Asymptotic Complexity Estimates for Probabilistic Programs and Their VASS Abstractions

Michal Ajdarów
Masaryk University, Brno, Czech Republic

Antonín Kučera
Masaryk University, Brno, Czech Republic

Abstract

The standard approach to analyzing the asymptotic complexity of probabilistic programs is based on studying the asymptotic growth of certain expected values (such as the expected termination time) for increasing input size. We argue that this approach is not sufficiently robust, especially in situations when the expectations are infinite. We propose new estimates for the asymptotic analysis of probabilistic programs with non-deterministic choice that overcome this deficiency. Furthermore, we show how to efficiently compute/analyze these estimates for selected classes of programs represented as Markov decision processes over vector addition systems with states.

2012 ACM Subject Classification Theory of computation → Models of computation

Keywords and phrases Probabilistic programs, asymptotic complexity, vector addition systems

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2023.12

Related Version *Full Version*: `https://arxiv.org/abs/2307.04707` [2]

Funding The work is supported by the Czech Science Foundation, Grant No. 21-24711S.

1 Introduction

Vector Addition Systems with States (VASS) [11] are a model for discrete systems with multiple unbounded resources expressively equivalent to Petri nets [20]. Intuitively, a VASS with $d \geq 1$ counters is a finite directed graph where the transitions are labeled by d-dimensional vectors of integers representing *counter updates*. A computation starts in some state for some initial vector of non-negative counter values and proceeds by selecting transitions non-deterministically and performing the associated counter updates. Since the counters cannot assume negative values, transitions that would decrease some counter below zero are disabled.

In program analysis, VASS are used as abstractions for programs operating over unbounded integer variables. Input parameters are represented by initial counter values, and more complicated arithmetical functions, such as multiplication, are modeled by VASS gadgets computing these functions in a weak sense (see, e.g., [17]). Branching constructs, such as **if-then-else**, are usually replaced with non-deterministic choice. VASS are particularly useful for evaluating the *asymptotic complexity* of infinite-state programs, i.e., the dependency of the running time (and other complexity measures) on the size of the program input [21, 22]. Traditional VASS decision problems such as reachability, liveness, or boundedness are computationally hard [9, 18, 19], and other verification problems such as equivalence-checking [12] or model-checking [10] are even undecidable. In contrast to this, decision problems related to the asymptotic growth of VASS complexity measures are solvable with low complexity and sometimes even in *polynomial time* [4, 23, 15, 16, 1]; see [14] for a recent overview.

34th International Conference on Concurrency Theory (CONCUR 2023).
Editors: Guillermo A. Pérez and Jean-François Raskin; Article No. 12; pp. 12:1–12:16
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The existing results about VASS asymptotic analysis are applicable to programs with non-determinism (in *demonic* or *angelic* form, see [5]), but cannot be used to analyze the complexity of *probabilistic programs.* This motivates the study of Markov decision process over VASS (VASS MDPs) with both non-deterministic and probabilistic states, where transitions in probabilistic states are selected according to fixed probability distributions. Here, the problems of asymptotic complexity analysis become even more challenging because VASS MDPs subsume infinite-state stochastic models that are notoriously hard to analyze. So far, the only existing result about asymptotic VASS MDP analysis is [3] where the linearity of expected termination time is shown decidable in polynomial time for VASS MDPs with DAG-like MEC decomposition.

Our Contribution: We study the problems of asymptotic complexity analysis for probabilistic programs and their VASS abstractions.

For non-deterministic programs, termination complexity is a function $\mathcal{L}_{\max}$ assigning to every $n \in \mathbb{N}$ the length of the longest computation initiated in a configuration with each counter set to n. A natural way of generalizing this concept to probabilistic programs is to define a function $\mathcal{L}_{\text{exp}}$ such that $\mathcal{L}_{\text{exp}}(n)$ is the maximal *expected length* of a computation initiated in a configuration of size n, where the maximum is taken over all strategies resolving non-determinism. The same approach is applicable to other complexity measures. We show that this natural idea is generally *inappropriate*, especially in situations when $\mathcal{L}_{\text{exp}}(n)$ is *infinite* for a sufficiently large n. By "inappropriate" we mean that this form of asymptotic analysis can be misleading. For example, if $\mathcal{L}_{\text{exp}}(n) = \infty$ for all $n \geq 1$, one may conclude that the computation takes a very long time independently of n. However, this is not necessarily the case, as demonstrated in a simple example of Fig. 1 (we refer to Section 3 for a detailed discussion). Therefore, we propose new notions of *lower/upper/tight complexity estimates* and demonstrate their advantages over the expected values. These notions can be adapted to other models of probabilistic programs, and constitute the main conceptual contribution of our work.

Then, we concentrate on algorithmic properties of the complexity estimates in the setting of VASS MDPs. Our first result concerns *counter complexity.* We show that for every VASS MDP with DAG-like MEC decomposition and every counter c, there are only two possibilities:

- The function n is a *tight estimate* of the asymptotic growth of the maximal c-counter value assumed along a computation initiated in a configuration of size n.
- The function n^2 is a *lower estimate* of the asymptotic growth of the maximal c-counter value assumed along a computation initiated in a configuration of size n.

Furthermore, it is decidable in *polynomial time* which of these alternatives holds.

Since the termination and transition complexities can be easily encoded as the counter complexity for a fresh "step counter", the above result immediately extends also to these complexities. To some extent, this result can be seen as a generalization of the result about termination complexity presented in [3]. See Section 4 for more details.

Our next result is a full classification of asymptotic complexity for one-dimensional VASS MDPs. We show that for every one-dimensional VASS MDP

- the counter complexity is either unbounded or n is a tight estimate;
- termination complexity is either unbounded or one of the functions n, n^2 is a tight estimate.
- transition complexity is either unbounded, or bounded by a constant, or one of the functions n, n^2 is a tight estimate.

