

Multi-Dimensional Long-Run Average Problems for Vector Addition Systems with States

Krishnendu Chatterjee 

IST Austria, Klosterneuburg, Austria
<https://pub.ist.ac.at/~kchatterjee/>
krish.chat@ist.ac.at

Thomas A. Henzinger

IST Austria, Klosterneuburg, Austria
<http://pub.ist.ac.at/~tah/>
tah@ist.ac.at

Jan Otop 

University of Wrocław, Poland
<https://sites.google.com/a/cs.uni.wroc.pl/jan-otop/>
jotop@cs.uni.wroc.pl

Abstract

A vector addition system with states (VASS) consists of a finite set of states and counters. A transition changes the current state to the next state, and every counter is either incremented, or decremented, or left unchanged. A state and value for each counter is a configuration; and a computation is an infinite sequence of configurations with transitions between successive configurations. A probabilistic VASS consists of a VASS along with a probability distribution over the transitions for each state. Qualitative properties such as state and configuration reachability have been widely studied for VASS. In this work we consider multi-dimensional long-run average objectives for VASS and probabilistic VASS. For a counter, the cost of a configuration is the value of the counter; and the long-run average value of a computation for the counter is the long-run average of the costs of the configurations in the computation. The multi-dimensional long-run average problem given a VASS and a threshold value for each counter, asks whether there is a computation such that for each counter the long-run average value for the counter does not exceed the respective threshold. For probabilistic VASS, instead of the existence of a computation, we consider whether the expected long-run average value for each counter does not exceed the respective threshold. Our main results are as follows: we show that the multi-dimensional long-run average problem (a) is NP-complete for integer-valued VASS; (b) is undecidable for natural-valued VASS (i.e., nonnegative counters); and (c) can be solved in polynomial time for probabilistic integer-valued VASS, and probabilistic natural-valued VASS when all computations are non-terminating.

2012 ACM Subject Classification Theory of computation → Automata over infinite objects; Theory of computation → Quantitative automata

Keywords and phrases vector addition systems, mean-payoff, multidimension, probabilistic semantics

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2020.23

Related Version A full version of the paper is available at <http://arxiv.org/abs/2007.08917>.

Funding *Krishnendu Chatterjee*: The Austrian Science Fund (FWF) NFN grant S11407-N23 (RiSE/SHiNE).

Thomas A. Henzinger: The Austrian Science Fund (FWF) grants S11402-N23 (RiSE/ShiNE) and Z211-N23 (Wittgenstein Award).

Jan Otop: The National Science Centre (NCN), Poland under grant 2017/27/B/ST6/00299.

Acknowledgements We want to thank the anonymous reviewers for their helpful comments, which contributed to the final version of this paper.



© Krishnendu Chatterjee, Thomas A. Henzinger, and Jan Otop;
licensed under Creative Commons License CC-BY

31st International Conference on Concurrency Theory (CONCUR 2020).

Editors: Igor Konnov and Laura Kovács; Article No. 23; pp. 23:1–23:22

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Vector Addition System with States (VASS) and probabilistic VASS. *Vector Addition Systems (VASs)* provide a powerful framework for analysis of parallel processes [16]. They are equivalent to the well-studied model of Petri Nets [25]. The generalization of VASs with a finite-state transition system gives *Vector Addition Systems with States (VASS)*. The model of VASS is as follows: there is a finite set of control states with transitions between them, and a set of k counters, where at every transition between the control states each counter is either incremented, decremented, or remains unchanged. For a VASS, a *configuration* is a control state and a valuation of each counter, and the transitions of the VASS determines the transitions between the configurations. Thus a VASS is a finite description of an infinite-state transition system between the configurations. The class of VASS where the counters can hold all possible integer values, are referred to as integer-valued VASS; and the class of VASS where the counters can hold only non-negative values, are referred to as natural-valued VASS. A probabilistic VASS consists of a VASS along with probability distribution over the transitions for every state.

VASS Framework in Verification. VASS are an elegant mathematical framework for concurrent processes [16], and have been widely studied in performance analysis of concurrent processes [14, 20, 23, 24]. They have also been used in several other contexts, such as: (a) analysis of parametrized systems [3], (b) abstract models for programs for bounds analysis [34], (c) interactions between components of an API in component-based synthesis [18]. The probabilistic VASS provide a natural model for problems mentioned above with stochasticity in the system [6]. Thus VASS and probabilistic VASS provide a rich framework for many problems in verification and program analysis.

Previous results for VASS. A computation (run) in a VASS is an infinite sequence of configurations with transitions between successive configurations. The classical problems studied for VASS are as follows: (a) *control-state reachability* where given a set of target control states a computation satisfies the objective if a target state is reached; (b) *configuration reachability* where given a set of target configurations a computation satisfies the objective if a target configuration reached. For natural-valued VASS, (a) the control-state reachability problem is EXPSpace-complete: the EXPSpace-hardness is shown in [15, 30] and the upper bound follows from [33]; and (b) the configuration reachability problem is decidable [26, 27, 28, 31], and a recent breakthrough result establishes non-elementary hardness [13]. For integer-valued VASS, (a) the control-state reachability problem is NLOGSpace-complete (by reduction to graph reachability); (b) the configuration reachability problem is NP-complete. In probabilistic VASS, for the natural-valued class, even defining the probability measure over infinite computations is a challenging and complex problem [6], as computations that violate the non-negativity condition terminate as finite computations.

Long-run average objective and multi-dimensional long-run average problem. The classical problems for VASS consider qualitative (or Boolean) objectives where each computation is either satisfactory or not. In this work we consider multi-dimensional long-run average objective. For a counter, we consider the cost of a configuration as the value of the counter. For a computation, the long-run average of the costs of the configurations of the computation is the long-run average value for the respective counter. The multi-dimensional long-run average problem given a VASS and a threshold value for each counter, asks whether there is

a computation such that for each counter the long-run average value for the counter does not exceed the respective threshold. For integer-valued probabilistic VASS, instead of the existence of a computation, we consider whether the expected long-run average value for each counter does not exceed the respective threshold. For natural-valued probabilistic VASS, the presence of terminating runs makes even defining the probability measure complex. We consider two variants: (a) *strict semantics* that require all computations to be non-terminating; (b) *relaxed semantics* where we consider the conditional probability with respect to non-terminating runs.

Motivating examples. We present some motivating examples for the problems we consider. First, consider a VASS where the counters represent different queue lengths, and each queue consumes a resource type (e.g., energy or memory or time delay) proportional to its length. The multi-dimensional long-run average problem asks that the average consumption of each resource does not exceed a desired threshold. Second, consider a system that uses two different batteries, and the counters represent the charge levels. At different states, different batteries are used, and we are interested in the long-run average charge of each battery. This is again modeled as the multi-dimensional long-run average problem.

Our contributions. Our main contributions are as follows:

1. For non-probabilistic VASS we show that the multi-dimensional long-run average problem (a) is NP-complete for integer-valued VASS, and (b) is undecidable for natural-valued VASS.
2. For probabilistic integer-valued VASS, we show that the multi-dimensional long-run average problem can be solved in polynomial time. For natural-valued VASS, we show that the multi-dimensional problem can be solved in polynomial-time for (a) the strict semantics, and (b) the relaxed semantics for strongly connected VASS such that the expected multi-dimensional long-run average is finite. For the relaxed semantics and general natural-valued VASS, we show EXPSpace-hardness, and the exact decidability and complexity remain open.

Related works. For probabilistic VASS the long-run average behavior problem has been studied [6], as well as for other infinite-state models such as pushdown automata and games [1, 11, 12]. However, these works consider that costs are associated with the transitions of the finite-state system and do not depend on the counter values; moreover, they do not consider the multi-dimensional problem. In contrast, we consider costs that depend on the counter values, and hence on the configurations. Costs based on configurations, specifically the content of the stack in pushdown automata, have been considered in [32]. Quantitative asymptotic bounds for polynomial-time termination in VASS have also been studied [5, 29], however, these works do not consider long-run average property. Finally, a related model of automata with monitor counters with long-run average property have been considered in [7, 8]. However, there is a crucial difference: in automata with monitor counters, counters are reset once the value is used. Moreover, the complexity results for automata with monitor counters are quite different from the results we establish. Finally a recent work considers long-run average problem for VASS [9]. However, the cost is always single-dimensional with a linear combination of the counter values, and moreover, probabilistic VASS have not been considered in [9].

2 Preliminaries

For a sequence w , we define $w[i]$ as the $(i + 1)$ -th element of w (we start with 0) and $w[i, j]$ as the subsequence $w[i]w[i + 1] \dots w[j]$. We allow j to be ∞ for infinite sequences. For a finite sequence w , we denote by $|w|$ its length; and for an infinite sequence the length is ∞ . We use the same notation for vectors. For a vector $\vec{x} \in \mathbb{R}^k$ (resp., \mathbb{Q}^k , \mathbb{Z}^k or \mathbb{N}^k), we define $x[i]$ as the i -th component of \vec{x} .

2.1 Vector addition systems with states (VASS)

A k -dimensional vector addition system with states (VASS) over \mathbb{Z} (resp., over \mathbb{N}), referred to as $\text{VASS}(\mathbb{Z}, k)$ (resp., $\text{VASS}(\mathbb{N}, k)$), is a tuple $\mathcal{A} = \langle Q, Q_0, \delta \rangle$, where (1) Q is a finite set of states, (2) $Q_0 \subseteq Q$ is a set of initial states, and (3) $\delta \subseteq Q \times Q \times \mathbb{Z}^k$ is a transition relation. In a transition (q, q', \vec{y}) , the vector \vec{y} is called a *counter update* as we refer to k dimensions of a VASS as *counters*. We often omit the dimension in VASS and write $\text{VASS}(\mathbb{Z})$, $\text{VASS}(\mathbb{N})$ if a definition or an argument is uniform w.r.t. the dimension.

We define the size of a VASS in a standard way assuming binary encoding of counter updates. Formally, the size of a VASS $\langle Q, Q_0, \delta \rangle$ is defined as $|Q| + \sum_{(q, q', \vec{y}) \in \delta} \text{len}(\vec{y})$, where $\text{len}(\vec{y})$ is the length of the binary representation of \vec{y} .

Configurations and computations. A *configuration* of a $\text{VASS}(\mathbb{Z}, k)$ \mathcal{A} is a pair from $Q \times \mathbb{Z}^k$, which consists of a state and a valuation of the counters. A *computation* of \mathcal{A} is an infinite sequence π of configurations such that (a) $\pi[0] \in Q_0 \times \{\vec{0}\}$, and (b) for every $i \geq 0$, there exists $(q, q', \vec{y}) \in \delta$ such that $\pi[i] = (q, \vec{x})$ and $\pi[i + 1] = (q', \vec{x} + \vec{y})$. Note that, without loss of generality, we assume that the initial counter valuation is $\vec{0}$. We can encode any initial configuration in the VASS itself.

A computation of a $\text{VASS}(\mathbb{N}, k)$ \mathcal{A} is a computation π of \mathcal{A} considered as a $\text{VASS}(\mathbb{Z}, k)$ such that the values of all counters are non-negative, i.e., for all i we have $\pi[i] \in Q \times \mathbb{N}^k$. Transitions of a $\text{VASS}(\mathbb{N})$ that make the value of some counter negative are disabled.

We call a finite sequence ρ a *subcomputation* of a $\text{VASS}(\mathbb{Z}, k)$ (resp., $\text{VASS}(\mathbb{N}, k)$) \mathcal{A} , if it satisfies condition (b), i.e., all configurations are consistent with some transitions of \mathcal{A} , and all configurations belong to $Q \times \mathbb{Z}^k$ (resp., $Q \times \mathbb{N}^k$).

