On the Existence and Computability of Long-Run Average Properties in Probabilistic VASS

Antonín Kučera*

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

Abstract. We present recent results about the long-run average properties of probabilistic vector additions systems with states (pVASS). Interestingly, for probabilistic pVASS with two or more counters, long-run average properties may take several different values with positive probability even if the underlying state space is strongly connected. This contradics the previous results about stochastic Petri nets established in 80s. For pVASS with three or more counters, it may even happen that the long-run average properties are undefined (i.e., the corresponding limits do not exist) for almost all runs, and this phenomenon is stable under small perturbations in transition probabilities. On the other hand, one can effectively approximate eligible values of long-run average properties and the corresponding probabilities for some sublasses of pVASS. These results are based on new exponential tail bounds achieved by designing and analyzing appropriate martingales. The paper focuses on explaining the main underlying ideas.

1 Introduction

Probabilistic vector addition systems with states (pVASS) are a stochastic extension of ordinary VASS obtained by assigning a positive integer weight to every rule. Every pVASS determines an infinite-state Markov chain where the states are pVASS configurations and the probability of a transition generated by a rule with weight ℓ is equal to ℓ/T, where T is the total weight of all enabled rules. A closely related model of stochastic Petri nets (SPN) has been studied since early 80s [10, 2] and the discrete-time variant of SPN is expressively equivalent to pVASS.

In this paper we give a summary of recent results about the long-run average properties of runs in pVASS achieved in [5, 4]. We show that long-run average properties may take several different values with positive probability even if the state-space of a given pVASS is strongly connected. It may even happen that these properties are undefined (i.e., the corresponding limits do not exist) for

* Based on a joint work with Tomáš Brázdil, Joost-Pieter Katoen, Stefan Kiefer, and Petr Novotný. The author is supported by the Czech Science Foundation, grant No. 15-17564S.

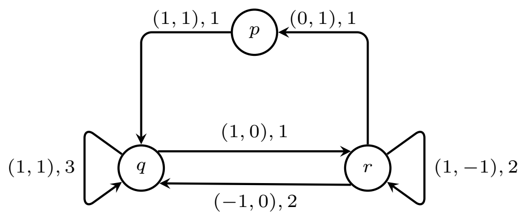

Fig. 1. An example of a two-dimensional pVASS.

almost all runs. These results contradict the corresponding claims about SPNs published in 80s (see Section 2 for more comments). On the other hand, we show that long-run average properties of runs in one-counter pVASS are defined almost surely and can be approximated up to an arbitrarily small relative error in *polynomial time*. This result is obtained by applying several deep observations about one-counter probabilistic automata that were achieved only recently. Further, we show that long-run average properties of runs in two-counter pVASS can also be effectively approximated under some technical (and effectively checkable) assumption about the underlying pVASS which prohibits some phenomena related to (possible) null-recurrency of the analyzed Markov chains.

2 Preliminaries

We use $\mathbb{Z}$, $\mathbb{N}$, $\mathbb{N}^+$, and $\mathbb{R}$ to denote the set of all integers, non-negative integers, positive integers, and real numbers, respectively. We assume familiarity with basic notions of probability theory (probability space, random variable, expected value, etc.). In particular, *Markov chains* are formally understood as pairs of the form $M = (S, \rightarrow)$, where S is a finite or countably infinite set of states and $\rightarrow \subseteq S \times (0,1] \times S$ is a transition relation such that for every $s \in S$ we have that $\sum_{s \xrightarrow{x} t} x$ is equal to one. Every state $s \in S$ determines the associated probability space over all runs (infinite paths) initiated in s in the standard way.

2.1 Probabilistic Vector Addition Systems with States

A *probabilistic Vector Addition System with States (pVASS)* with $d \geq 1$ counters is a finite directed graph whose edges are labeled by pairs κ, ℓ, where $\kappa \in \mathbb{Z}^d$ is a vector of *counter updates* and $\ell \in \mathbb{N}$ is a *weight*. A simple example of a two-counter pVASS is shown in Fig. 1. Formally, a pVASS is a triple $\mathcal{A} = (Q, \gamma, W)$, where Q is a finite set of *control states*, $\gamma \subseteq Q \times \mathbb{Z}^d \times Q$ is a set of *rules*, and $W : \gamma \rightarrow \mathbb{N}^+$ is a *weight assignment*. In the following, we 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 $W((p, \kappa, q)) = \ell$.