Furthermore, it is decidable in *polynomial time* which of the above cases hold.

Since the complexity of the considered problems remains low, the results are encouraging. On the other hand, they require non-trivial insights, indicating that establishing a full and effective classification of the asymptotic complexity of multi-dimensional VASS MDPs is a challenging problem.

Missing proofs can be found in a full version of this paper [2].

2 Preliminaries

We use $\mathbb{N}$, $\mathbb{Z}$, $\mathbb{Q}$, and $\mathbb{R}$ to denote the sets of non-negative integers, integers, rational numbers, and real numbers. Given a function $f \colon \mathbb{N} \to \mathbb{N}$, we use $O(f)$ and $\Omega(f)$ to denote the sets of all $g \colon \mathbb{N} \to \mathbb{N}$ such that $g(n) \leq a \cdot f(n)$ and $g(n) \geq b \cdot f(n)$ for all sufficiently large $n \in \mathbb{N}$, where a, b are some positive constants. If $h \in O(f)$ and $h \in \Omega(f)$, we write $h \in \Theta(f)$.

Let A be a finite index set. The vectors of $\mathbb{R}^A$ are denoted by bold letters such as $\mathbf{u}, \mathbf{v}, \mathbf{z}, \ldots$. The component of $\mathbf{v}$ of index $i \in A$ is denoted by $\mathbf{v}(i)$. If the index set is of the form $A = \{1, 2, \ldots, d\}$ for some positive integer d, we write $\mathbb{R}^d$ instead of $\mathbb{R}^A$. For every $n \in \mathbb{N}$, we use $\mathbf{n}$ to denote the constant vector where all components are equal to n. The other standard operations and relations on $\mathbb{R}$ such as $+$, $\leq$, or $<$ are extended to $\mathbb{R}^d$ in the component-wise way. In particular, $\mathbf{v} < \mathbf{u}$ if $\mathbf{v}(i) < \mathbf{u}(i)$ for every index i.

A *probability distribution* over a finite set A is a vector $\nu \in [0,1]^A$ such that $\sum_{a \in A} \nu(a) = 1$. We say that ν is *rational* if every $\nu(a)$ is rational, and *Dirac* if $\nu(a) = 1$ for some $a \in A$.

2.1 VASS Markov Decision Processes

▶ **Definition 1.** *Let $d \geq 1$. A d-dimensional* VASS MDP *is a tuple $\mathcal{A} = (Q, (Q_n, Q_p), T, P)$, where*

- *$Q \neq \emptyset$ is a finite set of* states *split into two disjoint subsets Q_n and Q_p of* nondeterministic *and* probabilistic *states,*
- *$T \subseteq Q \times \mathbb{Z}^d \times Q$ is a finite set of* transitions *such that, for every $p \in Q$, the set $Out(p) \subseteq T$ of all transitions of the form $(p, \mathbf{u}, q)$ is non-empty.*
- *P is a function assigning to each $t \in Out(p)$ where $p \in Q_p$ a positive rational probability so that $\sum_{t \in T(p)} P(t) = 1$.*

The encoding size of $\mathcal{A}$ is denoted by $\|\mathcal{A}\|$, where the integers representing counter updates are written in binary and probability values are written as fractions of binary numbers. For every $p \in Q$, we use $In(p) \subseteq T$ to denote the set of all transitions of the form $(q, \mathbf{u}, p)$. The update vector $\mathbf{u}$ of a transition $t = (p, \mathbf{u}, q)$ is also denoted by $\mathbf{u}_t$.

A *finite path* in $\mathcal{A}$ of length $n \geq 0$ is a finite sequence of the form $p_0, \mathbf{u}_1, p_1, \mathbf{u}_2, \ldots, \mathbf{u}_n, p_n$ where $(p_i, \mathbf{u}_{i+1}, p_{i+1}) \in T$ for all $i < n$. We use $len(\alpha)$ to denote the length of α. If there is a finite path from p to q, we say that q is *reachable* from p. An *infinite path* in $\mathcal{A}$ is an infinite sequence $\pi = p_0, \mathbf{u}_1, p_1, \mathbf{u}_2, \ldots$ such that every finite prefix of π ending in a state is a finite path in $\mathcal{A}$.

A *strategy* is a function σ assigning to every finite path $p_0, \mathbf{u}_1, \ldots, p_n$ such that $p_n \in Q_n$ a probability distribution over $Out(p_n)$. A strategy is *Markovian (M)* if it depends only on the last state p_n, and *deterministic (D)* if it always returns a Dirac distribution. The set of all strategies is denoted by $\Sigma_{\mathcal{A}}$, or just Σ when $\mathcal{A}$ is understood. Every initial state $p \in Q$ and every strategy σ determine the probability space over infinite paths initiated in p in the standard way. We use $\mathbb{P}_p^\sigma$ to denote the associated probability measure.

A *configuration* of $\mathcal{A}$ is a pair $p\mathbf{v}$, where $p \in Q$ and $\mathbf{v} \in \mathbb{Z}^d$. If some component of $\mathbf{v}$ is negative, then $p\mathbf{v}$ is *terminal*. The set of all configurations of $\mathcal{A}$ is denoted by $C(\mathcal{A})$.

input N
repeat
    random choice:
        $0.5: \quad N := N + 1;$
        $0.5: \quad N := N - 1;$
until $N = 0$

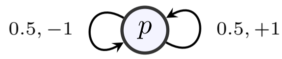

$\mathcal{A}$

Figure 1 A probabilistic program with infinite expected running time for every $N \geq 1$, and its 1-dimensional VASS MDP model $\mathcal{A}$.