Paths and cycles. A path $\mathbf{p} = (q_0, q'_0, \vec{y}_0), (q_1, q'_1, \vec{y}_1), \dots$ in a $\text{VASS}(\mathbb{Z})$ (resp., $\text{VASS}(\mathbb{N})$) \mathcal{A} is a (finite or infinite) sequence of transitions (from δ) such that for all $0 \leq i < |\mathbf{p}|$ we have $q'_i = q_{i+1}$. A finite path \mathbf{p} is a cycle if $\mathbf{p} = (q_0, q'_0, \vec{y}_0), \dots, (q_m, q'_m, \vec{y}_m)$ and $q_0 = q'_m$. Every computation in a $\text{VASS}(\mathbb{Z})$ (resp., $\text{VASS}(\mathbb{N})$) corresponds to the unique infinite path. Conversely, every infinite path in a $\text{VASS}(\mathbb{Z})$ \mathcal{A} starting with $q_0 \in Q_0$ defines a computation in \mathcal{A} . However, if \mathcal{A} is a $\text{VASS}(\mathbb{N}, k)$, some paths do not correspond to valid computations due to non-negativity restriction posed on the counters.

Cycle characteristics. For a path \mathbf{p} we define $\text{GAIN}(\mathbf{p})$ as the vector of total counter change upon \mathbf{p} . Formally, for \mathbf{p} of length n with counter updates $\vec{y}_1, \dots, \vec{y}_n$ we define $\text{GAIN}(\mathbf{p}) = \sum_{i=1}^n \vec{y}_i$.

2.2 Probabilistic semantics

Markov chains. A *Markov chain* is a tuple $\langle \Sigma, Q, Q_0, \delta, P, \mu \rangle$ such that (1) Σ is a (finite) set of labels, (2) Q is a (finite) set of states, (3) Q_0 is a set of initial states, (4) $\delta \subseteq Q \times Q \times \Sigma$ is a transition relation, (5) $P: \delta \rightarrow (0, 1]$ is a probability distribution over transitions such that for every $s \in S$ we have $\sum_{s' \in S, a \in \Sigma} p(s, s', a) = 1$, and (6) $\mu: Q_0 \rightarrow [0, 1]$ is an initial distribution such that $\sum_{q \in Q_0} \mu(q) = 1$.

Probability measures defined by Markov chains. For a finite path \mathbf{p} in a Markov chain \mathcal{M} , we define the probability of \mathbf{p} , denoted by $\mathbb{P}_{\mathcal{M}}(\mathbf{p})$, as the product of probabilities of transitions along \mathbf{p} . For any $n > 0$, the probability $\mathbb{P}_{\mathcal{M}}(\cdot)$ is indeed a probability measure over paths of length n . We extend this probability measure to infinite paths in the standard fashion. Let X be the set of all infinite paths in \mathcal{M} . For a basic open set $\mathbf{p} \cdot X$, which is the set of all paths with the common prefix \mathbf{p} , we define $\mathbb{P}_{\mathcal{M}}(\mathbf{p} \cdot X) = \mathbb{P}_{\mathcal{M}}(\mathbf{p})$, and then the probability measure over infinite paths defined by \mathcal{M} is the unique extension of the above measure (by Carathéodory's extension theorem [17]). We will denote the unique probability measure defined by \mathcal{M} as $\mathbb{P}_{\mathcal{M}}$.

Probabilistic VASS. Probabilistic VASS generalize both VASS and Markov chains. A probabilistic VASS is a VASS, in which transitions are labeled with probabilities. It can be also considered to be an infinite-state Markov chain over the set of states $Q \times \mathbb{Z}^k$ (resp., $Q \times \mathbb{N}^k$) and Σ is a singleton. Formally, a probabilistic VASS is a tuple $\mathcal{A} = \langle Q, Q_0, \delta, P, \mu \rangle$ such that (1) $\langle Q, Q_0, \delta \rangle$ is a VASS (VASS(\mathbb{Z}) or VASS(\mathbb{N})), (2) $P: \delta \rightarrow (0, 1]$ is the probability distribution over transitions, which for every $q \in Q$ satisfies $\sum_{(q, q', \vec{y}) \in \delta} P(q, q', \vec{y}) = 1$, and (3) $\mu: Q_0 \rightarrow [0, 1]$ is the initial distribution, which satisfies $\sum_{q \in Q_0} \mu(q) = 1$.

Probability measures defined by probabilistic VASS. A probabilistic VASS(\mathbb{Z}) (resp. VASS(\mathbb{N})) defines the probability measure over its computations. First, a probabilistic VASS(\mathbb{Z}) (resp., VASS(\mathbb{N})) \mathcal{A} defines the probability measure over its infinite paths in the same way as a Markov chain does. In VASS(\mathbb{Z}), every path corresponds to a computation and hence the probability measure over infinite paths carries over to computations.

- We define $\mathbb{P}_{\mathcal{A}}$ as the probability measure on computations carried over from infinite paths.

However, in VASS(\mathbb{N}) some paths may not correspond to valid computations. For that reason, defining the probability measure over computations poses difficulties [6]. We consider two possible solutions: the *strict* and the *relaxed* semantics.

- Under the strict semantics, we require that all paths correspond to valid computations and then we define the probability measure $\mathbb{P}_{\mathcal{A}}^s$ over computations as in the VASS(\mathbb{Z}) case.
- Under the relaxed semantics, we require the set of paths corresponding to valid computations to have a non-zero probability, and we define the probability measure $\mathbb{P}_{\mathcal{A}}^r$ over computations as the conditional probability under the condition being the set of all paths that correspond to valid computations.

Random computations. To indicate that we consider a computation picked at random, we denote by ξ computations considered as random events.

► **Remark 1.** Under the strict semantics we require that every path corresponds to a valid computation, i.e., no counter gets a negative value. Note that relaxing *all* to *almost all* (i.e., with probability 1) gives us the same notion. Being a valid computation is a safety property and hence if the set of paths corresponding to valid computations has probability 1, then it is the set of all paths.

3 Problems

In this section, we define *the multi-dimensional average problem* and *the expected multi-dimensional average problem*, which we study in this paper. We define the averages over selected positions; averages are parametrized by a set of states S , called *selected states*, which determines meaningful configurations over which we compute the average, while skipping other configurations. This allows us to specify properties based on desired events (from S) rather than steps.

Averages and limit-averages over selecting states. Let $\mathcal{A} = \langle Q, Q_0, \delta \rangle$ be a $\text{VASS}(\mathbb{Z}, k)$ (resp., $\text{VASS}(\mathbb{N}, k)$) and $S \subseteq Q$ be a set of *selecting states*. Fix a counter $i \in \{1, \dots, k\}$. For a finite subcomputation ρ of \mathcal{A} , which contains at least one configuration from $S \times \mathbb{Z}^k$ (resp., $S \times \mathbb{N}^k$), we define the *average value of counter i (over S)*, denoted by $\text{AVG}_S^i(\rho)$, as the average over values of counter i over configurations with the state belonging to S , i.e., we first pick a subsequence $(s_1, \vec{x}_1), \dots, (s_m, \vec{x}_m)$ consisting of all configurations (s, \vec{x}) such that $s \in S$, and then take the average of the values of counter i : $\text{AVG}_S^i(\rho) = \frac{1}{m} \sum_{j=1}^m \vec{x}_j[i]$. If ρ has no configurations with states from S , then $\text{AVG}_S^i(\rho)$ is undefined. For an infinite sequence π of configurations, which contains infinitely many configurations from $S \times \mathbb{Z}^k$ (resp., $S \times \mathbb{N}^k$), we define the *limit-average value of the counter i (over S)*, denoted by $\text{LIMAVG}_S^i(\pi)$, as $\text{LIMAVG}_S^i(\pi) = \liminf_{k \rightarrow \infty} \text{AVG}_S^i(\pi[0, k-1])$. If π does not contain infinitely many configurations from $S \times \mathbb{Z}^k$ (resp., $S \times \mathbb{N}^k$), then $\text{LIMAVG}_S^i(\pi)$ is undefined.

Multi-dimensional averages and limit-averages over selecting states. We extend averages and limit-averages to multiple dimensions. Let $\vec{S} = (S[1], \dots, S[k])$ be a k -tuple of the subsets of Q . For a (finite) subcomputation ρ and an (infinite) computation π , we define

$$\begin{aligned} \text{AVG}_{\vec{S}}(\rho) &= (\text{AVG}_{S[1]}^1(\rho), \dots, \text{AVG}_{S[k]}^k(\rho)) \\ \text{LIMAVG}_{\vec{S}}(\pi) &= (\text{LIMAVG}_{S[1]}^1(\pi), \dots, \text{LIMAVG}_{S[k]}^k(\pi)) \end{aligned}$$

if all their components are defined. If any component of $\text{AVG}_{\vec{S}}(\rho)$ (resp., $\text{LIMAVG}_{\vec{S}}(\pi)$) is undefined, the whole vector is undefined.

► **Definition 2** (The multi-dimensional average problem for VASS). *Given a $\text{VASS}(\mathbb{N}, k)$ (resp., $\text{VASS}(\mathbb{Z}, k)$) \mathcal{A} , $\vec{S} \in (2^Q)^k$ and $\vec{\lambda} \in \mathbb{Q}^k$, the **(multi-dimensional) average problem** asks whether there exists a computation π such that $\text{LIMAVG}_{\vec{S}}(\pi)$ is defined and $\text{LIMAVG}_{\vec{S}}(\pi) \leq \vec{\lambda}$, i.e., the limit-averages of counter values over \vec{S} are component-wise bounded by $\vec{\lambda}$.*

Expected limit-averages over selecting states. Consider a probabilistic $\text{VASS}(\mathbb{Z}, k)$ (resp., $\text{VASS}(\mathbb{N}, k)$), which defines a probability measure $\mathbb{P}_{\mathcal{A}}$ (resp., $\mathbb{P}_{\mathcal{A}}^s$ or $\mathbb{P}_{\mathcal{A}}^r$) over its computations. Let $\vec{S} = (S[1], \dots, S[k])$ be a k -tuple of the subsets of Q . The function $\xi \mapsto \text{LIMAVG}_{S[i]}^i(\xi)$ is a random variable w.r.t. $\mathbb{P}_{\mathcal{A}}$ (resp., $\mathbb{P}_{\mathcal{A}}^s$ or $\mathbb{P}_{\mathcal{A}}^r$) and we define $\mathbb{E}_{\mathcal{A}}(\text{LIMAVG}_{S[i]}^i(\xi))$ as the expected value of this random variable. If the set of computations ξ , at which $\text{LIMAVG}_{S[i]}^i(\xi)$ is undefined, has a non-zero probability, then the expected value is undefined as well. We extend the expectation to vectors and define the expected multi-dimensional limit-average as

$$\mathbb{E}_{\mathcal{A}}(\text{LIMAVG}_{\vec{S}}) = (\mathbb{E}_{\mathcal{A}}(\text{LIMAVG}_{S[1]}^1), \dots, \mathbb{E}_{\mathcal{A}}(\text{LIMAVG}_{S[k]}^k)).$$

As above, the expected value $\mathbb{E}_{\mathcal{A}}(\text{LIMAVG}_{\vec{S}})$ is defined only if all components are defined.

► **Definition 3** (The expected (multi-dimensional) average problem for VASS). *Given a probabilistic VASS \mathcal{A} , $\vec{S} \in (2^Q)^k$, the **expected multi-dimensional average** problem asks to compute the expected limit-averages over \vec{S} , i.e., $\mathbb{E}_{\mathcal{A}}(\text{LIMAVG}_{\vec{S}})$.*

► **Remark 4.** In all complexity results for the multi-dimensional average and the expected multi-dimensional average problems, we consider VASS where the counter updates are encoded in binary.