A *configuration* of a pVASS $\mathcal{A}$ is a pair $p\boldsymbol{v}$ where $p \in Q$ is the current control state and $\boldsymbol{v} \in \mathbb{N}^d$ is the vector of current counter values. A rule $p \xrightarrow{\kappa} q$ is *enabled* in a configuration $p\boldsymbol{v}$ iff $\boldsymbol{v} + \kappa \in \mathbb{N}^d$, i.e., the counters remain non-negative when applying the counter change κ to $\boldsymbol{v}$. The semantics of $\mathcal{A}$ is defined by its associated infinite-state Markov chain $M_{\mathcal{A}}$ whose states are the configurations of $\mathcal{A}$ and $p\boldsymbol{v} \xrightarrow{x} q\boldsymbol{u}$ if there is a rule $p \xrightarrow{\kappa} q$ with weight ℓ enabled in $p\boldsymbol{v}$ such that $\boldsymbol{u} = \boldsymbol{v} + \kappa$ and $x = \ell/T$, where T is the total weight of all rules enabled in $p\boldsymbol{v}$. If there is no rule enabled in $p\boldsymbol{v}$, then $p\boldsymbol{v}$ has only one outgoing transition $p\boldsymbol{v} \xrightarrow{1} p\boldsymbol{v}$. For example, if $\mathcal{A}$ is the pVASS of Fig. 1, then $r(3,0) \xrightarrow{1/3} p(3,1)$.

2.2 Patterns and Pattern Frequencies

Let $\mathcal{A} = (Q, \gamma, W)$ be a pVASS, and let $Pat_{\mathcal{A}}$ be the set of all *patterns* of $\mathcal{A}$, i.e., pairs of the form $p\alpha$ where $p \in Q$ and $\alpha \in \{0,+\}^d$. To every configuration $p\boldsymbol{v}$ we associate the pattern $p\alpha$ such that $\alpha_i = +$ iff $\boldsymbol{v}_i > 0$. Thus, every run $w = p_0\boldsymbol{v}_0, p_1\boldsymbol{v}_1, p_2\boldsymbol{v}_2, \ldots$ in the Markov chain $M_{\mathcal{A}}$ determines the unique sequence of patterns $p_0\alpha_0, p_1\alpha_1, p_2\alpha_2, \ldots$ For every $n \geq 1$, let $\mathcal{F}^n(w) : Pat_{\mathcal{A}} \to \mathbb{R}$ be the pattern frequency vector computed for the first n configurations of w, i.e., $\mathcal{F}^n(w)(p\alpha) = \#^n_{p\alpha}(w)/n$, where $\#^n_{p\alpha}(w)$ is the total number of all $0 \leq j < n$ such that $p_j\alpha_j = p\alpha$. The *limit* pattern frequency vector, denoted by $\mathcal{F}(w)$, is defined by $\mathcal{F}(w) = \lim_{n\to\infty} \mathcal{F}^n(w)$. If this limit does not exist, we put $\mathcal{F}(w) = \perp$.

Note that $\mathcal{F}$ is a random variable over $Run(p\boldsymbol{v})$. The very basic questions about $\mathcal{F}$ include the following:

- Do we have that $\mathcal{P}[\mathcal{F}{=}\perp] = 0$?
- Is $\mathcal{F}$ a discrete random variable?
- If so, is the set of values taken by $\mathcal{F}$ with positive probability finite?
- Can we compute these values and the associated probabilities?

Since the set of rules enabled in a configuration $p\boldsymbol{v}$ is fully determined by the associated pattern $p\alpha$, the frequency of patterns also determines the frequency of rules. More precisely, almost all runs that share the same pattern frequency also share the same frequency of rules performed along these runs, and the rule frequency is easily computable from the pattern frequency.

The above problems have been studied already in 80s for a closely related model of stochastic Petri nets (SPN). In [8], Section IV.B, is stated that if the state-space of a given SPN (with arbitrarily many unbounded places) is strongly connected, then the firing process is ergodic. 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. This result is closely related to the questions formulated above. Unfortunately, this claim is invalid. In Fig. 2, there is an example of a SPN (with weighted transitions) with two counters (places) and strongly connected state space where the limit frequency of transitions takes two eligible values (each with probability 1/2). Intuitively, if both places/counters are positive, then both of them have a tendency to decrease, i.e., a configuration where one of the counters is empty is reached almost surely. When we reach a configuration where, e.g., the first place/counter is zero and the second place/counter is positive, then the second place/counter starts to *increase*, i.e., it never becomes zero again with some positive probability. The first place/counter stays zero for most of the time, because when it

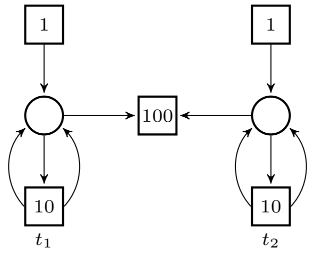

Fig. 2. A discrete-time SPN $\mathcal{N}$.