Every infinite path $p_0, \mathbf{u}_1, p_1, \mathbf{u}_2, \ldots$ and every initial vector $\mathbf{v} \in \mathbb{Z}^d$ determine the corresponding *computation* of $\mathcal{V}$, i.e., the sequence of configurations $p_0\mathbf{v}_0, p_1\mathbf{v}_1, p_2\mathbf{v}_2, \ldots$ such that $\mathbf{v}_0 = \mathbf{v}$ and $\mathbf{v}_{i+1} = \mathbf{v}_i + \mathbf{u}_{i+1}$. Let $\mathit{Term}(\pi)$ be the least j such that $p_j\mathbf{v}_j$ is terminal. If there is no such j, we put $\mathit{Term}(\pi) = \infty$.

Note that every computation uniquely determines its underlying infinite path. We define the probability space over all computations initiated in a given $p\mathbf{v}$, where the underlying probability measure $\mathbb{P}^\sigma_{p\mathbf{v}}$ is obtained from $\mathbb{P}^\sigma_p$ in an obvious way. For a measurable function X over computations, we use $\mathbb{E}^\sigma_{p\mathbf{v}}[X]$ to denote the expected value of X.

3 Asymptotic Complexity Measures for VASS MDPs

In this section, we introduce asymptotic complexity estimates applicable to probabilistic programs with non-determinism and their abstract models (such as VASS MDPs). We also explain their relationship to the standard measures based on the expected values of relevant random variables.

Let us start with a simple motivating example. Consider the simple probabilistic program of Fig. 1. The program inputs a positive integer N and then repeatedly increments/decrements N with probability 0.5 until $N = 0$. One can easily show that for every $N \geq 1$, the program terminates with probability one, and the expected termination time is *infinite*. Based on this, one may conclude that the execution takes a very long time, independently of the initial value of N. However, this conclusion is *not* consistent with practical experience gained from trial runs[1]. The program tends to terminate “relatively quickly” for small N, and the termination time *does* depend on N. Hence, the function assigning ∞ to every $N \geq 1$ is *not* a faithful characterization of the asymptotic growth of termination time. We propose an alternative characterization based on the following observations[2]:

- For every $\varepsilon > 0$, the probability of all runs terminating after more than $n^{2+\varepsilon}$ steps (where n is the initial value of N) approaches *zero* as $n \to \infty$.
- For every $\varepsilon > 0$, the probability of all runs terminating after more than $n^{2-\varepsilon}$ steps (where n is the initial value of N) approaches *one* as $n \to \infty$.

Since the execution time is “squeezed” between $n^{2-\varepsilon}$ and $n^{2+\varepsilon}$ for an arbitrarily small $\varepsilon > 0$ as $n \to \infty$, it can be characterized as “asymptotically quadratic”. This analysis is in accordance with experimental outcomes.

[1] For $N = 1$, about 95% of trial runs terminate after at most 1000 iterations of the **repeat-until** loop. For $N = 10$, only about 75% of all runs terminate after at most 1000 iterations, but about 90% of them terminate after at most 10000 iterations.

[2] Formal proofs of these observations are simple; in Section 5, we give a full classification of the asymptotic behaviour of one-dimensional VASS MDPs subsuming the trivial example of Fig. 1.

3.1 Complexity of VASS Runs

We recall the complexity measures for VASS runs used in previous works [4, 23, 15, 16, 1]. These functions can be seen as variants of the standard time/space complexities for Turing machines.

Let $\mathcal{A} = (Q, (Q_n, Q_p), T, P)$ be a d-dimensional VASS MDP, $c \in \{1, \ldots, d\}$, and $t \in T$. For every computation $\pi = p_0\mathbf{v}_0, p_1\mathbf{v}_1, p_2\mathbf{v}_2, \ldots$, we put

$$\begin{aligned}\mathcal{L}(\pi) &= \mathit{Term}(\pi)\\ \mathcal{C}[c](\pi) &= \sup\{\mathbf{v}_i(c) \mid 0 \leq i < \mathit{Term}(\pi)\}\\ \mathcal{T}[t](\pi) &= \text{the total number of all } 0 \leq i < \mathit{Term}(\pi) \text{ such that } (p_i, \mathbf{v}_{i+1}-\mathbf{v}_i, p_{i+1}) = t\end{aligned}$$

We refer to the functions $\mathcal{L}$, $\mathcal{C}[c]$, and $\mathcal{T}[t]$ as *termination*, *c-counter*, and *t-transition complexity*, respectively.

Let $\mathcal{F}$ be one of the complexity functions defined above. In VASS abstractions of computer programs, the input is represented by initial counter values, and the input size corresponds to the maximal initial counter value. The existing works on *non-probabilistic* VASS concentrate on analyzing the asymptotic growth of the functions $\mathcal{F}_{\max} : \mathbb{N} \to \mathbb{N}_\infty$ where

$$\mathcal{F}_{\max}(n) = \max\{\mathcal{F}(\pi) \mid \pi \text{ is a computation initiated in } p\mathbf{n} \text{ where } p \in Q\}$$

For VASS MDP, we can generalize $\mathcal{F}_{\max}$ into $\mathcal{F}_{\exp}$ as follows:

$$\mathcal{F}_{\exp}(n) = \max\{\mathbb{E}^{\sigma}_{p\mathbf{n}}[\mathcal{F}] \mid \sigma \in \Sigma_{\mathcal{A}}, p \in Q\}$$

Note that for non-probabilistic VASS, the values of $\mathcal{F}_{\max}(n)$ and $\mathcal{F}_{\exp}(n)$ are the same. However, the function $\mathcal{F}_{\exp}$ suffers from the deficiency illustrated in the motivating example at the beginning of Section 3. To see this, consider the one-dimensional VASS MDP $\mathcal{A}$ modeling the simple probabilistic program (see Fig. 1). For every $n \geq 1$ and the only (trivial) strategy σ, we have that $\mathbb{P}^{\sigma}_{p\mathbf{n}}[\mathit{Term} < \infty] = 1$ and $\mathcal{L}_{\exp}(n) = \infty$. However, the practical experience with trial runs of $\mathcal{A}$ is the same as with the original probabilistic program (see above).