4 Results on integer-valued VASS

4.1 The multi-dimension average problem

Consider a VASS (\mathbb{Z}, k) $\mathcal{A} = \langle Q, Q_0, \delta \rangle$, a vector $\vec{S} \in (2^Q)^k$ and thresholds $\vec{\lambda} \in \mathbb{Q}^k$. For simplicity, we assume that $Q_0 = \{q_0\}$ and hence $(q_0, \vec{0})$ is the initial configuration.

We present sufficient and necessary conditions for the existence of a computation π with $\text{LIMAVG}_{\vec{S}}(\pi) \leq \vec{\lambda}$. These conditions are expressed in terms of simple cycles in \mathcal{A} , i.e., they stipulate that for each counter i there exist (a) a simple cycle that can be *iterated* to ensure that limit average infimum is consistent with the threshold $\vec{\lambda}[i]$, and (b) a path to *access* this cycle, and then to switch back to another cycle. These conditions can be checked in NP. We present main ideas assuming that \mathcal{A} is strongly connected.

Assume that \mathcal{A} is strongly connected, i.e., it is strongly connected as a labeled graph. We distinguish two types of counters based on their behavior in \mathcal{A} : *bounded* and *unbounded*. We first assume that for every counter i there is a cycle \mathbf{c}_i such that iterating this cycle decreases this counter's value, i.e., $\text{GAIN}(\mathbf{c}_i)[i] < 0$. In such a case all counters are *unbounded* and for any $\vec{\lambda}$ there exists a computation π such that $\text{LIMAVG}_{\vec{S}}(\pi) \leq \vec{\lambda}$.

The all-unbounded case. We assume that all counters are unbounded. Fix some $\vec{\lambda} \in \mathbb{Q}^k$. We construct π such that $\text{LIMAVG}_{\vec{S}}(\pi) \leq \vec{\lambda}$ by interleaving strategies for each counter i to make its partial average below $\vec{\lambda}[i]$. More precisely, we define the path \mathbf{p} of the form

$$\mathbf{p} = \mathbf{s}_1^1 \mathbf{s}_2^1 \dots \mathbf{s}_k^1 \mathbf{s}_1^2 \dots \mathbf{s}_k^2 \dots$$

such that for every prefix $\mathbf{s}_1^1 \dots \mathbf{s}_i^j$ of \mathbf{p} , the subcomputation ρ_i^j corresponding to that prefix satisfies $\text{AVG}_{S[i]}(\rho_i^j) \leq \vec{\lambda}[i]$, i.e., the partial average over $S[i]$ is bounded by $\vec{\lambda}[i]$. We can construct such \mathbf{p} as follows. Suppose that a prefix of \mathbf{p} has been defined as above, and we need to construct \mathbf{s}_i^j . There are two cases: if the cycle \mathbf{c}_i with $\text{GAIN}(\mathbf{c}_i)[i] < 0$ contains a selecting state from $S[i]$, then $\mathbf{s}_i^j = (\mathbf{c}_i)^m$ for some large m , i.e., we iterate \mathbf{c}_i long enough such that the average of the whole prefix computation is below $\vec{\lambda}[i]$. If \mathbf{c}_i does not contain any state from $S[i]$, then there exists \mathbf{c}'_i that contains a selecting state and $\text{GAIN}(\mathbf{c}'_i)[i] < 0$. Indeed, let \mathbf{d} be a cycle from the initial state of \mathbf{c}_i to itself that contains a selecting state. Then, $\mathbf{d} \cdot \mathbf{c}_i^N$ contains a selecting state and $\text{GAIN}(\mathbf{d} \cdot \mathbf{c}_i^N)[i] < 0$ for some N .

Now, let π be the computation corresponding to \mathbf{p} . For every counter i there are infinitely many positions at which the partial average over $S[i]$ is at most $\vec{\lambda}[i]$ and hence $\text{LIMAVG}_{\vec{S}}(\pi) \leq \vec{\lambda}$.

The some-bounded case. Assume that for a counter j , there is no cycle such that iterating it decreases the value of counter j . In other words, for all cycles c we have $\text{GAIN}(c)[j] \geq 0$. We classify such a counter as *bounded*. It is clearly lower bounded and for the limit average to be finite it has to be upper bounded. In consequence, in any computation π with finite limit-average, all cycles c that occur infinitely often satisfy $\text{GAIN}(c)[j] = 0$. This in turn restricts the set of cycles that can appear infinitely often in the considered paths, which

makes other counters *bounded*. We iterate this process until we reach a fixed point \mathbf{B} , which is the set of all bounded counters. The complement of \mathbf{B} , denoted by \mathbf{U} , is the set of *unbounded* counters.

Note that for each unbounded counter $i \in \mathbf{U}$ there is a cycle \mathbf{c}_i such that:

- (U1) we have $\text{GAIN}(\mathbf{c}_i)[i] < 0$ and it contains a selecting state, and
- (U2) for each bounded counter $j \in \mathbf{B}$, we have $\text{GAIN}(\mathbf{c}_i)[j] = 0$.

It follows that similarly to the *all-unbounded* case, we can make sure that the partial averages of unbounded counters are arbitrarily low.

The limit average of a bounded counter depends on its initial value. Indeed, in the extreme case, if the value of a counter i does not change in any transition, then it is bounded and in every computation the limit average of counter i is precisely its initial value. However, to characterize cycles that witness low limit-averages of bounded counters it is more convenient to refer to a configuration that occurs infinitely often rather than the initial configuration. Therefore, we consider a *recurring configuration* (s_0, \vec{x}) that is: (a) reachable from the initial configuration $(q_0, \vec{0})$, (b) there are infinitely many configurations (s_0, \vec{y}) such that \vec{y} and \vec{x} agree on bounded counters. We now drop the strongly-connected assumption on \mathcal{A} .

Observe that switching between cycles for different (bounded or unbounded) counters may affect values of bounded counters. Therefore, for a bounded counter $i \in \mathbf{B}$ we require that there is a cycle \mathbf{c}_i , which (a) can be *accessed* with an appropriate path, and (b) its average together with the initial value are bounded by $\vec{\lambda}[i]$. To make it more precise: there exist a cycle \mathbf{c}_i and paths $\mathbf{in}_i, \mathbf{out}_i$ such that

- (B1) we have $\text{AVG}_{S[i]}(\rho) \leq \vec{\lambda}[i]$, where ρ is the subcomputation corresponding to the cycle \mathbf{c}_i starting from the configuration reached from (s_0, \vec{x}) over the path \mathbf{in}_i , and
- (B2) $\mathbf{in}_i, \mathbf{out}_i$ are from s_0 to some $s \in \mathbf{c}_i$ and from the same s to s_0 respectively, and for each bounded counter $j \in \mathbf{B}$, we have $\text{GAIN}(\mathbf{c}_i)[j] = 0$ and $\text{GAIN}(\mathbf{in}_i \mathbf{out}_i)[j] = 0$.

Finally, we require that for all unbounded counters $i \in \mathbf{U}$ there exist access paths $\mathbf{in}_i, \mathbf{out}_i$ as in condition (B2), i.e., paths $\mathbf{in}_i, \mathbf{out}_i$ satisfy:

- (U3) $\mathbf{in}_i, \mathbf{out}_i$ are from s_0 to some $s \in \mathbf{c}_i$ and from the same s to s_0 respectively, and for each bounded counter $j \in \mathbf{B}$, we have $\text{GAIN}(\mathbf{in}_i \mathbf{out}_i)[j] = 0$.

Condition (U3) is necessary as otherwise, switching between cycles for unbounded cycles and bounded cycles could change values of bounded counters. Observe that conditions (U2) and (U3) together are the same as (B2). We unify these conditions into a single one denoted (BU).

A witness for $\text{LimAvg}_{\vec{s}} \leq \vec{\lambda}$. A *witness* for $\text{LIMAVG}_{\vec{s}} \leq \vec{\lambda}$ is a tuple consisting of (a) a (recurring) configuration (s_0, \vec{x}) reachable from the initial configuration $(q_0, \vec{0})$, (b) a partition of counters into \mathbf{B} and \mathbf{U} , and (c) cycles \mathbf{c}_i and access paths $\mathbf{in}_i, \mathbf{out}_i$, for all i , which all satisfy conditions (U1), (B1) and (BU).

First, we show that the existence of a witness for $\text{LIMAVG}_{\vec{s}} \leq \vec{\lambda}$ is sufficient for the existence of a computation π with $\text{LIMAVG}_{\vec{s}}(\pi) \leq \vec{\lambda}$.

Key ideas. Using a witness, we construct a computation π satisfying $\text{LIMAVG}_{\vec{s}}(\pi) \leq \vec{\lambda}$ in a similar way as in the *all-unbounded* case. The only difference here is that we use access paths to switch between cycles for different counters so that we switch between the counters in the state s_0 , where the values of bounded counters are the same as in the configuration (s_0, \vec{x}) . Due to condition (BU), we do not require \mathcal{A} to be strongly connected. The full proof is relegated to the appendix. In consequence, we have the following:

► **Lemma 5.** *Let \mathcal{A} be a VASS(\mathbb{Z}). If it has a witness for $\text{LIMAVG}_{\vec{g}} \leq \vec{\lambda}$, then there exists a computation π such that $\text{LIMAVG}_{\vec{g}}(\pi) \leq \vec{\lambda}$.*

We show that the existence of a witness for $\text{LIMAVG}_{\vec{g}} \leq \vec{\lambda}$ is necessary for the existence of a computation π with $\text{LIMAVG}_{\vec{g}}(\pi) \leq \vec{\lambda}$. We present key ideas, while the full proof is relegated to the appendix.

Key ideas. Consider a computation π such that $\text{LIMAVG}_{\vec{g}}(\pi) \leq \vec{\lambda}$ and let \mathbf{p} be the infinite path corresponding to π . Let s_0 be a state that occurs infinitely often. We can decompose \mathbf{p} into simple cycles greedily always picking the first occurring simple cycle. Now, consider all simple cycles $D = \{\mathbf{d}_1, \dots, \mathbf{d}_m\}$ that occur infinitely often. Based on these cycles we define \mathbf{B} as the set of counters j such that for all cycles $\mathbf{d} \in D$ we have $\text{GAIN}(\mathbf{d})[j] = 0$, and $\mathbf{U} = \{1, \dots, k\} \setminus \mathbf{B}$. It follows that there is a configuration (s_0, \vec{x}) in π such that for all (s_0, \vec{y}) past that configuration \vec{x} and \vec{y} agree on the values of bounded counters. This follows from the fact that eventually all cycles from s_0 to itself in \mathbf{p} can be decomposed into cycles from D , which have zero gain for all bounded counters. We pick (s_0, \vec{x}) as the recurrent configuration in the witness.

Observe that for the limit-average to be finite, for every $i \in \mathbf{U}$ there has to be a simple cycle \mathbf{c}_i with $\text{GAIN}(\mathbf{c}_i)[i] < 0$. We can also enforce that \mathbf{c}_i contains a selecting state; such a cycle with a selecting state is of the form $\mathbf{f}\mathbf{c}_i^N$, where \mathbf{f}, \mathbf{c}_i are simple cycles, but $N = \text{GAIN}(\mathbf{f})[i]$ can be exponential. Still, this cycle can be represented polynomially by the pair $(\mathbf{f}, \mathbf{c}_i)$. Therefore, condition **(U1)** is satisfied. For counter $i \in \mathbf{B}$ there is a subsequence on which the limit average of i converges to a value less or equal to $\vec{\lambda}[i]$. It follows that there has to be a cycle from D satisfying **(B1)**. Otherwise, the limit averages of all computations exceed $\vec{\lambda}[i]$.