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 positive 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 these two scenarios, and therefore there are two limit frequencies of transitions, each of which is taken with probability 1/2. This possibility of reversing the "global" trend of the counters after hitting zero in some counter was not considered in [8]. Further, there exists a three-counter pVASS $\mathcal{A}$ with strongly connected state-space where the limit frequency of transitions is undefined for almost all runs, and this property is preserved for all ε-perturbations in transition weights for some fixed $\varepsilon > 0$ (see [4]). So, we must unfortunately conclude that the results of [8] are invalid for fundamental reasons.

In the next sections, we briefly summarize the results of [4] about pattern frequency vector in pVASS of dimension one and two. From now on, we assume that

- every counter is changed at most by one when performing a single rule, i.e., the vector of counter updates ranges over $\{-1, 0, 1\}^d$;
- for every pair of control states p, q, there is at most one rule of the form $p \xrightarrow{\kappa} q$.

These assumptions are not restrictive, but they have some impact on complexity, particularly when the counter updates are encoded in binary.

3 Pattern Frequency in One-Counter pVASS

For one-counter pVASS, we have the following result [4]:

Theorem 1. *Let $p(1)$ be an initial configuration of a one-counter pVASS $\mathcal{A}$. Then*

- $\mathcal{P}[\mathcal{F}{=}\bot] = 0$;

– $\mathcal{F}$ *is a discrete random variable;*
– *there are at most* $2|Q|-1$ *pairwise different vectors* F *such that* $\mathcal{P}(\mathcal{F}{=}F) > 0$*;*
– *these vectors and the associated probabilities can be approximated up to an arbitrarily small relative error* $\varepsilon > 0$ *in polynomial time.*

Since pattern frequencies and the associated probabilities may take irrational values, they cannot be computed precisely in general; in this sense, Theorem 1 is the best result achievable.

A proof of Theorem 1 is not too complicated, but it builds on several deep results that have been established only recently. As a running example, consider the simple one-counter pVASS of Fig. 3 (top), where $p(1)$ is the initial configuration. The first important step in the proof of Theorem 1 is to classify the runs initiated in $p(1)$ according to their *footprints*. A footprint of a run w initiated in $p(1)$ is obtained from w by deleting all intermediate configurations in all maximal subpaths that start in a configuration with counter equal to one, end in a configuration with counter equal to zero, and the counter stays positive in all intermediate configurations (here, the first/last configurations of a finite path are not considered as intermediate). For example, let w be a run of the form

$$\underline{p(1)}, p(2), r(2), r(1), s(1), \underline{s(0)}, \underline{r(0)}, \underline{s(0)}, \underline{s(1)}, r(1), s(1), r(1), \underline{r(0)}, \ldots$$

Then the footprint of w starts with the underlined configurations

$$p(1), s(0), r(0), s(0), s(1), r(0), \ldots$$

Note that a configuration $q(\ell)$, where $\ell > 1$, is preserved in the footprint of w iff all configurations after $q(\ell)$ have positive counter value. Further, for all $p, q \in Q$, let

– $[p{\downarrow}q]$ be the probability of all runs that start with a finite path from $p(1)$ to $q(0)$ where the counter stays positive in all intermediate configurations;
– $[p{\uparrow}] = 1 - \sum_{q \in Q} [p{\downarrow}q]$.

Almost every footprint can be seen as a run in a *finite-state* Markov chain $X_{\mathcal{A}}$ where the set of states is $\{q_0, q_1, q_{\uparrow} \mid q \in Q\}$ and the transitions are determined as follows:

– $p_0 \xrightarrow{x} q_\ell$ in $X_{\mathcal{A}}$ if $x > 0$ and $p(0) \xrightarrow{x} q(\ell)$ in $M_{\mathcal{A}}$;
– $p_1 \xrightarrow{x} q_0$ in $X_{\mathcal{A}}$ if $x = [p{\downarrow}q] > 0$;
– $p_1 \xrightarrow{x} p_{\uparrow}$ in $X_{\mathcal{A}}$ if $x = [p{\uparrow}] > 0$;
– $p_{\uparrow} \xrightarrow{1} p_{\uparrow}$.

The structure of $X_{\mathcal{A}}$ for the one-counter pVASS of Fig. 3 (top) is shown in Fig. 3 (down). In particular, note that since $r_{\uparrow} = s_{\uparrow} = 0$, there are no transitions $r_1 \to r_{\uparrow}$ and $s_1 \to s_{\uparrow}$ in $X_{\mathcal{A}}$.