3.2 Asymptotic Complexity Estimates

In this section, we introduce asymptotic complexity estimates allowing for a precise analysis of the asymptotic growth of the termination, c-counter, and t-transition complexity, especially when their expected values are infinite for a sufficiently large input. For the sake of readability, we first present a simplified variant applicable to *strongly connected* VASS MDPs.

Let $\mathcal{F}$ be one of the complexity functions for VASS computations defined in Section 3.1, and let $f : \mathbb{N} \to \mathbb{N}$. We say that f is a *tight estimate of* $\mathcal{F}$ if, for arbitrarily small $\varepsilon > 0$, the value of $\mathcal{F}(n)$ is "squeezed" between $f^{1-\varepsilon}(n)$ and $f^{1+\varepsilon}(n)$ as $n \to \infty$. More precisely, for every $\varepsilon > 0$,

- there exist $p \in Q$ and strategies $\sigma_1, \sigma_2, \ldots$ such that $\liminf_{n\to\infty} \mathbb{P}^{\sigma_n}_{p\mathbf{n}}[\mathcal{F} \geq (f(n))^{1-\varepsilon}] = 1$;
- for all $p \in Q$ and strategies $\sigma_1, \sigma_2, \ldots$ we have that $\limsup_{n\to\infty} \mathbb{P}^{\sigma_n}_{p\mathbf{n}}[\mathcal{F} \geq (f(n))^{1+\varepsilon}] = 0$.

The above definition is adequate for strongly connected VASS MDPs because tight estimates tend to exist in this subclass. Despite some effort, we have not managed to construct an example of a strongly connected VASS MDP where an $\mathcal{F}$ with some upper polynomial estimate does *not* have a tight estimate (see Conjecture 3). However, if the underlying graph of $\mathcal{A}$ is *not* strongly connected, then the asymptotic growth of $\mathcal{F}$ can differ for computations visiting a different sequence of maximal end components (MECs) of $\mathcal{A}$, and the asymptotic growth of $\mathcal{F}$ can be "squeezed" between $f^{1-\varepsilon}(n)$ and $f^{1+\varepsilon}(n)$ only for the subset of computations visiting the same sequence of MECs. This explains why we need a more general definition of complexity estimates presented below.

An *end component (EC)* of $\mathcal{A}$ is a pair (C, L) where $C \subseteq Q$ and $L \subseteq T$ such that the following conditions are satisfied:

- $C \neq \emptyset$;
- if $p \in C \cap Q_n$, then at least one outgoing transition of p belongs to L;
- if $p \in C \cap Q_p$, then all outgoing transitions of p belong to L;
- if $(p, \mathbf{u}, q) \in L$, then $p, q \in C$;
- for all $p, q \in C$ we have that q is reachable from p and vice versa.

Note that if (C, L) and (C', L') are ECs such that $C \cap C' \neq \emptyset$, then $(C \cup C', L \cup L')$ is also an EC. Hence, every $p \in Q$ either belongs to a unique *maximal end component* (MEC), or does not belong to any EC. Also observe that each MEC can be seen as a strongly connected VASS MDP. We say that $\mathcal{A}$ has *DAG-like MEC decomposition* if for every pair M, M' of different MECs such that the states of M' are reachable from the states of M we have that the states of M are not reachable from the states of M'.

For every infinite path π of $\mathcal{A}$, let $mecs(\pi)$ be the unique sequence of MECs visited by π. Observe that $mecs(\pi)$ disregards the states that do not belong to any EC; intuitively, this is because the transitions executed in such states do not influence the asymptotic growth of $\mathcal{F}$. Observe that the length of $mecs(\pi)$, denoted by $len(mecs(\pi))$, can be finite or infinite. The first possibility corresponds to the situation when an infinite suffix of π stays within the same MEC. Furthermore, for all $\sigma \in \Sigma$ and $p \in Q$, we have that $\mathbb{P}_p^\sigma[len(mecs) = \infty] = 0$, and the probability $\mathbb{P}_p^\sigma[len(mecs) \geq k]$ decays exponentially in k (these folklore results are easy to prove). All of these notions are lifted to computations in an obvious way.

Observe that if a strategy σ aims at maximizing the growth of $\mathcal{F}$, we can safely assume that σ eventually stays in a *bottom* MEC that cannot be exited (intuitively, σ can always move from a non-bottom MEC to a bottom MEC by executing a few extra transitions that do not influence the asymptotic growth of $\mathcal{F}$, and the bottom MEC may allow increasing $\mathcal{F}$ even further). On the other hand, the maximal asymptotic growth of $\mathcal{F}$ may be achievable along some "minimal" sequence of MECs, and this information is certainly relevant for understanding the behaviour of a given probabilistic program. This leads to the following definition:

▶ **Definition 2.** *A* type *is a finite sequence β of MECs such that $mecs(\pi) = \beta$ for some infinite path π.*

We say that f is a lower estimate of $\mathcal{F}$ for a type β *if for every $\varepsilon > 0$ there exist $p \in Q$ and a sequence of strategies $\sigma_1, \sigma_2, \ldots$ such that $\mathbb{P}_{p\mathbf{n}}^{\sigma_n}[mecs = \beta] > 0$ for all $n \geq 1$ and*

$$\liminf_{n\to\infty} \; \mathbb{P}_{p\mathbf{n}}^{\sigma_n}[\mathcal{F} \geq (f(n))^{1-\varepsilon} \mid mecs{=}\beta] \;=\; 1\,.$$