Each cycle \mathbf{d} from D is contained in some cycle $\tilde{\mathbf{d}}$ from s_0 to itself, i.e., $\tilde{\mathbf{d}} = \mathbf{in} \cdot \mathbf{d} \cdot \mathbf{out}$. We can pick such $\tilde{\mathbf{d}}$ that occurs infinitely often and hence can be decomposed into cycles $\mathbf{e}_1, \dots, \mathbf{e}_l$ from D , which by the definition of \mathbf{B} , satisfy $\text{GAIN}(\mathbf{e}_p)[j] = 0$ for each $j \in \mathbf{B}$. Therefore, paths $\mathbf{in}_i, \mathbf{out}_i$ for \mathbf{c}_i satisfy condition **(BU)**.

Finally, we need to show that if there is a witness, then there is a witness of polynomial size. Cycles \mathbf{c}_i (for $i \in \mathbf{B}$) as well as paths $\mathbf{in}_i, \mathbf{out}_i$ (for all i) can be picked to have polynomial length and for $i \in \mathbf{U}$ cycles \mathbf{c}_i can be picked to have polynomial length representation. Now, for the recurrent configuration (s_0, \vec{x}) , observe that the components of \vec{x} are relevant only for bounded counters, whose average is bounded by α , which has the binary representation of polynomial length in $|\mathcal{A}|$. Therefore, we need to find any reachable (s_0, \vec{x}) such that for all bounded i we have $\vec{x}[i]$ is less than $-\alpha + \vec{\lambda}[i]$. Furthermore, reachable configurations in VASS(\mathbb{Z}) are semilinear sets [4] defined by equations of polynomial size in $|\mathcal{A}|$ and hence if there is any configuration (s_0, \vec{x}) then there is also one with the binary representation of polynomial length in $|\mathcal{A}| + |\vec{\lambda}|$. In consequence, we have:

► **Lemma 6.** *For all VASS(\mathbb{Z}, k) \mathcal{A} and $\vec{\lambda} \in \mathbb{Q}^k$ the following holds: if there is a computation π such that $\text{LIMAVG}_{\vec{g}}(\pi) \leq \vec{\lambda}$, then there exists a witness for $\text{LIMAVG}_{\vec{g}} \leq \vec{\lambda}$, which has a polynomial size in $|\mathcal{A}| + |\vec{\lambda}|$.*

Finally, a polynomial-size witness for $\text{LIMAVG}_{\vec{g}}(\pi) \leq \vec{\lambda}$ can be non-deterministically picked and verified in polynomial time. More precisely, in the definition of a witness for $\text{LIMAVG}_{\vec{g}} \leq \vec{\lambda}$ condition (a) can be checked in NP as reachability for VASS(\mathbb{Z}) is NP-complete [4], and conditions (b) and (c) can be checked in polynomial time. In consequence, we have:

► **Lemma 7.** *The multi-dimensional average problem for VASS(\mathbb{Z}) is in NP.*

For hardness of the multi-dimensional average problem, consider configuration-reachability for VASS(\mathbb{Z}), which is NP-complete. Configuration-reachability is mutually reducible to coverability for VASS(\mathbb{Z}) [21], which in turn is equivalent to dual coverability, i.e, the problem, given a VASS(\mathbb{Z}) and two configurations $(s, \vec{0})$ and (t, \vec{x}) , decide whether there is a (finite) subcomputation from $(s, \vec{0})$ to some (t, \vec{y}) , where $\vec{y} \leq \vec{x}$. The dual coverability straightforwardly reduces to the multi-dimensional average problem as follows. We construct \mathcal{A}' from \mathcal{A} by adding a fresh state t^* and two transitions labeled with $\vec{0}$: from t to t^* and a self-loop over t^* . Observe that there is a subcomputation from $(s, \vec{0})$ to (t, \vec{y}) where $\vec{y} \leq \vec{x}$ in \mathcal{A} if and only if there is a computation from $(s, \vec{0})$ that eventually reaches t^* and the multi-dimensional limit-averages are bounded by \vec{x} . To enforce that a computation eventually reaches t^* , we use an additional counter that is 0 in the configurations of \mathcal{A} and it changes to -1 upon moving to t^* . Requirement that the limit average of this counter is less or equal to -1 forces the computation to move to t^* . In consequence, the dual coverability for VASS(\mathbb{Z}) reduces to the multi-dimensional average problem for VASS(\mathbb{Z}) and hence the latter problem is NP-complete.

► **Theorem 8.** *The multi-dimensional average problem for VASS(\mathbb{Z}) is NP-complete.*

4.2 The expected average problem

Observe that the expected average problem for probabilistic VASS(\mathbb{Z}) is modular and each dimension can be considered separately. This follows from the fact that each path in a VASS(\mathbb{Z}) corresponds to a computation, which is not the case for VASS(\mathbb{N}). Furthermore, in this problem we compute the expected value over all computations and hence it can be considered for each dimension separately. Therefore, we consider VASS that are single-dimensional. We first discuss the strongly-connected case and then generalize our results to all VASS(\mathbb{Z}).

4.2.1 The strongly-connected case

Let \mathcal{A} be a single-dimensional probabilistic VASS($\mathbb{Z}, 1$), which is strongly connected as a labeled graph. We additionally assume that it has a single initial configuration $(q_0, 0)$. The case of any initial distribution follows easily. First, we define the expected gain of \mathcal{A} , which corresponds to the expected trend of the counter.

The expected gain. The graph of \mathcal{A} can be considered as a Markov chain and using standard methods we compute for each state q its long-run frequency x_q [2, 19]. More precisely, the *frequency* of q in a subcomputation $\xi[1, n]$ is the number of configurations with the state q in $\xi[1, n]$ divided by n . The Ergodic Theorem for Markov chains implies that with probability 1 over a random computation ξ , for every state q , the frequency of q in $\xi[1, n]$ converges to x_q as n tends to infinity. Based on frequencies x_q we define *the expected gain* $\mathbb{E}(\text{Gain})$ as the expected counter update provided that the initial state q is picked at random according to the frequencies x_q and the outgoing transition is picked at random according to the distribution at q , that is:

$$\mathbb{E}(\text{Gain}) = \sum_{(q, q', y) \in \delta} x_q \cdot P(q, q', y) \cdot y$$

The classification based on $\mathbb{E}(\text{Gain})$. We show that if $\mathbb{E}(\text{Gain})$ is positive (resp., negative), then the limit-average is infinite (resp. minus infinity). However, if the expected gain is zero there are two cases based on boundedness of configurations. Either the gain of every cycle is actually zero, or cycles with a positive gain balance cycles with a negative gain so that the expected gain is zero. We discuss these cases below.

We say that a VASS($\mathbb{Z}, 1$) is *totally bounded* if the gain of each cycle is zero. This property does not depend on the probability distribution over transitions and we extend it straightforwardly to probabilistic VASS($\mathbb{Z}, 1$). Note that if a VASS($\mathbb{Z}, 1$) is strongly connected and totally bounded, then in each reachable configuration, the state uniquely determines the counter's value. Otherwise, there exists a cycle with a non-zero gain. This observation allows us to reduce the expected limit-average problem for such VASS to computing the expected long-run reward for Markov chains [2, Chapter 10.5].

Consider a VASS($\mathbb{Z}, 1$) with $\mathbb{E}(\text{Gain})$ being zero and at least one cycle with a non-zero gain. Observe that $\mathbb{E}(\text{Gain})$ being 0 implies that there is at least one cycle with a positive gain and a cycle with a negative gain. Let us consider the simplest probabilistic VASS, which has a single state q_0 and two self loops labeled with 1 and -1 , both with probability 0.5. The distribution of the counter's gain in n transitions, denoted S_n , is related to the binomial distribution $B(n, 0.5)$ in the following way: $S_n \sim 2 \cdot B(n, 0.5) - n$. It follows that with probability 1 over a random computation ξ , the counter in ξ is neither lower nor upper bounded. Furthermore, we show that with probability 1, a random computation ξ has two subsequences such that the averages on one sequence tend to ∞ , and on the other tend to $-\infty$. To state this formally we define:

$$\text{LIMAVGINF}_S(\pi) = \liminf_{k \rightarrow \infty} \text{AVG}_S(\rho[1, k])$$

$$\text{LIMAVGSUP}_S(\pi) = \limsup_{k \rightarrow \infty} \text{AVG}_S(\rho[1, k])$$

Now, we present the lemma summarizing the above discussion.

► **Lemma 9.** *Let \mathcal{A} be a strongly-connected probabilistic VASS($\mathbb{Z}, 1$). One of the following conditions holds:*

- (1) $\mathbb{E}(\text{Gain}) > 0$, and $\text{LIMAVGINF}_S(\xi) = \text{LIMAVGSUP}_S(\xi) = \infty$ with probability 1 (over ξ),
- (2) $\mathbb{E}(\text{Gain}) < 0$, and $\text{LIMAVGINF}_S(\xi) = \text{LIMAVGSUP}_S(\xi) = -\infty$ with probability 1,
- (3) \mathcal{A} is totally bounded, and for some $x \in \mathbb{Q}$, with probability 1 over ξ we have $\text{LIMAVGINF}_S(\xi) = \text{LIMAVGSUP}_S(\xi) = x$, and
- (4) $\mathbb{E}(\text{Gain}) = 0$, \mathcal{A} is not totally bounded, and with probability 1 over ξ we have $\text{LIMAVGINF}_S(\xi) = -\infty$ and $\text{LIMAVGSUP}_S(\xi) = \infty$.

Proof's sketch. To show conditions (1) and (2) observe that the Ergodic Theorem for Markov chains implies that with probability 1 (over ξ) the value of the counter in configuration $\xi[n]$ divided by n converges to $\mathbb{E}(\text{Gain})$ as n tends to infinity and hence the average counter value converges to ∞ if $\mathbb{E}(\text{Gain}) > 0$ and it converges to $-\infty$ if $\mathbb{E}(\text{Gain}) < 0$.

Condition (3) follows from the above discussion and the fact that in a strongly-connected Markov chain there is a value x such that the long-run reward is almost surely x .

We present the full proof of conditions (1), (2) and (3) in the appendix. Now, we focus on condition (4). We only show that $\text{LIMAVGINF}_S(\xi) = -\infty$ holds with probability 1 over ξ , as the proof of $\text{LIMAVGSUP}_S(\xi) = \infty$ is symmetric. Observe that in a strongly-connected VASS(\mathbb{Z}), the event $\text{LIMAVGINF}_S(\xi) = -\infty$ is a tail event. Therefore, due to Kolmogorov's 0-1 law [17] it has either probability 0 or 1. In consequence, it suffices to show that it has a positive probability.

First, we show that with a positive probability $\text{LIMAVGINF}_S(\xi)$ is upper bounded.

23:12 Multi-Dimensional Long-Run Average Problems for VASS

► **Lemma 10.** *Consider a probabilistic VASS(\mathbb{Z}) \mathcal{A} as in (4) of Lemma 9. There exist $c \in \mathbb{Q}$ and $\delta > 0$ such that $\text{LIMAVGINF}_S(\xi) < c$ holds with probability greater than δ .*

Proof's ideas. We observe that $\text{LIMAVGINF}_S(\xi) = \infty$ implies that the values of the counter converge to infinity, which implies $\mathbb{E}(\text{Gain}) > 0$. The full proof is relegated to the appendix. ◀

For $c \in \mathbb{Q}$, we define X_c as the set of computations π such that $\text{LIMAVGINF}_S(\pi) < c$. Lemma 10 states that there are $c \in \mathbb{Q}$ and $\delta > 0$ such that $\mathbb{P}(X_c) = \delta$. We show that for every $d \in \mathbb{Q}$, $\text{LIMAVGINF}_S(\xi) < d$ holds with probability at least δ .

