_1, R_2]$ reducible to \mathcal{E} . ■

V. SOME NOTES ON THREE-COUNTER PVASS

In this section we give an example of a 3-dimensional pVASS \mathcal{A} such that \mathcal{M}_A is strongly connected, and the pattern frequency vector seems to take the \perp value with




Fig. 2: A 3-dimensional pVASS \mathcal{A} . For suitable weights $P, Q, R > 0$, we have that $F_{\mathcal{A}} = \perp$ almost surely.

probability one (this intuition is confirmed by Monte Carlo simulations, see below). Further, the example is insensitive to small changes in rule weights, and it also shows that the method of Section IV based on constructing pVASS of smaller dimension by “forgetting” one of the counters and then studying the “trend” of this counter in the smaller pVASS is insufficient for three (or more) counters.

The pVASS \mathcal{A} is shown in Fig. 2. Some rules increase the counter by more than 1, so these should be formally replaced by several rules using auxiliary control states. Intuitively, \mathcal{A} behaves in the following way. Suppose we start in an initial configuration $p(m, 0, 0)$, where m is “large”. Then, \mathcal{A} starts to decrease the first counter and increase the second one. On average, the value of the second counter becomes $2m$ when the first counter is decreased to zero, and the third counter is kept “small”. So, “on average” we eventually reach a configuration $p(0, 2m, 0)$ in about $2m$ transitions. Then, the second counter is decreased and the third counter is increased, where the value is again doubled “on average”, using $4m$ transitions. Thus, we reach a configuration $p(0, 0, 4m)$. Then, we “pump” the tokens from the third counter to the first one, reaching $p(8m, 0, 0)$ in about $8m$ transitions. And so on. Observe that the k -th phase takes about 2^k transitions, and so at the end of each phase, about half of the time was spent in configurations with the “current” pattern. Hence, the pattern frequency oscillates.

A precise formulation of this phenomenon, and a formal proof that almost all runs really behave in the above indicated way, are technically demanding and we do not provide them in this paper. For the reader’s convenience, we have implemented a simple Maple sheet which can be used to perform Monte Carlo simulations of \mathcal{A} and observe the above described phenomenon in practice².

Note that the oscillation of \mathcal{A} is insensitive to small changes in rule weights. However, if we modify \mathcal{A} into \mathcal{A}' so that the counter value is *decreased* on average in each phase (e.g., we start in $p(m, 0, 0)$, and then reach $p(0, m-1, 0)$, $p(0, 0, m-2)$, $p(m-3, 0, 0)$, etc., on average), then the *sum* of the counters has a tendency to decrease and $\mathcal{M}_{\mathcal{A}'}$ has a finite attractor. This means that the pattern frequency vector is well defined for almost all runs of \mathcal{A}' . Still, the behaviour of all two-counter

machines $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ obtained from \mathcal{A} by “forgetting” the first, the second, and the third counter, is essentially similar to the behaviour of $\mathcal{B}'_1, \mathcal{B}'_2$, and \mathcal{B}'_3 obtained from \mathcal{A}' in the same way (for example, both in \mathcal{B}_1 and \mathcal{B}'_1 , the second counter has a tendency to increase and the third has a tendency to decrease). Hence, we cannot distinguish between the behaviour of \mathcal{A} and \mathcal{A}' just by studying the “trends” in the two-counter pVASS obtained by “forgetting” one of the counters. This indicates that the study of 3-dimensional pVASS requires different (and perhaps more advanced) methods than those presented in this paper.

ACKNOWLEDGEMENT

Tomáš Brázdil and Antonín Kučera are supported by the Czech Science Foundation, Grant No. 15-17564S. The research leading to these results has received funding from the People Programme (Marie Curie Actions) of the European Union’s Seventh Framework Programme (FP7/2007-2013) under REA grant agreement no [291734].

REFERENCES

- [1] P. 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.
- [2] 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.
- [3] T. Brázdil, P. Jančar, and A. Kučera. Reachability games on extended vector addition systems with states. In *Proceedings of ICALP 2010, Part II*, volume 6199 of LNCS, pages 478–489. Springer, 2010.
- [4] T. Brázdil, S. Kiefer, and A. Kučera. Efficient analysis of probabilistic programs with an unbounded counter. *JACM*, 61(6), 2014.
- [5] T. Brázdil, S. Kiefer, A. Kučera, and P. Novotný. Long-run average behaviour of probabilistic vector addition systems. *CoRR*, abs/1505.02655, 2015.
- [6] J. Esparza, D. Hansel, P. Rossmanith, and S. Schwoon. Efficient algorithms for model checking pushdown systems. In *Proceedings of CAV 2000*, volume 1855 of LNCS, pages 232–247. Springer, 2000.
- [7] 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.
- [8] G. Florin and S. Natkin. Necessary and sufficient ergodicity condition for open synchronized queueing networks. *IEEE Trans. Software Eng.*, 15(4):367–380, 1989.
- [9] D. Grigoriev. Complexity of deciding Tarski algebra. *Journal of Symbolic Computation*, 5(1–2):65–108, 1988.
- [10] M. Kudlek. Probability in Petri nets. *Fundam. Inform.*, 67(1-3):121–130, 2005.
- [11] J. Leroux and G. Sutre. On flatness for 2-dimensional vector addition systems with states. In *Proceedings of CONCUR 2004*, volume 3170 of LNCS, pages 402–416. Springer, 2004.
- [12] M. K. Molloy. Performance analysis using stochastic Petri nets. *IEEE Trans. Computers*, 31(9):913–917, 1982.
- [13] J. Norris. *Markov Chains*. Cambridge University Press, 1998.
- [14] A. Stewart, K. Etessami, and M. Yannakakis. Upper bounds for Newton’s method on monotone polynomial systems, and P-time model checking of probabilistic one-counter automata. In *Proceedings of CAV 2013*, volume 8044 of LNCS, pages 495–510. Springer, 2013.

²Available at <http://www.cs.ox.ac.uk/people/stefan.kiefer/pVASS-simulation.txt>