Similarly, we say that f is an upper estimate of $\mathcal{F}$ for a type β *if for every $\varepsilon > 0$, every $p \in Q$, and every sequence of strategies $\sigma_1, \sigma_2, \ldots$ such that $\mathbb{P}_{p\mathbf{n}}^{\sigma_n}[mecs = \beta] > 0$ for all $n \geq 1$ we have that*

$$\limsup_{n\to\infty} \; \mathbb{P}_{p\mathbf{n}}^{\sigma_n}[\mathcal{F} \geq (f(n))^{1+\varepsilon} \mid mecs{=}\beta] \;=\; 0$$

If there is no upper estimate of $\mathcal{F}$ for a type β, we say that $\mathcal{F}$ is unbounded for β*. Finally, we say that f is a* tight estimate of $\mathcal{F}$ for β *if it is both a lower estimate and an upper estimate of $\mathcal{F}$ for β.*

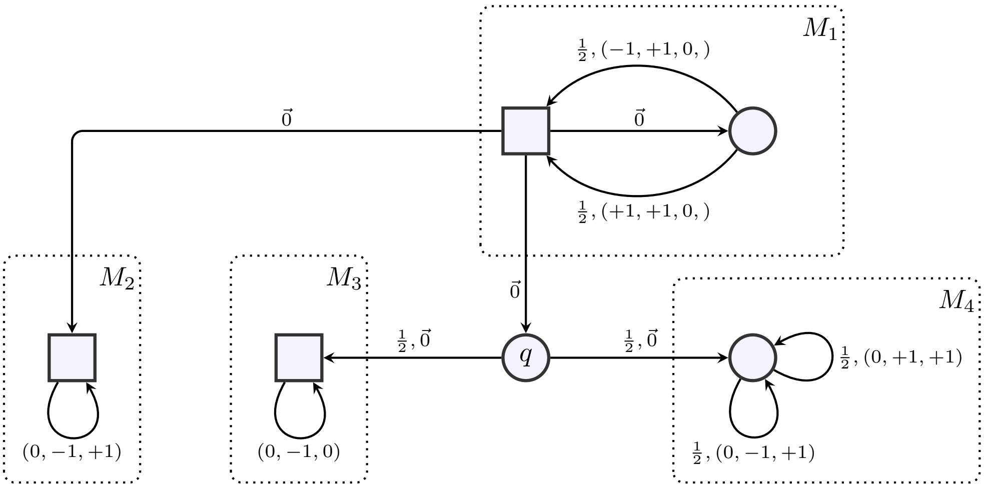

Figure 2 A VASS MDP $\mathcal{A}$ with four MECs and seven types.

Let us note that in the subclass of *non-probabilistic* VASS, MECs become strongly connected components (SCCs), and types correspond to paths in the directed acyclic graph of SCCs. Each such path determines the corresponding asymptotic increase of $\mathcal{F}$, as demonstrated in [1]. We conjecture that types play a similar role for VASS MDPs. More precisely, we conjecture the following:

▶ **Conjecture 3.** *If some polynomial is an upper estimate of $\mathcal{F}$ for β, then there exists a tight estimate f of $\mathcal{F}$ for β.*

Even if Conjecture 3 is proven wrong, there are interesting subclasses of VASS MDPs where it holds, as demonstrated in subsequent sections.

For every pair of MECs M, M', let $P(M, M')$ be the maximal probability (achievable by some strategy) of reaching a state of M' from a state of M in $\mathcal{A}$ without passing through a state of some other MEC M''. Note that $P(M, M')$ is efficiently computable by standard methods for finite-state MDPs. The *weight* of a given type $\beta = M_1, \ldots, M_k$ is defined as $\mathit{weight}(\beta) = \prod_{i=1}^{k-1} P(M_i, M_{i+1})$. Intuitively, $\mathit{weight}(\beta)$ corresponds to the maximal probability of "enforcing" the asymptotic growth of $\mathcal{F}$ according to the tight estimate f of $\mathcal{F}$ for β achievable by some strategy.

Generally, higher asymptotic growth of $\mathcal{F}$ may be achievable for types with smaller weights. Consider the following example to understand better the types, their weights, and the associated tight estimates.

▶ **Example 4.** Let $\mathcal{A}$ be the VASS MDP of Fig. 2. There are four MECs M_1, M_2, M_3, M_4 where M_2, M_3, M_4 are bottom MECs. Hence, there are four types of length one and three types of length two. Let us examine the types of length two initiated in M_1 for $\mathcal{F} \equiv \mathcal{C}[c]$ where c is the third counter.

Note that in M_1, the first counter is repeatedly incremented/decremented with the same probability $\frac{1}{2}$. The second counter "counts" these transitions and thus it is "pumped" to a *quadratic* value (cf. the VASS MDP of Fig. 1). Then, a strategy may decide to move to M_2, where the value of the second counter is transferred to the third counter. Hence, n^2 is the tight estimate of $\mathcal{C}[c]$ for the type M_1, M_2, and $\mathit{weight}(M_1, M_2) = 1$. Alternatively, a strategy may decide to move to the probabilistic state q. Then, either M_3 or M_4 is entered with the same probability $\frac{1}{2}$, which implies $\mathit{weight}(M_1, M_3) = \mathit{weight}(M_1, M_4) = \frac{1}{2}$. In M_3, the third counter is unchanged, and hence n is the tight estimate of $\mathcal{C}[c]$ for the type M_1, M_3. However, in M_4, the second counter previously pumped to a quadratic value is repeatedly incremented/decremented with the same probability $\frac{1}{2}$, and the third counter "counts" these transitions. This means that n^4 is a tight estimate of $\mathcal{C}[c]$ for the type M_1, M_4.