► **Lemma 11.** *Consider a probabilistic VASS(\mathbb{Z}) \mathcal{A} as in (4) of Lemma 9. Assume that $\mathbb{P}(X_c) = \delta > 0$. Then, for every d we have $\mathbb{P}(\text{LIMAVGINF}_S(\xi) < d) \geq \delta$.*

Proof's ideas. The main idea is to prepend to computations from X_c a subcomputation that decreases the initial counter's value to a . Then, the limit infimum of averages is $c + a$. Furthermore, we show that the set of such finite paths has probability 1.

More precisely, consider a subcomputation ρ from $(q_0, 0)$ to (q_0, a) and $\pi \in X_c$. We define the *join* of ρ and π , denoted by $\rho \bowtie \pi$, as the computation consisting of first ρ and then $\pi[1, \infty]$ (π with the first configuration removed) with a added to the counter of all following configurations of π . Observe that the join of ρ and π is indeed a computation. Moreover, the influence of the average of ρ on the whole computation diminishes and hence $\text{LIMAVGINF}_S(\rho \bowtie \pi) = \text{LIMAVGINF}_S(\pi) + a < c + a$.

Let Y be the set of (finite) subcomputations that start in $(q_0, 0)$ and terminate once they reach some configuration (q_0, b) where $b < d - c$. Since \mathcal{A} is strongly connected and not totally bounded, almost surely a random computation ξ reaches a configuration (q_0, b) where $b < d - c$. It follows that the set of all computations extending some subcomputation from Y has probability 1. Therefore, the set of all joins of subcomputations from Y with computations from X_c has probability at least δ and all such computations $\rho \bowtie \pi$ satisfy $\text{LIMAVGINF}_S(\rho \bowtie \pi) < d$, and hence Lemma 11 follows. ◀

Lemma 11 implies that the set of computations ξ such that $\text{LIMAVGINF}_S(\xi) = -\infty$ has a positive probability. Since $\text{LIMAVGINF}_S(\xi) = -\infty$ is a tail event in a strongly-connected VASS(\mathbb{Z}), Kolmogorov's 0-1 law [17] implies that its probability is 1, which concludes the proof of Lemma 9. ◀

Lemma 9 implies the following:

► **Lemma 12.** *The expected average problem for strongly-connected probabilistic VASS($\mathbb{Z}, 1$) can be solved in polynomial time.*

Proof's ideas. Consider a strongly-connected probabilistic VASS($\mathbb{Z}, 1$) \mathcal{A} . We can compute frequencies x_q of states of \mathcal{A} in polynomial time using standard methods [2, Chapter 10.5]. Having frequencies x_q , we can compute the expected gain $\mathbb{E}(\text{Gain})$ of \mathcal{A} in polynomial time from the definition.

Assume that $\mathbb{E}(\text{Gain}) = 0$. We can check whether \mathcal{A} is not totally bounded by checking whether it has a cycle with a non-zero gain, which can be done in polynomial time. Finally, if it is totally bounded, then each state of \mathcal{A} uniquely determines the value of each counter, and we can eliminate the counters and label states with counter values. Therefore, the problem of computing the long-run average of almost all computations, denoted by x , reduces to computing the expected long-run reward a Markov chain with rewards, which can be done in polynomial time [2, Chapter 10.5]. In consequence, we can check all the conditions of Lemma 9 in polynomial time and hence the result follows. ◀

4.2.2 The general case

Let \mathcal{A} be a probabilistic VASS $(\mathbb{Z}, 1)$. We show how to compute its expected limit-average in polynomial time. We identify all *bottom* SCCs (BSCCs) of \mathcal{A} (where an SCC B is *bottom* if all states reachable from B belong to B). If there is a BSCC that does not contain a state from S , then the expected limit-average is undefined. Assume that every BSCC contains a state from S and consider the following cases:

- If there are two BSCCs: (a) one with a positive expected gain, and (b) the other with a negative expected gain, then the expected limit average is undefined. The expected value is undefined for random variables that attain $+\infty$ and $-\infty$ with a positive probability [17].
- If there is a BSCC with a positive gain and every BSCC has (a) a positive gain, or (b) it has the zero gain and it is totally bounded, then the expected limit-average is ∞ .
- If there is a BSCC with (a) a negative gain, or (b) the zero gain and not totally bounded, and every BSCC has a non-positive gain, then the expected limit-average is $-\infty$.
- If all BSCCs have the zero gain and are totally bounded, the expected limit-average is finite and we discuss below how to compute it.

First, we compute all BSCCs B_1, \dots, B_m of \mathcal{A} . We pick in each of these components an initial state q_0^i . For each BSCC B_i with its initial configuration $(q_0^i, 0)$, we compute x_i , which is the expected limit-average in B_i . As we observed before, if we join a subcomputation from $(q_0, 0)$ to (q_0^i, y_i) and some computation from $(q_0^i, 0)$ with the limit-average x_i , then the limit-average of the resulting computation is $x_i + y_i$. Therefore, for each state q_0^i we compute the probability of reaching that state from the initial distribution, denoted p_i , and the expected counter's value y_i upon reaching q_0^i , i.e., the conditional expected counter's value under the condition that the state q_0^i is reached. Probabilities p_i can be computed using standard methods for Markov chains [2, Chapter 10.1]. The values y_i can be computed as well using standard methods for Markov chains with rewards [2, Chapter 10.5]. Observe that the expected limit-average of \mathcal{A} is given by the following formula:

$$\mathbb{E}_{\mathcal{A}}(\text{LIMAVG}_S) = \sum_{i=1}^m p_i \cdot (y_i + x_i)$$

Finally, as we discussed above, we can compute the expected value for each counter separately. In consequence we have the following:

► **Theorem 13.** *The expected average problem for probabilistic VASS (\mathbb{Z}) can be solved in polynomial time.*

5 Results on natural-valued VASS

5.1 The average problem in a single dimension

We first study the average problem for single-dimensional VASS $(\mathbb{N}, 1)$. For the lower bound observe that the reachability problem for VASS $(\mathbb{N}, 1)$, which is NP-complete [22], reduces to the average problem for VASS $(\mathbb{N}, 1)$. The reduction is straightforward and hence we omit it. To show the NP upper bound, we show the following:

► **Lemma 14.** *For all VASS $(\mathbb{N}, 1)$ \mathcal{A} the following holds: there exists a computation π with $\text{LIMAVG}_S(\pi) \leq \lambda$ if and only if there exist subcomputations ρ_0, ρ_c such that:*

- ρ_0 is from $(q_0, 0)$ to (s, x) , where $x \leq \lambda$, and
- ρ_c is a cycle from (s, x) to itself satisfying the following conditions:

- (a) $\text{AVG}_S(\rho_c) \leq \lambda$,
- (b) the number of configurations with selecting states in ρ_c , i.e., configurations from $S \times \mathbb{N}$, is $O(|S| \cdot |\lambda|^2)$, and
- (c) the value of the counter in each configuration of ρ_c from $S \times \mathbb{N}$ is $O(|S| \cdot |\lambda|^2)$.

We present the proof's sketch, while the full proof is relegated to the appendix.

Proof's sketch. Observe that having ρ_0, ρ_c as above, the computation $\pi = \rho_0(\rho_c)^\infty$ is a valid computation and it satisfies $\text{LIMAVG}_S(\pi) \leq \lambda$. Conversely, we observe that there has to be a cycle ρ in π with the average $\text{AVG}_S(\rho) \leq \lambda$. Furthermore, we pick a cycle of minimal length and observe that it satisfies all the conditions of the lemma, i.e., if some condition is violated, then the cycle can be shortened. \blacktriangleleft

We can check in NP whether there exist ρ_0, ρ_c satisfying the conditions from Lemma 14.

Key ideas. We non-deterministically pick all selecting configurations $(s_1, x_1), \dots, (s_m, x_m)$ from ρ_c . Then, we check reachability from $(q_0, 0)$ to (s_1, x_1) , and for each $i < m$ reachability over non-selecting configurations from (s_i, x_i) to (s_{i+1}, x_{i+1}) , and from (s_m, x_m) to (s_1, x_1) . All these reachability checks can be done in NP. Finally, we check $\frac{1}{m} \sum_{i=1}^m x_i \leq \lambda$. All these checks can be done in NP. The number of configurations m as well as the size of each configuration is polynomially bounded due to Lemma 14. In consequence, we have:

► **Theorem 15.** *The average problem for $\text{VASS}(\mathbb{N}, 1)$ is NP-complete.*

5.2 The multi-dimension average problem

We show that the (decision variant of the) multi-dimensional average problem for $\text{VASS}(\mathbb{N})$ is undecidable. A related problem, called the average-value problem, has been studied in [9]. In that problem, the values of all counters in each configuration (q, \vec{x}) are aggregated into a single number, called *the cost*, by computing dot-product of \vec{x} and a cost vector $\vec{c}_q \in \mathbb{N}^k$. A cost vector \vec{c}_q depends on the state q in the configuration. The average-value problem asks whether there exists a computation such that the limit average of costs is less or equal to a threshold λ . The problem for $\text{VASS}(\mathbb{N})$ with threshold 0 is undecidable [9, Theorem 24]. Threshold 0 in $\text{VASS}(\mathbb{N})$ means that whenever a cost vector is non-zero at component i (i.e., $\vec{c}_q[i] \neq 0$), then counter's i value should be 0. This constraint is expressed in the multi-dimensional average problem and hence we have:

► **Theorem 16.** *The decision variant of the multi-dimensional average problem for $\text{VASS}(\mathbb{N})$ is undecidable.*

5.3 The expected multi-dimensional average problem

We first study probabilistic $\text{VASS}(\mathbb{N})$ under the strict semantics and give the precise complexity. Next, we consider the relaxed semantics, where we have the exact complexity in the strongly-connected case and a hardness result in the general case.

5.3.1 Probabilistic natural-valued VASS under the strict semantics

Consider a probabilistic $\text{VASS}(\mathbb{N})$ \mathcal{A} under the strict semantics. To check whether every path of \mathcal{A} corresponds to a valid computation, we examine each counter i separately and check whether it can reach a negative value from some initial configuration. This can be done in polynomial time with the standard reachability analysis. If a negative value for some

counter is reachable, then the expected limit-average is undefined under the strict semantics. Otherwise, every path in \mathcal{A} corresponds to a valid computation and we can consider \mathcal{A} as a $\text{VASS}(\mathbb{Z})$ as the non-negativity restriction is vacuous for \mathcal{A} . Therefore, we apply Theorem 13 and compute the expected value for \mathcal{A} . In consequence, we have:

► **Theorem 17.** *The expected average problem for probabilistic $\text{VASS}(\mathbb{N})$ under the strict semantics can be solved in polynomial time.*

5.3.2 Probabilistic natural-valued VASS under the relaxed semantics

The finite strongly-connected case. Consider a probabilistic $\text{VASS}(\mathbb{N})$ \mathcal{A} , which is strongly connected. We show that if the expected limit-average is finite, then the strict and the relaxed semantics coincide. We first assume that \mathcal{A} is single-dimensional. Using the classification from Lemma 9 applied to \mathcal{A} considered as a $\text{VASS}(\mathbb{Z}, 1)$, we observe that:

- If $\mathbb{E}(\text{Gain}) < 0$ or $\mathbb{E}(\text{Gain}) = 0$ and \mathcal{A} is not totally bounded, then a random computation ξ (under the $\text{VASS}(\mathbb{Z})$ semantics) satisfies $\text{LIMAVGINF}_S(\xi) = -\infty$, and hence it is not a valid computation of the $\text{VASS}(\mathbb{N})$. Therefore, the expected limit-average under the relaxed semantics is undefined for \mathcal{A} .
- If $\mathbb{E}(\text{Gain}) > 0$, then $\text{LIMAVGINF}_S(\xi) = \infty$. Therefore, if the set of random computations ξ (under the $\text{VASS}(\mathbb{Z})$ semantics) that are also valid computations under the $\text{VASS}(\mathbb{N})$ semantics has a positive probability, then the expected limit-average under the relaxed semantics is defined and infinite. Otherwise, it is undefined.
- If $\mathbb{E}(\text{Gain}) = 0$ and \mathcal{A} is totally bounded, then (as we observe in Section 4.2) in each configuration, the state uniquely determines the counters value. Therefore, we can check whether counter values in all states are non-negative. If this is the case, then all paths correspond to valid computations, the expected limit-average is defined and finite, and we can compute it with Theorem 17. Otherwise, observe that in a strongly-connected \mathcal{A} every state is visited with probability 1 and hence the expected value is undefined.