For almost all runs w initiated in $p(1)$, the footprint of w determines a run in $X_{\mathcal{A}}$ initiated in p_1 in the natural way. In particular, if w contains only finitely many configurations with zero counter, then the footprint of w takes

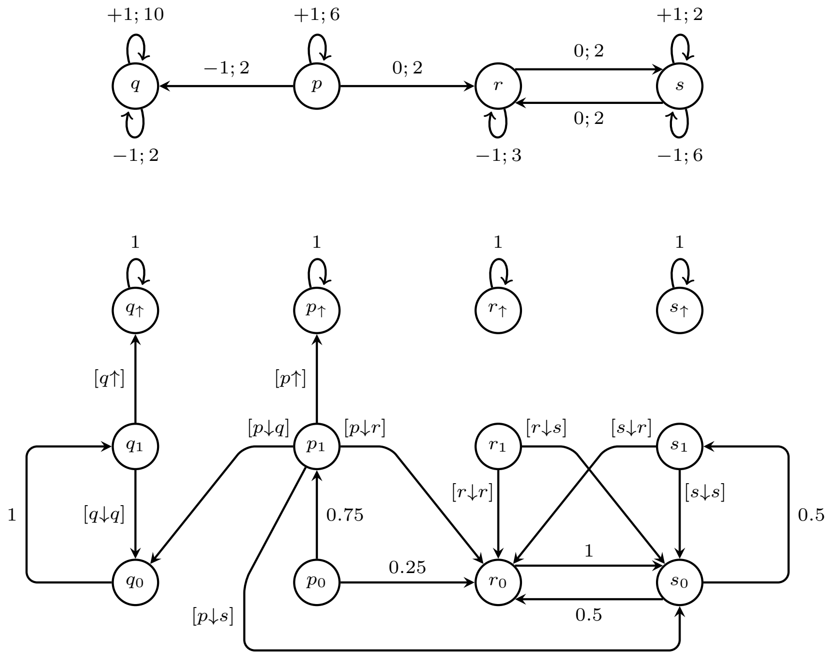

Fig. 3. A one-counter pVASS $\mathcal{A}$ (top) and its associated finite-state Markov chain $X_{\mathcal{A}}$ (down).

the form $u, s(0), r(1), v$, where $s(0)$ is the last configuration of w with zero counter and $r(1), v$ is an infinite suffix of w. This footprint corresponds to a run $u, s(0), r(1), r_\uparrow, r_\uparrow, \ldots$ of $X_{\mathcal{A}}$. In general, it may also happen that the footprint of w *cannot* be interpreted as a run in $X_{\mathcal{A}}$, but the total probability of all such w is equal to zero. As a concrete example, consider the run

$$p(1), r(1), s(0), s(1), s(2), s(3), s(4), \ldots$$

in $M_{\mathcal{A}}$ where the counter never reaches zero after the third configuration. The footprint of this run cannot be seen as a run in $X_{\mathcal{A}}$ (note that the infinite sequence $p_1, s_0, s_1, s_\uparrow, s_\uparrow, s_\uparrow, \ldots$ is *not* a run in $X_{\mathcal{A}}$).

Since almost every run in $M_{\mathcal{A}}$ initiated in $p(1)$ determines a run in $X_{\mathcal{A}}$ (via its footprint), we obtain that almost every run in $M_{\mathcal{A}}$ initiated in $p(1)$ visits a bottom strongly connected component (BSCC) of $X_{\mathcal{A}}$. Formally, for every BSCC B of $X_{\mathcal{A}}$ we define $\mathit{Run}(p(1), B)$ as the set of all runs w in $M_{\mathcal{A}}$ initiated in $p(1)$ such that the footprint of w determines a run in $X_{\mathcal{A}}$ that visits B. One is tempted to expect that almost all runs of $\mathit{Run}(p(1), B)$ share the same pattern frequency vector. This is true if (the underlying graph of) $\mathcal{A}$ has at most one *diverging* BSCC. To explain this, let us fix a BSCC D of $\mathcal{A}$ and define a Markov chain $\mathcal{D} = (D, \rightarrow)$ such that $s \xrightarrow{x} t$ in $\mathcal{D}$ iff $s \xrightarrow{\kappa,\ell} t$ and $x = \ell/T_s$, where T_s is the sum of the weights of all outgoing rules of s in $\mathcal{A}$. Further, we define the *trend of* D as follows:

$$t_d \;=\; \sum_{s \in D} \mu_D(s) \cdot \sum_{s \xrightarrow{\kappa,\ell} t} \ell / T_s \cdot \kappa \tag{1}$$

Here, μ_D is the invariant distribution of $\mathcal{D}$. Intuitively, the trend t_D corresponds to the expected change of the counter value per transition (mean payoff) in D when the counter updates are interpreted as payoffs. If t_D is positive/negative, then the counter has a tendency to increase/decrease. If $t_D = 0$, the situation is more subtle. One possibility is that the counter can never be emptied to zero when it reaches a sufficiently high value (in this case, we say that D is *bounded*). The other possibility is that the counter can always be emptied, but then the frequency of visits to a configuration with zero counter is equal to zero almost surely. We say that D is *diverging* if either $t_D > 0$, or $t_D = 0$ and D is bounded. For the one-counter pVASS of Fig. 3, we have that the BSCC $\{q\}$ is diverging, because its trend is positive. The other BSCC $\{r, s\}$ is not diverging, because its trend is negative.