This analysis provides detailed information about the asymptotic growth of $\mathcal{C}[c]$ in $\mathcal{A}$. Every type shows "how" the growth specified by the corresponding tight estimate is achievable, and its weight corresponds to the "maximal achievable probability of this growth". This information is completely lost when analyzing the maximal expected value of $\mathcal{C}[c]$ for computations initiated in configurations $p\mathbf{n}$ where p is a state of M_1, because these expectations are *infinite* for all $n \geq 1$.

Finally, let us clarify the relationship between the lower/upper estimates of $\mathcal{F}$ and the asymptotic growth of $\mathcal{F}_{\text{exp}}$. The following observation is easy to prove.

▶ **Observation 5.** *If $\mathcal{F}_{\text{exp}} \in O(f)$ where $f : \mathbb{N} \to \mathbb{N}$ is an unbounded function, then f is an upper estimate of $\mathcal{F}$ for every type. Furthermore, if $f : \mathbb{N} \to \mathbb{N}$ is a lower estimate of $\mathcal{F}$ for some type, then $\mathcal{F}_{\text{exp}} \in \Omega(f^{1-\epsilon})$ for each $\epsilon > 0$. However, if $\mathcal{F}_{\text{exp}} \in \Omega(f)$ where $f : \mathbb{N} \to \mathbb{N}$, then f is* not *necessarily a lower estimate of $\mathcal{F}$ for some type.*

Observation 5 shows that complexity estimates are generally more informative than the asymptotics of $\mathcal{F}_{\text{exp}}$ even if $\mathcal{F}_{\text{exp}} \in \Theta(f)$ for some "reasonable" function f. For example, it may happen that there are only two types β_1 and β_2 where n and n^3 are tight estimates of $\mathcal{L}$ for β_1 and β_2 with weights 0.99 and 0.01, respectively. In this case, $\mathcal{L}_{\text{exp}} \in \Theta(n^3)$, although the termination time is linear for 99% of computations.

4 A Dichotomy between Linear and Quadratic Estimates

In this section, we prove the following result:

▶ **Theorem 6.** *Let $\mathcal{A}$ be a VASS MDP with DAG-like MEC decomposition and $\mathcal{F}$ one of the complexity functions $\mathcal{L}$, $\mathcal{C}[c]$, or $\mathcal{T}[t]$. For every type β, we have that either n is a tight estimate of $\mathcal{F}$ for β, or n^2 is a lower estimate of $\mathcal{F}$ for β. It is decidable in polynomial time which of the two cases holds.*

Theorem 6 can be seen as a generalization of the linear/quadratic dichotomy results previously achieved for non-deterministic VASS [4] and for the termination complexity in VASS MDPs [3].

It suffices to prove Theorem 6 for the *counter complexity*. The corresponding results for the termination and transition complexities then follow as simple consequences. To see this, observe that we can extend a given VASS MDP with a fresh "step counter" sc that is incremented by every transition (in the case of $\mathcal{L}$) or the transition t (in the case of $\mathcal{T}[t]$) and thus "emulate" $\mathcal{L}$ and $\mathcal{T}[t]$ as $\mathcal{C}[sc]$.

We first consider the case when $\mathcal{A}$ is strongly connected and then generalize the obtained results to VASS MDPs with DAG-like MEC decomposition. So, let $\mathcal{A}$ be a strongly connected d-dimensional VASS MDP and c a counter of $\mathcal{A}$. The starting point of our analysis is the dual constraint system designed in [23] for non-probabilistic strongly connected VASS. We generalize this system to strongly connected VASS MDPs in the way shown in Figure 3 (the original system of [23] can be recovered by disregarding the probabilistic states).

Note that solutions of both (I) and (II) are closed under addition. Therefore, both (I) and (II) have solutions maximizing the specified objectives, computable in polynomial time. For clarity, let us first discuss an intuitive interpretation of these solutions, starting with simplified variants obtained for non-probabilistic VASS in [23].

Constraint system (I):

Find $\mathbf{x} \in \mathbb{Z}^T$ such that

$$\sum_{t \in T} \mathbf{x}(t)\mathbf{u}_t \geq \vec{0}$$

$$\mathbf{x} \geq \vec{0}$$

and for each $p \in Q$

$$\sum_{t \in Out(p)} \mathbf{x}(t) = \sum_{t \in In(p)} \mathbf{x}(t)$$

and for all $p \in Q_p$, $t \in Out(p)$

$$\mathbf{x}(t) = P(t) \cdot \sum_{t' \in Out(p)} \mathbf{x}(t')$$

Objective: *Maximize*

- the number of valid inequalities of the form
$$\sum_{t \in T} \mathbf{x}(t)\mathbf{u}_t(c) > 0,$$
- the number of valid inequalities of the form $\mathbf{x}(t) > 0$.

Constraint system (II):

Find $\mathbf{y} \in \mathbb{Z}^d, \mathbf{z} \in \mathbb{Z}^Q$ such that

$$\mathbf{y} \geq \vec{0}$$

$$\mathbf{z} \geq \vec{0}$$

and for each $(p, \mathbf{u}, q) \in T$ where $p \in Q_n$

$$\mathbf{z}(q) - \mathbf{z}(p) + \sum_{i=1}^{d} \mathbf{u}(i)\mathbf{y}(i) \leq 0$$

and for each $p \in Q_p$

$$\sum_{t=(p,\mathbf{u},q) \in Out(p)} P(t)\big(\mathbf{z}(q) - \mathbf{z}(p) + \sum_{i=1}^{d} \mathbf{u}_t(i)\mathbf{y}(i)\big) \leq 0$$

Objective: *Maximize*

- the number of valid inequalities of the form $\mathbf{y}(c) > 0$,
- the number of transitions $t = (p, \mathbf{u}, q)$ such that $p \in Q_n$ and
$$\mathbf{z}(q) - \mathbf{z}(p) + \sum_{i=1}^{d} \mathbf{u}(i)\mathbf{y}(i) < 0,$$
- the number of states $p \in Q_p$ such that
$$\sum_{t=(p,\mathbf{u},q) \in Out(p)} P(t)\big(\mathbf{z}(q) - \mathbf{z}(p) + \sum_{i=1}^{d} \mathbf{u}(i)\mathbf{y}(i)\big) < 0\,.$$

Figure 3 Constraint systems for strongly connected VASS MDPs.