Therefore, for the expected limit-average under the relaxed semantics to be defined and finite, the expected gain w.r.t. every counter has to be 0 and it has to be totally bounded. Furthermore, we check for each counter independently whether every path corresponds to a valid computation in $\text{VASS}(\mathbb{N})$. Since we consider all paths, we can make these checks independently for all counters. In consequence we have the following:

► **Theorem 18.** *Deciding whether the expected limit-average is defined and finite over strongly-connected probabilistic $\text{VASS}(\mathbb{N})$ under the relaxed semantics can be solved in polynomial time. Furthermore, if it is it can be computed in polynomial time.*

The general case. We present only a hardness result. The coverability problem for $\text{VASS}(\mathbb{N})$, which is EXPSpace -complete [30], reduces to (the decision version of) the expected limit-average problem for $\text{VASS}(\mathbb{N})$. The reduction is rather straightforward with minor technical difficulties (we need to ensure that the expected value of each counter is finite). In consequence, we have:

► **Theorem 19.** *The problem, given a probabilistic $\text{VASS}(\mathbb{N}, k)$ \mathcal{A} under the relaxed semantics, $S \subseteq Q$ and $\vec{x} \in \mathbb{Q}^k$, decide whether $\mathbb{E}_{\mathcal{A}}(\text{LIMAVG}_{\vec{S}}) < \vec{x}$ is EXPSpace -hard.*

The proof has been relegated to the appendix.

► **Remark 20.** The decidability of the problem from Theorem 19 is open.

References

- 1 Parosh Aziz Abdulla, Mohamed Faouzi Atig, Piotr Hofman, Richard Mayr, K. Narayan Kumar, and Patrick Totzke. Infinite-state energy games. In *CSL-LICS 2014*, pages 7:1–7:10, 2014. doi:10.1145/2603088.2603100.
- 2 Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT Press, 2008.
- 3 Roderick Bloem, Swen Jacobs, Ayrat Khalimov, Igor Konnov, Sasha Rubin, Helmut Veith, and Josef Widder. Decidability in parameterized verification. *SIGACT News*, 47(2):53–64, 2016. doi:10.1145/2951860.2951873.
- 4 Michael Blondin, Alain Finkel, Stefan Göller, Christoph Haase, and Pierre McKenzie. Reachability in two-dimensional vector addition systems with states is pspace-complete. In *LICS 2015*, pages 32–43. IEEE Computer Society, 2015. doi:10.1109/LICS.2015.14.
- 5 Tomás Brázdil, Krishnendu Chatterjee, Antonín Kucera, Petr Novotný, Dominik Velan, and Florian Zuleger. Efficient algorithms for asymptotic bounds on termination time in VASS. In *LICS 2018*, pages 185–194, 2018. doi:10.1145/3209108.3209191.
- 6 Tomás Brázdil, Stefan Kiefer, Antonín Kucera, and Petr Novotný. Long-run average behaviour of probabilistic vector addition systems. In *LICS 2015*, pages 44–55, 2015. doi:10.1109/LICS.2015.15.
- 7 Krishnendu Chatterjee, Thomas A. Henzinger, and Jan Otop. Nested weighted limit-average automata of bounded width. In *MFCS 2016*, pages 24:1–24:14, 2016. doi:10.4230/LIPIcs.MFCS.2016.24.
- 8 Krishnendu Chatterjee, Thomas A. Henzinger, and Jan Otop. Quantitative monitor automata. In *SAS 2016*, pages 23–38, 2016. doi:10.1007/978-3-662-53413-7_2.
- 9 Krishnendu Chatterjee, Thomas A. Henzinger, and Jan Otop. Long-run average behavior of vector addition systems with states. In *CONCUR 2019*, pages 27:1–27:16, 2019. doi:10.4230/LIPIcs.CONCUR.2019.27.
- 10 Krishnendu Chatterjee, Thomas A. Henzinger, and Jan Otop. Quantitative automata under probabilistic semantics. *Logical Methods in Computer Science*, 15(3), 2019. doi:10.23638/LMCS-15(3:16)2019.
- 11 Krishnendu Chatterjee and Yaron Velner. The complexity of mean-payoff pushdown games. *J. ACM*, 64(5):34:1–34:49, 2017. doi:10.1145/3121408.
- 12 Krishnendu Chatterjee and Yaron Velner. Hyperplane separation technique for multidimensional mean-payoff games. *J. Comput. Syst. Sci.*, 88:236–259, 2017. doi:10.1016/j.jcss.2017.04.005.
- 13 Wojciech Czerwinski, Slawomir Lasota, Ranko Lazic, Jérôme Leroux, and Filip Mazowiecki. The reachability problem for petri nets is not elementary. In *STOC 2019*, pages 24–33, 2019. doi:10.1145/3313276.3316369.
- 14 Emanuele D’Osualdo, Jonathan Kochems, and C.-H. Luke Ong. Automatic verification of erlang-style concurrency. In *SAS 2013*, pages 454–476, 2013. doi:10.1007/978-3-642-38856-9_24.
- 15 Javier Esparza. Decidability and complexity of petri net problems—an introduction. *Lectures on Petri nets I: Basic models*, pages 374–428, 1998.
- 16 Javier Esparza and Mogens Nielsen. Decidability issues for petri nets - a survey. *Bull. EATCS*, 52:244–262, 1994.
- 17 W. Feller. *An introduction to probability theory and its applications*. Wiley, 1971.
- 18 Yu Feng, Ruben Martins, Yuepeng Wang, Isil Dillig, and Thomas W. Reps. Component-based synthesis for complex apis. In *POPL 2017*, pages 599–612, New York, NY, USA, 2017. ACM. doi:10.1145/3009837.3009851.
- 19 Jerzy Filar and Koos Vrieze. *Competitive Markov decision processes*. Springer, 1996.
- 20 Pierre Ganty and Rupak Majumdar. Algorithmic verification of asynchronous programs. *ACM Trans. Program. Lang. Syst.*, 34(1):6:1–6:48, May 2012. doi:10.1145/2160910.2160915.
- 21 Christoph Haase and Simon Halfon. Integer vector addition systems with states. In *RP 2014*, pages 112–124, 2014. doi:10.1007/978-3-319-11439-2_9.

- 22 Christoph Haase, Stephan Kreutzer, Joël Ouaknine, and James Worrell. Reachability in succinct and parametric one-counter automata. In *CONCUR 2009*, pages 369–383, 2009. doi:10.1007/978-3-642-04081-8_25.
- 23 Alexander Kaiser, Daniel Kroening, and Thomas Wahl. Dynamic cutoff detection in parameterized concurrent programs. In *CAV 2010*, pages 645–659, 2010. doi:10.1007/978-3-642-14295-6_55.
- 24 Alexander Kaiser, Daniel Kroening, and Thomas Wahl. Efficient coverability analysis by proof minimization. In Maciej Koutny and Irek Ulidowski, editors, *CONCUR 2012*, pages 500–515, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. doi:10.1007/978-3-642-32940-1_35.
- 25 Richard M. Karp and Raymond E. Miller. Parallel program schemata. *J. Comput. Syst. Sci.*, 3(2):147–195, 1969. doi:10.1016/S0022-0000(69)80011-5.
- 26 S. Rao Kosaraju. Decidability of reachability in vector addition systems (preliminary version). In *Proceedings of the 14th Annual ACM Symposium on Theory of Computing, May 5-7, 1982, San Francisco, California, USA*, pages 267–281, 1982. doi:10.1145/800070.802201.
- 27 Jean-Luc Lambert. A structure to decide reachability in petri nets. *Theoretical Computer Science*, 99(1):79–104, 1992. doi:10.1016/0304-3975(92)90173-D.
- 28 Jérôme Leroux. Vector addition systems reachability problem (A simpler solution). In *Turing-100 - The Alan Turing Centenary, Manchester, UK, June 22-25, 2012*, pages 214–228, 2012. URL: <https://easychair.org/publications/paper/Blr>.
- 29 Jérôme Leroux. Polynomial vector addition systems with states. In *ICALP 2018*, pages 134:1–134:13, 2018. doi:10.4230/LIPIcs.ICALP.2018.134.
- 30 Richard Lipton. The reachability problem is exponential-space hard. *Department of Computer Science, Yale University, Tech. Rep.*, 62, 1976.
- 31 Ernst W. Mayr. An Algorithm for the General Petri Net Reachability Problem. In *STOC 1981*, pages 238–246, 1981. doi:10.1145/800076.802477.
- 32 Jakub Michaliszyn and Jan Otop. Average stack cost of büchi pushdown automata. In *FSTTCS 2017*, pages 42:1–42:13, 2017. doi:10.4230/LIPIcs.FSTTCS.2017.42.
- 33 Charles Rackoff. The covering and boundedness problems for vector addition systems. *Theoretical Computer Science*, 6(2):223–231, 1978. doi:10.1016/0304-3975(78)90036-1.
- 34 Moritz Sinn, Florian Zuleger, and Helmut Veith. A simple and scalable static analysis for bound analysis and amortized complexity analysis. In *CAV 2014*, pages 745–761, 2014. doi:10.1007/978-3-319-08867-9_50.

A Appendix

► **Lemma 5.** *Let \mathcal{A} be a VASS(\mathbb{Z}). If it has a witness for $\text{LIMAVG}_{\vec{s}} \leq \vec{\lambda}$, then there exists a computation π such that $\text{LIMAVG}_{\vec{s}}(\pi) \leq \vec{\lambda}$.*

Proof. Let $(q_0, \vec{0})$ be the initial configuration and (s_0, \vec{x}) be the recurrent configuration of the witness. We assume that \mathbf{B} is non-empty. Otherwise, the construction presented in the all-bounded case essentially works. The only difference is that we use paths $\mathbf{in}_i, \mathbf{out}_i$ to switch between cycles.

We define the path \mathbf{p} of the form

$$\mathbf{p} = s_0 s_1^1 s_2^1 \dots s_k^1 s_1^2 \dots s_k^2 \dots$$

such that for every prefix $s_0 s_1^1 \dots s_i^j$ of \mathbf{p} , the precomputation ρ_i^j corresponding to that prefix satisfies:

(a) $\text{AVG}_{S[i]}(\rho_i^j) \leq \vec{\lambda}[i] + \frac{1}{j}$, and

(b) ρ_i^j terminates in (s_0, \vec{y}) , where for every $i \in \mathbf{B}$ we have $\vec{y}[i] = \vec{x}[i]$.

Having such a path \mathbf{p} , consider the computation π that corresponds to \mathbf{p} . Observe that for every counter i , condition (a) implies that for every $\epsilon > 0$ there are infinitely many positions k such that the average $\text{AVG}_{S[i]}(\pi[1, k])$ is less than $\vec{\lambda}[i] + \epsilon$ and hence $\text{LIMAVG}_{S[i]}(\pi) \leq \vec{\lambda}[i]$. It follows that $\text{LIMAVG}_{\vec{S}}(\pi) \leq \vec{\lambda}$.