Let us suppose that $\mathcal{A}$ has at most one diverging BSCC, and let B be a BSCC of $X_\mathcal{A}$. If $B = \{q_\uparrow\}$ for some $q \in Q$, then almost all runs of $Run(p(1), \{q_\uparrow\})$ share the same pattern frequency vector F where $F(s(+)) = \mu_D(s)$ for all $s \in D$, and $F(pat) = 0$ for all of the remaining patterns *pat*. In the example of Fig. 3, we have that almost all runs of $Run(p(1), \{q_\uparrow\})$ and $Run(p(1), \{p_\uparrow\})$ share the same pattern frequency vector F such that $F(q(+)) = 1$. Now let B be a BSCC of $X_\mathcal{A}$ which is *not* of the form $\{q_\uparrow\}$. As an example, consider the following BSCC of the chain $X_\mathcal{A}$ given in Fig. 3:

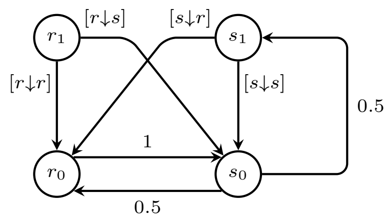

For all $r_1, s_0 \in B$ and all $t \in S$, let

- $E\langle r{\downarrow}s\rangle$ be the conditional expected length of a path from $r(1)$ to $s(0)$ under the condition that $s(0)$ is reached from $r(1)$ via a finite path where the counter stays positive in all configurations preceding $s(0)$;
- $E_{\#_t}\langle r{\downarrow}s\rangle$ be the conditional expected number of visits to a configuration with control state t along a path from $r(1)$ to $s(0)$ (where the visit to $s(0)$ does not count) under the condition that $s(0)$ is reached from $r(1)$ via a finite path where the counter stays positive in all configurations preceding $s(0)$;

If some $E\langle r{\downarrow}s\rangle$ is infinite, then the frequency of visits to configurations with zero counter is equal to zero for almost all runs of $Run(p(1), B)$. Further, there is a non-bounded BSCC D of $\mathcal{A}$ with zero trend such that $r \in D$, and one can easily show that almost all runs of $Run(p(1), B)$ share the same pattern frequency vector F where $F(s(+)) = \mu_D(s)$ for all $s \in D$, and $F(pat) = 0$ for all of the remaining patterns pat.

Now suppose that all $E\langle r{\downarrow}s\rangle$ are finite (which implies that all $E_{\#_t}\langle r{\downarrow}s\rangle$ are also finite). Recall that every transition of $X_{\mathcal{A}}$ represents a finite subpath of a run in $M_{\mathcal{A}}$. The expected length of a subpath represented by a transition of B is given by

$$E[L] \;=\; \sum_{s_0 \in B} \mu_B(s_0) \cdot 1 \;+\; \sum_{r_1 \in B} \mu_B(r_1) \cdot \sum_{r_1 \xrightarrow{x} s_0} x \cdot E\langle r{\downarrow}s\rangle$$

where μ_B is the invariant distribution of B. Similarly, we can define the expected number of visits to a configuration $t(k)$, where $k > 0$, along a subpath represented by a transition of B by

$$E[t(+)] \;=\; \sum_{r_1 \in B} \mu(r_1) \cdot \sum_{r_1 \xrightarrow{x} s_0} x \cdot E_{\#_t}\langle r{\downarrow}s\rangle \,.$$

The expected number of visits to a configuration $t(0)$ along a subpath represented by a transition of B (where the last configuration of a subpath does not count) is given by $E[t(0)] \;=\; \mu_B(t_0)$. It follows that almost all runs of $Run(p(1), B)$ share the same pattern frequency vector F where

$$F(t(+)) = \frac{E[t(+)]}{E[L]}, \qquad F(t(0)) = \frac{E[t(0)]}{E[L]}$$

for all $t \in S$.

To sum up, for each BSCC B of $X_{\mathcal{A}}$ we need to approximate the probability of $Run(p(1), B)$ and the associated pattern frequency vector up to a given relative error $\varepsilon > 0$. To achieve that, we need to approximate the following numbers up to a sufficient relative precision, which is determined by a straightforward error propagation analysis:

- the probabilities of the form $[p{\downarrow}q]$ and $[p{\uparrow}]$;
- the conditional expectations $E\langle r{\downarrow}s\rangle$ and $E_{\#_t}\langle r{\downarrow}s\rangle$

The algorithms that approximate the above values are non-trivial and have been obtained only recently. More detailed comments are given in the next subsections.