In the non-probabilistic case, a solution of (I) can be interpreted as a *weighted multicycle*, i.e., as a collection of cycles $M_1, \ldots, M_k$ together with weights $a_1, \ldots, a_k$ such that the total effect of the multicycle, defined by $\sum_{i=1}^{k} a_i \cdot \mathit{effect}(M_i)$, is non-negative for every counter. Here, $\mathit{effect}(M_i)$ is the effect of M_i on the counters. The objective of (I) ensures that the multicycle includes as many transitions as possible, and the total effect of the multicycle is positive on as many counters as possible. For VASS MDPs, the $M_1, \ldots, M_k$ should not be interpreted as cycles but as Markovian strategies for some ECs, and $\mathit{effect}(M_i)$ corresponds to the vector of expected counter changes per transition in M_i. The objective of (I) then maximizes the number of transitions used in the strategies $M_1, \ldots, M_k$, and the number of counters where the expected effect of the "multicycle" is positive.

A solution of (II) for non-probabilistic VASS can be interpreted as a ranking function for configurations defined by $\mathit{rank}(p\mathbf{v}) = \mathbf{z}(p) + \sum_{i=1}^{d} \mathbf{y}(i)\mathbf{v}(i)$, such that the value of rank cannot increase when moving from a configuration $p\mathbf{v}$ to a configuration $q\mathbf{u}$ using a transition $t = (p, \mathbf{u} - \mathbf{v}, q)$. The objective of (II) ensures that as many transitions as possible decrease the value of rank, and rank depends on as many counters as possible. For VASS MDPs, this interpretation changes only for the outgoing transitions $t = (p, \mathbf{u}, q)$ of probabilistic states. Instead of considering the change of *rank* caused by such t, we now consider the expected change of *rank* caused by executing a step from p. The objective ensures that *rank* depends on as many counters as possible, the value of *rank* is decreased by as many outgoing transitions of non-deterministic states as possible, and the expected change of *rank* caused by performing an step is negative in as many probabilistic states as possible.

The key tool for our analysis is the following dichotomy:

▶ **Lemma 7.** *Let* $\mathbf{x}$ *be a (maximal) solution to the constraint system (I) and* $\mathbf{y}, \mathbf{z}$ *be a (maximal) solution to the constraint system (II). Then, for each counter* c *we have that either* $\mathbf{y}(c) > 0$ *or* $\sum_{t\in T} \mathbf{x}(t)\mathbf{u}_t(c) > 0$*, and for each transition* $t = (p, \mathbf{u}, q) \in T$ *we have that*

- *if* $p \in Q_n$ *then either* $\mathbf{z}(q) - \mathbf{z}(p) + \sum_{i=1}^{d} \mathbf{u}(i)\mathbf{y}(i) < 0$ *or* $\mathbf{x}(t) > 0$*;*
- *if* $p \in Q_p$ *then either*

$$\sum_{t'=(p,\mathbf{u}',q')\in Out(p)} P(t')\big(\mathbf{z}(q') - \mathbf{z}(p) + \sum_{i=1}^{d} \mathbf{u}'(i)\mathbf{y}(i)\big) < 0$$

or $\mathbf{x}(t) > 0$.

For the rest of this section, we fix a maximal solution $\mathbf{x}$ of (I) and a maximal solution $\mathbf{y}, \mathbf{z}$ of (II), such that the smallest non-zero element of $\mathbf{y}, \mathbf{z}$ is at least 1. We define a ranking function $rank : C(\mathcal{A}) \to \mathbb{N}$ as $rank(s\mathbf{v}) = \mathbf{z}(s) + \sum_{i=1}^{d} \mathbf{v}(i)\mathbf{y}(i)$.

▶ **Theorem 8.** *For each counter* c*, if* $\mathbf{y}(c) > 0$ *then* n *is a tight estimate of* $\mathcal{C}[c]$ *(for the only type of* $\mathcal{A}$*). Otherwise, i.e., when* $\mathbf{y}(c) = 0$*, the function* n^2 *is a lower estimate of* $\mathcal{C}[c]$*.*

Note that Theorem 8 implies Theorem 6 for strongly connected VASS MDPs. A proof is obtained by combining the following lemmata.

▶ **Lemma 9.** *For every counter* c *such that* $\mathbf{y}(c) > 0$*, every* $\varepsilon > 0$*, every* $p \in Q$*, and every* $\sigma \in \Sigma$*, there exists* n_0 *such that for all* $n \geq n_0$ *we have that* $\mathbb{P}^{\sigma}_{p\mathbf{n}}(\mathcal{C}[c] \geq n^{1+\varepsilon}) \leq kn^{-\varepsilon}$ *where* k *is a constant depending only on* $\mathcal{A}$*.*

For *Targets* $\subseteq C(\mathcal{A})$ and $m \in \mathbb{N}$, we use $\mathrm{Reach}^{\leq m}(\mathit{Targets})$ to denote the set of all computations $\pi = p_0\mathbf{v}_0, p_1\mathbf{v}_1, \ldots$ such that $p_i\mathbf{v}_i \in \mathit{Targets}$ for some $i \leq m$.