Now, we discuss how to construct such \mathbf{p} . First, \mathbf{s}_0 is a path that corresponds to a computation from $(q_0, \vec{0})$ to (s_0, \vec{x}) . Second, suppose that a prefix \mathbf{p}_1 of \mathbf{p} has been defined as above, and we need to construct \mathbf{s}_i^j . Observe that the already constructed subcomputation ends in (s_0, \vec{y}) such that for every $k \in \mathbf{B}$ we have $\vec{y}[i] = \vec{x}[i]$.

There exist paths $\mathbf{in}_i, \mathbf{out}_i$ and a cycle \mathbf{c}_i satisfying **(BU)**, and **(U1)** (if $i \in \mathbf{U}$) or **(B1)** (if $i \in \mathbf{B}$). Consider \mathbf{s}_i^j of the form $\mathbf{in}_i \mathbf{c}_i^N \mathbf{out}_i$ from some $N > 0$ fixed later. First, due to condition **(BU)**, for every $k \in \mathbf{B}$ we have $\text{GAIN}(\mathbf{in}_i \mathbf{out}_i) = 0$ and $\text{GAIN}(\mathbf{c}_i) = 0$, and hence $\text{GAIN}(\mathbf{s}_i^j) = 0$. It follows that **(b)** holds.

Now, to see that **(a)** holds for N big enough we consider two cases. If counter i is bounded, then **(B1)** implies that for ρ_N being the computation corresponding to $\mathbf{p}_1 \mathbf{in}_i \mathbf{c}_i^N \mathbf{out}_i$, the average of the part corresponding to \mathbf{c}_i^N is x , which is less or equal to $\vec{\lambda}[i]$. Therefore, $\text{AVG}_{S[i]}(\rho_N)$ tends to x as $N \rightarrow \infty$ and hence there is N such that $\text{AVG}_{S[i]}(\rho_N) \leq \vec{\lambda}[i] + \frac{1}{j}$.

If counter i is unbounded, then $\text{GAIN}(\mathbf{c}_i)[i] < 0$ and \mathbf{c}_i contains a selecting state. It follows that the values of counter i tend to $-\infty$, and hence $\text{AVG}_{S[i]}(\rho_N)$ tends to $-\infty$ as $N \rightarrow \infty$. Therefore, there exists N such that $\text{AVG}_{S[i]}(\rho_N) \leq \vec{\lambda}[i] + \frac{1}{j}$. \blacktriangleleft

► **Lemma 6.** *For all VASS(\mathbb{Z}, k) \mathcal{A} and $\vec{\lambda} \in \mathbb{Q}^k$ the following holds: if there is a computation π such that $\text{LIMAVG}_{\vec{S}}(\pi) \leq \vec{\lambda}$, then there exists a witness for $\text{LIMAVG}_{\vec{S}} \leq \vec{\lambda}$, which has a polynomial size in $|\mathcal{A}| + |\vec{\lambda}|$.*

Proof. Consider a computation π such that $\text{LIMAVG}_{\vec{S}}(\pi) \leq \vec{\lambda}$ and let \mathbf{p} be the infinite path corresponding to π . We decompose \mathbf{p} into simple cycles greedily always picking the first occurring simple cycle. Now, consider all simple cycles that occur infinitely often as well as all rotations of these cycles $D = \{\mathbf{d}_1, \dots, \mathbf{d}_m\}$. Based on these cycles, we define \mathbf{B} as the set of counters j such that for all cycles $\mathbf{d} \in D$ we have $\text{GAIN}(\mathbf{d})[j] = 0$, and $\mathbf{U} = \{1, \dots, k\} \setminus \mathbf{B}$.

Let s_0 be a state that occurs infinitely often in \mathbf{p} . Note that eventually, past some position K , all transitions belong to cycles from D . Therefore, we pick the first configuration (s_0, \vec{x}) past position K and observe that for all successive configurations (s_0, \vec{y}) , for every counter $i \in \mathbf{B}$, the gain between these configurations is 0 and hence $\vec{x}[i] = \vec{y}[i]$. The length of description of (s_0, \vec{x}) is unbounded, but we show at the end of the proof that it can be chosen to be polynomial in $|\mathcal{A}| + |\vec{\lambda}|$. First, we show that there is any witness.

Consider a counter j such that all cycles $\mathbf{d} \in D$ satisfy $\text{GAIN}(\mathbf{d})[j] \geq 0$. We observe that for all cycles $\mathbf{d} \in D$ we have $\text{GAIN}(\mathbf{d})[j] = 0$ and hence $j \in \mathbf{B}$. Indeed, if there is a cycle $\mathbf{d} \in D$ that satisfies $\text{GAIN}(\mathbf{d})[j] > 0$, then values of counter j in π tend to ∞ and hence $\text{LIMAVG}_{S[j]}(\pi) = \infty > \vec{\lambda}[j]$. It follows that for every $i \in \mathbf{U}$, there is a cycle \mathbf{c}_i such that $\text{GAIN}(\mathbf{c}_i)[i] < 0$.

The unbounded-counter case. Consider $i \in \mathbf{U}$. Let $\tilde{\mathbf{c}}_i \in D$ be such that $\text{GAIN}(\tilde{\mathbf{c}}_i)[i] < 0$ and let q be the first state of $\tilde{\mathbf{c}}_i$. Observe that there exist cycles $\mathbf{f}_1, \mathbf{f}_2 \in D$ such that \mathbf{f}_1 is from s_0 to itself and contains q , and \mathbf{f}_2 is from q to itself and contains some selecting state from S_i . Indeed, q and s_0 occur infinite often in \mathbf{p} . Consider disjoint cycles $\mathbf{e}_1, \mathbf{e}_2, \dots$ each from s_0 to itself that contains q . For each \mathbf{e}_l we remove from it iteratively simple cycles from D such that the resulting \mathbf{e}'_l does not contain any simple cycle from D . Observe that only finitely many \mathbf{e}'_l are non-empty as otherwise there would be another simple cycle that occurs infinitely often and does not belong to D . Now, let \mathbf{e}_i be a cycle that can be decomposed

into simple cycles from D . Let us remove iteratively simple cycles to leave the ends s_0 of \mathbf{e}_l and a single occurrence of q . The resulting cycle consists of one or two simple cycles from D . The proof for \mathbf{f}_2 is similar.

Now, for $N = |\text{GAIN}(\mathbf{f}_2)[i]| + 1$ we define $\mathbf{c}_i = \mathbf{f}_2 \tilde{\mathbf{c}}_i^N$. Then, $\text{GAIN}(\mathbf{c}_i^N)[i] < 0$ and \mathbf{c}_i contains a selecting state from $S[i]$. Therefore, condition **(U1)** holds.

The cycle \mathbf{f}_1 can be decomposed into paths $\mathbf{in}_i, \mathbf{out}_i$, respectively from s_0 to q , and from q to s_0 . Furthermore, since $\mathbf{f}_2, \mathbf{c}_i, \mathbf{in}_i, \mathbf{out}_i$ can be decomposed into cycles from D , then by definition of \mathbf{B} , for all $j \in \mathbf{B}$ we have $\text{GAIN}(\mathbf{f}_2)[j] = \text{GAIN}(\mathbf{c}_i)[j] = \text{GAIN}(\mathbf{in}_i, \mathbf{out}_i)[j] = 0$ and hence $\text{GAIN}(\mathbf{f}_2 \mathbf{c}_i^N)[j] = 0$. Therefore, condition **(BU)** holds. Note that $\mathbf{in}_i, \mathbf{out}_i$ have the lengths bounded by $2 \cdot |\mathcal{A}|$ and \mathbf{c}_i can be represented by the pair of cycles: $(\tilde{\mathbf{c}}_i, \mathbf{f}_2)$ of the length at most $2 \cdot |\mathcal{A}|$. Thus, all have polynomial-size representation.

The bounded-counter case. Let $i \in \mathbf{B}$. Since every cycle $\mathbf{d} \in D$ satisfies $\text{GAIN}(\mathbf{d})[i] = 0$, from some position K onwards the gain of each cycle is 0 and hence the value of counter i on any two positions past K with the same state are the same. Therefore, we associate with each state the value of counter i and eliminate values of counters. Furthermore, we associate with each cycle $\mathbf{d} \in D$ its average value over $S[i]$, which is uniquely defined. Finally, if for all cycles $\mathbf{d} \in D$ the average exceeds $\vec{\lambda}[i]$, then $\text{LIMAVG}_{S[i]}(\pi) > \vec{\lambda}[i]$. Therefore, there exists a cycle $\mathbf{c}_i \in D$ with the average value less or equal to $\vec{\lambda}[i]$.

Since s_0 occurs infinitely often, in particular it occurs past position K . Therefore, we show as in the unbounded-counter case that there exist $\mathbf{in}_i, \mathbf{out}_i$ such that \mathbf{in}_i leads from s_0 to some state q of \mathbf{c}_i and \mathbf{out}_i from q to s_0 such that $\mathbf{c}_i, \mathbf{in}_i, \mathbf{out}_i$ satisfy **(BU)**. Finally, observe that $\mathbf{c}_i, \mathbf{in}_i, \mathbf{out}_i$ satisfy **(U1)**. Note that $\mathbf{c}_i, \mathbf{in}_i, \mathbf{out}_i$ can be picked to have the lengths at most $2 \cdot |\mathcal{A}|$.

A witness with polynomial recurrent configuration. We have shown that there exists a witness for $\text{LIMAVG}_{\vec{g}} \leq \vec{\lambda}$ with (s_0, \vec{x}) . We show that there exists \vec{z} such that (a) the witness for $\text{LIMAVG}_{\vec{g}} \leq \vec{\lambda}$ with (s_0, \vec{x}) replaced by (s_0, \vec{z}) remains a witness for $\text{LIMAVG}_{\vec{g}} \leq \vec{\lambda}$, and (b) \vec{z} has the binary representation of polynomial length in $|\mathcal{A}| + |\vec{\lambda}|$.

For (a) we need to show that (i) **(B1)** is satisfied with (s_0, \vec{z}) , and (ii) (s_0, \vec{z}) is reachable from $(q_0, \vec{0})$. Recall that **(B1)** states that for every $i \in \mathbf{B}$, the subcomputation ρ corresponding to the cycle \mathbf{c}_i starting from the configuration reached from (s_0, \vec{z}) over the path \mathbf{in}_i satisfies $\text{AVG}_{S[i]}(\rho) \leq \vec{\lambda}[i]$. Note that the lengths of $\mathbf{in}_i, \mathbf{c}_i$ are bounded by $2 \cdot |\mathcal{A}|$ (because $i \in \mathbf{B}$) and hence there exists α_i with the binary representation of polynomial-length in $|\mathcal{A}|$ such that $\text{AVG}_{S[i]}(\rho) = \alpha_i + \vec{z}[i]$. Therefore, any \vec{z} such that $\vec{z}[i] < -\alpha_i + \vec{\lambda}[i]$ for all $i \in \mathbf{B}$ satisfies (i). For (ii) observe that reachable configurations in $\text{VASS}(\mathbb{Z})$ are semilinear sets [4] represented by polynomial-size equations (where coefficients are given in binary). Therefore, we can find a vector \vec{z} satisfying (i) and (ii) whose binary representation has polynomial length in $|\mathcal{A}| + |\vec{\lambda}|$. ◀

► **Lemma 9.** *Let \mathcal{A} be a strongly-connected probabilistic $\text{VASS}(\mathbb{Z}, 1)$. One of the following conditions holds:*

- (1) $\mathbb{E}(\text{Gain}) > 0$, and $\text{LIMAVGINF}_S(\xi) = \text{LIMAVGSUP}_S(\xi) = \infty$ with probability 1 (over ξ),
- (2) $\mathbb{E}(\text{Gain}) < 0$, and $\text{LIMAVGINF}_S(\xi) = \text{LIMAVGSUP}_S(\xi) = -\infty$ with probability 1,
- (3) \mathcal{A} is totally bounded, and for some $x \in \mathbb{Q}$, with probability 1 over ξ we have $\text{LIMAVGINF}_S(\xi) = \text{LIMAVGSUP}_S(\xi) = x$, and
- (4) $\mathbb{E}(\text{Gain}) = 0$, \mathcal{A} is not totally bounded, and with probability 1 over ξ we have $\text{LIMAVGINF}_S(\xi) = -\infty$ and $\text{LIMAVGSUP}_S(\xi) = \infty$.