Finally, let us note that the genarel case when $\mathcal{A}$ has more than one diverging BSCC does not cause any major difficulties; for each diverging BSCC D, we construct a one-counter pVASS $\mathcal{A}_D$ where the other diverging BSCCs of $\mathcal{A}$ are modified so that their trend becomes negative. The analysis of $\mathcal{A}$ is thus reduced to the analysis of several one-counter pVASS where the above discussed method applies.

3.1 Approximating $[p{\downarrow}q]$ and $[p{\uparrow}]$

In this subsection we briefly indicate how to approximate the probabilities of the form $[p{\downarrow}q]$ and $[p{\uparrow}]$ up to a given relative error $\varepsilon > 0$.

Let $x_{\min}$ be the least positive transition probability in $M_{\mathcal{A}}$. It is easy to show that if $[p{\downarrow}q] > 0$, then $[p{\downarrow}q] > x_{\min}^{|Q|^3}$ (one can easily bound the length of a path from $p(1)$ to $q(0)$). Hence, it suffices to show how to approximate $[p{\downarrow}q]$ up to a given *absolute* error $\varepsilon > 0$.

The vector of all probabilities of the form $[p{\downarrow}q]$ is the least solution (in $[0,1]^{|Q|^2}$) of a simple system of recursive non-linear equations constructed as follows:

$$[p{\downarrow}q] \;=\; \sum_{p(1)\xrightarrow{x}q(0)} x \;+\; \sum_{p(1)\xrightarrow{x}t(1)} x\cdot[t{\downarrow}q] \;+\; \sum_{p(1)\xrightarrow{x}t(2)} x\cdot\sum_{s}[t{\downarrow}s]\cdot[s{\downarrow}r]$$

These equations are intuitive, and can be seen as a special case of the equations designed for a more general model of probabilistic pushdown automata [6, 7]. A solution to this system (even for pPDA) can be approximated by a decomposed Newton method [7] which produces one bit of precision per iteration after exponentially many initial iterations [9]. For one-counter pVASS, this method produces one bit of precision per iteration after *polynomially* many iterations. By implementing a careful rounding, a *polynomial time approximation algorithm* for $[p{\downarrow}q]$ was designed in [11].

Since $[p{\uparrow}] = 1 - \sum_{q\in Q}[p{\downarrow}q]$, we can easily approximate $[p{\uparrow}]$ up to an arbitrarily small *absolute* error in polynomial time. To approximate $[p{\uparrow}]$ up to a given *relative* error, we need to establish a reasonably large lower bound for a positive $[p{\uparrow}]$. Such a bound was obtained in [3] by designing and analyzing an appropriate martingale. More precisely, in [3] it was shown that if $[p{\uparrow}] > 0$, then one of the following possibilities holds:

- There is $q \in Q$ such that $[q{\uparrow}] = 1$ and $p(1)$ can reach a configuration $q(k)$ for some $k > 0$. In this case, $[p{\uparrow}] \geq x_{\min}^{|Q|^2}$.
- There is a BSCC D of $\mathcal{A}$ such that $t_D > 0$ and

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

3.2 Approximating $E\langle r{\downarrow}s\rangle$ and $E_{\#_t}\langle r{\downarrow}s\rangle$

The conditional expectations of the form $E\langle r{\downarrow}s\rangle$ satisfy a simple system of *linear* recursive equations constructed in the following way:

$$\begin{aligned} E\langle q{\downarrow}r\rangle \;=\; & \sum_{q(1)\xrightarrow{x}r(0)} \frac{x}{[q{\downarrow}r]} \;+\; \sum_{q(1)\xrightarrow{x}t(1)} \frac{x\cdot[t{\downarrow}r]}{[q{\downarrow}r]}\cdot(1+E\langle t{\downarrow}r\rangle) \\ & +\; \sum_{q(1)\xrightarrow{x}t(2)} x\cdot\sum_{s}\frac{[t{\downarrow}s]\cdot[s{\downarrow}r]}{[q{\downarrow}r]}\cdot(1+E\langle t{\downarrow}s\rangle+E\langle s{\downarrow}r\rangle) \end{aligned}$$

The only problem is that the coefficients are fractions of probabilities of the form $[p{\downarrow}q]$, which may take irrational values and cannot be computed precisely in general. Still, we can approximate these coefficients up to an arbitrarily small relative error in polynomial time by applying the results of the previous subsection. Hence, the very core of the problem is to determine a sufficient precision for these coefficients such that the approximated linear system still has a unique solution which approximates the vector of conditional expectations up to a given relative error. This was achieved in [3] by developing an upper bound on $E\langle r{\downarrow}s\rangle$, which was then used to analyze the condition number of the matrix of the linear system.

The same method is applicable also to $E_{\#_t}\langle r{\downarrow}s\rangle$ (the system of linear equation presented above must be slightly modified).