Proof. We prove only (1),(2),(3), as (4) is presented with details in the paper.

We show conditions (1) and (2). Assume that $\mathbb{E}(\text{Gain}) \neq 0$. The Ergodic Theorem for Markov chains implies that with probability 1 (over ξ) for every state q , the frequency of configurations with the state q converges to $x[q]$. For every transition (q, q', y) , the frequency of this transition converges to $x_q \cdot P(q, q', y)$. Now, we multiply the frequency of each transition (q, q', y) by its update value y and get the value of the counter in $\xi[n]$ divided by n . On the other hand, this value converges to $\mathbb{E}(\text{Gain})$ as n tends to infinity. It follows that with probability 1 over ξ the counter's value at a position n equals $\mathbb{E}(\text{Gain}) \cdot n \pm o(n)$. Therefore, if $\mathbb{E}(\text{Gain}) > 0$ we have $\text{LIMAVGINF}_S(\xi) = \text{LIMAVGSUP}_S(\xi) = \infty$. Similarly, if $\mathbb{E}(\text{Gain}) < 0$, then $\text{LIMAVGINF}_S(\xi) = \text{LIMAVGSUP}_S(\xi) = -\infty$ with probability 1.

Now, we show condition (3). Assume that \mathcal{A} is totally bounded. In every computation π if there are two configurations with the same state $(s, x_1), (s, x_2)$, then the counter's value is the same $x_1 = x_2$. To see that, consider a subcomputation from (s, x_1) to (s, x_2) and let \mathbf{p} be the path that corresponds to that subcomputation. Then, $0 = \text{GAIN}(\mathbf{p}) = x_2 - x_1$. Furthermore, since \mathcal{A} is strongly connected and $(q_0, 0)$ is the initial configuration for all computations, then in all computations the state determines the value of the counter.

It follows that we can eliminate the counter and consider \mathcal{A} as a Markov chain with the limit-average objective with silent moves [10]. In a Markov chain with silent moves, transitions are weighted with rational numbers and a special value \perp , which is skipped in the computation of partial averages. Similarly to Markov chains, in strongly-connected Markov chains with silent moves, the expected limit-average in the Markov chain is actually the limit-average of almost all paths and it is our value x . Moreover, the expected value can be computed in polynomial time [10].

More precisely, let $h(q)$ be the value of the counter in the state q . We define a Markov chain \mathcal{M} corresponding to \mathcal{A} as follows. We define $\mathcal{M} = \langle \{a\}, Q, \{q_0\}, \delta', P', \mu' \rangle$ such that $\delta'(q, q', a)$ holds if and only if $(q, q', x) \in \delta$ for some $x \in \mathbb{Z}$, $P'(q, q', a) = \sum_{x \in \mathbb{Z}} P(q, q', x)$, and $\mu'(q_0) = 1$. We consider the weighted Markov chain with silent moves $\langle \mathcal{M}, \mathbf{c} \rangle$ such that $\mathbf{c}: \delta' \rightarrow \mathbb{Z} \cup \{\perp\}$ is defined as $\mathbf{c}(q, q', a) = h(q)$, if $q \in S$ is a selecting state, and $\mathbf{c}(q, q', a) = \perp$ (is silent) otherwise. Observe that for every computation π and the corresponding \mathbf{p} (without counter updates), we have $\text{LIMAVG}_S(\pi)$ is precisely the limit average of costs \mathbf{c} of transitions along \mathbf{p} . Since for almost all paths \mathbf{p} in $\langle \mathcal{M}, \mathbf{c} \rangle$, the limit average cost of \mathbf{p} is the expected cost x of $\langle \mathcal{M}, \mathbf{c} \rangle$, almost all computations in \mathcal{A} have the limit-average equal to x . \blacktriangleleft

► Lemma 10. *Consider a probabilistic VASS(\mathbb{Z}) \mathcal{A} as in (4) of Lemma 9. There exist $c \in \mathbb{Q}$ and $\delta > 0$ such that $\text{LIMAVGINF}_S(\xi) < c$ holds with probability greater than δ .*

Proof. Let \mathcal{A}' results from \mathcal{A} by assuming that the all states are initial, i.e., $Q_0 = Q$, and the initial distribution over states μ coincides with the long-run frequencies of states, i.e., $\mu(q) = x_q$.

Suppose that $\text{LIMAVGINF}_S(\xi) = \infty$ with probability 1 w.r.t. \mathcal{A} . Then, it also holds with probability 1 w.r.t. \mathcal{A}' . Then, the average counter value at the n -th position converges to ∞ ($\lim_{n \rightarrow \infty} \text{AVG}_S(\xi[1, n]) = \infty$) with probability 1 in \mathcal{A}' . Therefore, the expected average counter value up to position n , $\mathbb{E}_{\mathcal{A}'}(\text{AVG}_S(\xi[1, n]))$, converges to ∞ . However, the expected gain is 0, which implies that in \mathcal{A}' , at every position n , the expected value of the counter (in \mathcal{A}') is 0. It follows that $\mathbb{E}_{\mathcal{A}'}(\text{AVG}_S(\xi[1, n]))$ is 0. A contradiction. \blacktriangleleft

► Lemma 14. *For all VASS($\mathbb{N}, 1$) \mathcal{A} the following holds: there exists a computation π with $\text{LIMAVG}_S(\pi) \leq \lambda$ if and only if there exist subcomputations ρ_0, ρ_c such that:*

- ρ_0 is from $(q_0, 0)$ to (s, x) , where $x \leq \lambda$, and
- ρ_c is a cycle from (s, x) to itself satisfying the following conditions:

- (a) $\text{AVG}_S(\rho_c) \leq \lambda$,
- (b) the number of configurations with selecting states in ρ_c , i.e., configurations from $S \times \mathbb{N}$, is $O(|S| \cdot |\lambda|^2)$, and
- (c) the value of the counter in each configuration of ρ_c from $S \times \mathbb{N}$ is $O(|S| \cdot |\lambda|^2)$.

Proof. Observe that having ρ_0, ρ_c as above, the computation $\pi = \rho_0(\rho_c)^\infty$ is a valid computation and it satisfies $\text{LIMAVG}_S(\pi) \leq \lambda$.

Conversely, assume that there is a computation π with $\text{LIMAVG}_S(\pi) \leq \lambda$. Consider $\epsilon > 0$ and pick a cycle subcomputation ρ from π of the average value at most $\lambda + \epsilon$ of the minimal length (all shorter subcomputations have higher average). Such a cycle exists as there has to be a configuration (s, y) with $y \leq \lambda + \epsilon$ that occurs infinitely often. Otherwise, $\text{LIMAVG}_S(\pi) \geq \lambda + \epsilon$. Then, we divide π into cycles with ends with configuration (s, y) and there has to be a cycle with the average value at most $\lambda + \epsilon$.

We show that this minimal ρ has few *selecting configurations*, which are configurations with a selecting state. Let L be the number of selecting configurations in ρ with the counter's value at most $\lceil \lambda \rceil$ and let H be the number of selecting configurations with the counter's value at least $\lceil \lambda \rceil + 1$. We lower the average if we replace the value of the configurations of the first type by 0 and the second by $\lceil \lambda \rceil + 1$ and get

$$\frac{(\lceil \lambda \rceil + 1)H}{L + H} \leq \lambda + \epsilon$$

and hence $H \leq \frac{\lambda + \epsilon}{1 - \epsilon} \cdot L$. Assuming that $\epsilon \leq 0.5$, we can bound $H \leq 2 \cdot (\lceil \lambda \rceil + 1)L$.

Now, we give a bound on L . Due to minimality assumption on ρ , it cannot contain subcycles with the same properties. Suppose it has a subcycle τ . Due to minimality assumption, the subcycle has the average value exceeding $\lambda + \epsilon$. But then, ρ' obtained from ρ by removal of τ has a smaller average and a shorter length. A contradiction. It follows that for every state q and for every value $x \leq \lceil \lambda \rceil$ there is at most one configuration (q, x) in ρ . Therefore, $L \leq (\lceil \lambda \rceil + 1)|S|$. and hence

$$L + H \leq (\lceil \lambda \rceil + 1)|S| + 2 \cdot (\lceil \lambda \rceil + 1)^2|S| \leq 2 \cdot |S| \cdot (\lceil \lambda \rceil + 2)^2.$$

As previously observed the minimal value of a configuration from L is 0 and from H is $\lceil \lambda \rceil + 1$. Suppose that there is a single high value B in ρ and all other values take the minimal possible value. Then, we get:

$$\frac{B + (H - 1)\lceil \lambda \rceil}{L + H} \leq \lambda$$

thus

$$B \leq \lceil \lambda \rceil(L + 1) \leq (\lceil \lambda \rceil + 1)^2|Q|$$

and that is the bound on the maximal value of a selecting configuration. \blacktriangleleft

► Theorem 19. *The problem, given a probabilistic VASS(\mathbb{N}, k) \mathcal{A} under the relaxed semantics, $S \subseteq Q$ and $\vec{x} \in \mathbb{Q}^k$, decide whether $\mathbb{E}_{\mathcal{A}}(\text{LIMAVG}_{\vec{S}}) < \vec{x}$ is EXPSPACE-hard.*

Proof. Consider a VASS(\mathbb{N}, k) \mathcal{A} , an initial configuration (s, \vec{x}_1) , and a target configuration (t, \vec{x}_2) . Without loss of generality, we assume that $\vec{x}_1 = \vec{x}_2 = \vec{0}$ and the update of each counter is $-1, 0, 1$. We construct a probabilistic VASS(\mathbb{N}) \mathcal{A}^P based on \mathcal{A} by adding an additional counter $k + 1$ and a sink state r , which has only a single outgoing transition, which is a self-loop upon which counters do not change values, i.e., $(r, r, \vec{0})$. We add a transition

23:22 Multi-Dimensional Long-Run Average Problems for VASS

from t to r labeled with $\vec{0}$ and assign to it some positive probability. To make sure that the expected value of \mathcal{A}^P is defined and finite, we add to every state \mathcal{A}^P a transition to r labeled with $\vec{1}$ with probability $\frac{1}{2}$. We assign positive probabilities to the remaining transitions. Observe that a random computation reaches the sink r with probability 1, and the probability that it happens after more than n steps is bounded by $(\frac{1}{2})^n$. The value of the counter after n steps is at most n . Therefore, the expected limit-average of counters $1, \dots, k$ is bounded by $\sum_{j=1}^{\infty} (\frac{1}{2})^j \cdot j = 2$. Upon reaching r the counter $k + 1$ has value 0, if the previous state was t , and 1 otherwise. Therefore, the expected limit-average is strictly less than $(2 + \epsilon, \dots, 2 + \epsilon, 1)$ (for any $\epsilon > 0$) if and only if there is a computation from $(s, \vec{0})$ to (t, \vec{y}) (with any \vec{y}) in \mathcal{A} . ◀