4 Pattern Frequency in Two-Counter pVASS

The analysis of pattern frequencies in two-counter pVASS imposes new difficulties that cannot be solved by the methods presented in Section 3. Still, the results achieved for one-counter pVASS are indispensable, because in certain situations, one of the two counters becomes "irrelevant", and then we proceed by constructing and analyzing an appropriate one-counter pVASS.

The results achieved in [4] for two-counter pVASS are formulated in the next theorem.

Theorem 2. *Let $p\boldsymbol{v}$ be an initial configuration of a* stable *two-counter pVASS $\mathcal{A}$. Then*

- $\mathcal{P}[\mathcal{F}{=}\bot] = 0$*;*
- *$\mathcal{F}$ is a discrete random variable;*
- *there are only finitely many vectors F such that $\mathcal{P}(\mathcal{F}{=}F) > 0$;*
- *these vectors and the associated probabilities can be effectively approximated up to an arbitrarily small absolute/relative error $\varepsilon > 0$.*

The condition of *stability* (explained below) can be checked in exponential time and guarantees that certain infinite-state Markov chains that are used to analyze the pattern frequencies of $\mathcal{A}$ are not null-recurrent.

Let $p(1,1)$ be an initial configuration of $\mathcal{A}$, and let $\mathcal{A}_1$ and $\mathcal{A}_2$ be one-counter pVASS obtained from $\mathcal{A}$ by preserving the first and the second counter, and abstracting the other counter into a *payoff* which is assigned to the respective rule of $\mathcal{A}_1$ and $\mathcal{A}_2$, respectively. The analysis of runs initiated in $p(1,1)$ must take into account the cases when one or both counters become bounded, one or both counters cannot be emptied to zero anymore, etc. For simplicity, let us assume that

(1) the set of all configurations reachable from $p(1,1)$ is a strongly connected component of $M_{\mathcal{A}}$, and for every $k \in \mathbb{N}$ there is a configuration reachable from $p(1,1)$ where both counters are larger than k;

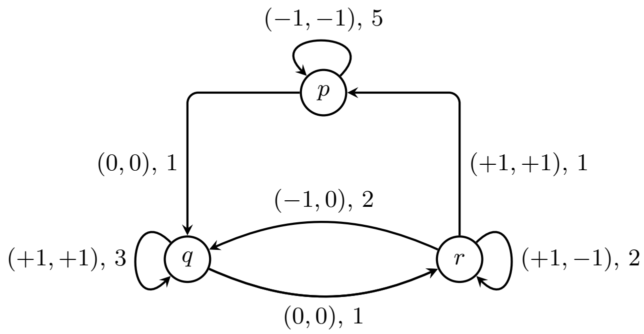

Fig. 4. A two-counter pVASS $\mathcal{A}$.

(2) the set of all configuration reachable from $p(1)$ in $\mathcal{A}_1$ and $\mathcal{A}_2$ is infinite and forms a strongly connected component of $M_{\mathcal{A}_1}$ and $M_{\mathcal{A}_2}$, respectively.

Note that Assumption (1) implies that the graph of $\mathcal{A}$ is strongly connected. An example of a two-counter pVASS satisfying these assumptions is given in Fig. 4. Now we define

- the *global trend* $t = (t_1, t_2)$, where t_i is the trend of $\mathcal{A}_i$ as defined by Equation (1) where $D = Q$ (recall that the graph of $\mathcal{A}_i$ is strongly connected);
- the *expected payoff* τ_i of a run initiated in a configuration of $\mathcal{A}_i$ reachable from $p(1)$. Since the set of all configurations reachable from $p(1)$ in $\mathcal{A}_i$ is strongly connected, it is not hard to show that τ_i is independent of the initial configuration and the mean payoff of a given run is equal to τ_i almost surely.

Intuitively, the global trend t specifies the average change of the counter values per transition if *both* counters are abstracted into payoffs. Note that if t is positive in both components, then almost all runs initiated in $p(1,1)$ "diverge", i.e., both counters remain positive from certain point on (here we need Assumption (1)). This means that almost all runs initiated in $p(1,1)$ share the same pattern frequency vector F where $F(q(+,+)) = \mu_{\mathcal{A}}$ and $F(pat) = 0$ for all of the remaining patterns pat (here $\mu_{\mathcal{A}}$ is the invariant distribution of $\mathcal{A}$; see the definition of μ_D in Equation (1) and recall that $\mathcal{A}$ is strongly connected).

Now suppose that t_2 is negative, and consider a configuration $q(k,0)$ reachable from $p(1,1)$, where k is "large". Obviously, a run initiated in $q(k,0)$ hits a configuration of the form $q'(k',0)$ without visiting a configuration with zero in the first counter with very high probability. Further, if $\tau_2 > 0$, then k' is larger than k "on average". Hence, the runs initiated in $q(k,0)$ will have a tendency to "diverge along the x-axis". If $\tau_2 < 0$, then k' is smaller than k on average, and the runs initiated in $q(k,0)$ will be moving towards the y-axis. A symmetric observation can be made when t_1 is negative. Hence, if *both* t_1 and t_2 are negative, we can distinguish three possibilities:

- $\tau_1 > 0$ and $\tau_2 > 0$. Then, almost all runs initiated in $p(1,1)$ will eventually diverge either along the x-axis or along the y-axis. That is, one of the counters eventually becomes irrelevant almost surely, and the pattern frequencies for $\mathcal{A}$ can be determined from the ones for $\mathcal{A}_1$ and $\mathcal{A}_2$ (since $\mathcal{A}_1$ and $\mathcal{A}_2$ are one-counter pVASS, we can apply the results of Section 3). The SPN of Fig. 2 is one concrete example of this scenario.
- $\tau_1 < 0$ and $\tau_2 > 0$. Then almost all runs initiated in $p(1,1)$ will eventually diverge along the x-axis. The case when $\tau_1 > 0$ and $\tau_2 < 0$ is symmetric.
- $\tau_1 < 0$ and $\tau_2 < 0$. In this case, there is a computable m such that the set of all configurations of the form $q(k,0)$ and $q(0,k)$, where $q \in Q$ and $k \leq m$, is a *finite eager attractor*. That is, this set of configurations is visited infinitely often by almost all runs initiated in $p(1,1)$, and the probability of revisiting this set in ℓ transitions decays (sub)exponentially in ℓ. The pattern frequencies for the runs initiated in $p(1,1)$ can be analyzed be generic methods for systems with a finite eager attractor developed in [1].

The cases when t_1, t_2, τ_1, or τ_2 is equal to zero are disregarded in [4], because the behaviour of $\mathcal{A}$, $\mathcal{A}_1$, or $\mathcal{A}_2$ can then exhibit strange effects caused by (possible) null recurrency of the underlying Markov chains, which requires different analytical methods (the *stability* condition in Theorem 2 requires that t_1, t_2, τ_1, and τ_2 are non-zero). We have not discussed the case when t_1 is negative and t_2 positive (or vice versa), because the underlying analysis is similar to the one presented above.

To capture the above explained intuition precisely, we need to develop an explicit lower bound for the probability of "diverging along the x-axis from a configuration $q(k,0)$" when $\tau_2 > 0$, examine the expected value of the second counter when hitting the y-axis by a run initiated in $q(k,0)$ when $\tau_2 < 0$, etc. These bounds are established in [4] by generalizing the martingales designed in [3] for one-counter pVASS.

5 Future Research

One open problem is to extend Theorem 2 so that it also covers two-counter pVASS that are not necessarily stable. Since one can easily construct a (non-stable) two-counter pVASS such that $\mathcal{P}(\mathcal{F}{=}F) > 0$ for infinitely many pairwise different vectors F, and there even exists a (non-stable) two-counter pVASS such that $\mathcal{P}(\mathcal{F}{=}\perp) = 1$, this task does not seem trivial.

Another challenge is to design algorithms for the analysis of long-run average properties in reasonably large subclasses of multi-dimensional pVASS. Such algorithms might be obtained by generalizing the ideas used for two-counter pVASS.

References

1. P. Abdulla, N.B. 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 Computer Society Press, 2006.
2. Marco Ajmone Marsan, Gianni Conte, and Gianfranco 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, S. Kiefer, and A. Kučera. Efficient analysis of probabilistic programs with an unbounded counter. *Journal of the Association for Computing Machinery*, 61(6), 2014.
4. T. Brázdil, S. Kiefer, A. Kučera, and P. Novotný. Long-run average behaviour of probabilistic vector addition systems. In *Proceedings of LICS 2015*, 2015.
5. T. Brázdil, S. Kiefer, A. Kučera, P. Novotný, and J.-P. Katoen. Zero-reachability in probabilistic multi-counter automata. In *Proceedings of CSL-LICS 2014*, pages 22:1–22:10. ACM Press, 2014.
6. J. Esparza, A. Kučera, and R. Mayr. Model-checking probabilistic pushdown automata. In *Proceedings of LICS 2004*, pages 12–21. IEEE Computer Society Press, 2004.
7. K. Etessami and M. Yannakakis. Recursive Markov chains, stochastic grammars, and monotone systems of nonlinear equations. *Journal of the Association for Computing Machinery*, 56, 2009.
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. S. Kiefer, M. Luttenberger, and J. Esparza. On the convergence of Newton's method for monotone systems of polynomial equations. In *Proceedings of STOC 2007*, pages 217–226. ACM Press, 2007.
10. Michael K. Molloy. Performance analysis using stochastic Petri nets. *IEEE Trans. Computers*, 31(9):913–917, 1982.
11. 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 *Lecture Notes in Computer Science*, pages 495–510. Springer, 2013.