▶ **Lemma 10.** *For each counter* c *such that* $\mathbf{y}(c) = 0$ *we have that* $\mathcal{C}_{\exp}[c] \in \Omega(n^2)$ *and* n^2 *is a lower estimate of* $\mathcal{C}[c]$*. Furthermore, for every* $\varepsilon > 0$ *there exist a sequence of strategies* $\sigma_1, \sigma_2, \ldots,$ *a constant* k*, and* $p \in Q$ *such that for every* $0 < \varepsilon' < \varepsilon$*, we have that*

$$\lim_{n\to\infty} \mathbb{P}^{\sigma_n}_{p\mathbf{n}}(\mathrm{Reach}^{\leq kn^{2-\varepsilon'}}(\mathit{Targets}_n)) = 1$$

where $\mathit{Targets}_n = \{q\mathbf{v} \in C(\mathcal{A}) \mid \mathbf{v}(c) \geq n^{2-\varepsilon}$ *for every counter* c *such that* $\mathbf{y}(c) = 0\}$*.*

It remains to prove Theorem 6 for VASS MDPs with DAG-like MEC decomposition. Here, we proceed by analyzing the individual MECs one by one, transferring the output of the previous MEC to the next one. We start in a top MEC with all counters initialized to n. Here we can directly apply Theorem 8 to determine which of the $\mathcal{C}[c]$ have a tight estimate n and a lower estimate n^2, respectively. It follows from Lemma 10 that all counters c such that n^2 is a lover estimate of $\mathcal{C}[c]$ can be simultaneously pumped to $n^{2-\varepsilon}$ with very high probability. However, this computation may decrease the counters c such that n is a tight estimate for $\mathcal{C}[c]$. To ensure that the value of these counters is still $\Omega(n)$ when entering the next MEC, we first divide the initial counter vector $\mathbf{n}$ into two halves, each of size $\lfloor \frac{\mathbf{n}}{2} \rfloor$, and then pump the counters c such that n^2 is a lower estimate for $\mathcal{C}[c]$ to the value $(\lfloor \frac{n}{2} \rfloor)^{2-\varepsilon}$. We show that the length of this computation is at most quadratic. The value of the other counters stays at least $\lfloor \frac{n}{2} \rfloor$. When analyzing the next MEC, we treat the counters previously pumped to quadratic values as "infinite" because they are sufficiently large so that they cannot prevent pumping additional counters to asymptotically quadratic values. Technically, this is implemented by modifying every counter update vector $\mathbf{u}$ so that $\mathbf{u}[c] = 0$ for every "quadratic" counter c. A precise formulation of these observations and the corresponding proofs are given in [2].

We conjecture that the dichotomy of Theorem 6 holds for *all* VASS MDPs, but we do not have a complete proof. If the MEC decomposition is not DAG-like, a careful analysis of computations revisiting the same MECs is required; such repeated visits may but do not have to enable additional asymptotic growth of $\mathcal{C}[c]$.

5 One-Dimensional VASS MDPs

In this section, we give a full and effective classification of tight estimates of $\mathcal{L}$, $\mathcal{C}[c]$, and $\mathcal{T}[t]$ for one-dimensional VASS MDPs. More precisely, we prove the following theorem:

▶ **Theorem 11.** *Let $\mathcal{A}$ be a one-dimensional VASS MDP. We have the following:*

- *Let c be the only counter of $\mathcal{A}$. Then one of the following possibilities holds:*
  - *There exists a type $\beta = M$ such that $\mathcal{C}[c]$ is unbounded for β.*
  - *n is a tight estimate of $\mathcal{C}[c]$ for every type.*
- *Let t be a transition of $\mathcal{A}$. Then one of the following possibilities holds:*
  - *There exists a type $\beta = M$ such that $\mathcal{T}[t]$ is unbounded for β.*
  - *There exists a type β such that $weight(\beta) > 0$ and $\mathcal{T}[t]$ is unbounded for β.*
  - *There exists a type $\beta = M$ such that n^2 is a tight estimate of $\mathcal{T}[t]$ for β.*
  - *The transition t occurs in some MEC M, n is a tight estimate of $\mathcal{T}[t]$ for every type β containing the MEC M, and 0 is a tight estimate of $\mathcal{T}[t]$ for every type β not containing the MEC M.*
  - *The transition t does not occur in any MEC, and for every type β of length k we have that k is an upper estimate of $\mathcal{T}[t]$ for β.*
- *One of the following possibilities holds:*
  - *There exists a type $\beta = M$ such that $\mathcal{L}$ is unbounded for β.*
  - *There exists a type $\beta = M$ such that n^2 is a tight estimate of $\mathcal{L}$ for β.*
  - *n is a tight estimate of $\mathcal{L}$ for every type.*

It is decidable in polynomial time which of the above cases hold.

Note that some cases are mutually exclusive and some may hold simultaneously. Also recall that $weight(\beta) = 1$ for every type β of length one, and $weight(\beta)$ decays exponentially in the length of β. Hence, if a transition t does not occur in any MEC, there is a constant $\kappa < 1$ depending only on $\mathcal{A}$ such that $\mathbb{P}^{\sigma}_{p\mathbf{v}}[\mathcal{T}[t] \geq i] \leq \kappa^i$ for every $\sigma \in \Sigma$ and $p\mathbf{v} \in C(\mathcal{A})$.

For the rest of this section, we fix a one-dimensional VASS MDP $\mathcal{A} = (Q, (Q_n, Q_p), T, P)$ and some linear ordering $\sqsubseteq$ on Q. A proof of Theorem 11 is obtained by analyzing bottom strongly connected components (BSCCs) in a Markov chain obtained from $\mathcal{A}$ by "applying" some MD strategy σ (we use Σ_{MD} to denote the class of all MD strategies for $\mathcal{A}$). Recall that σ selects the same outgoing transition in every $p \in Q_n$ whenever p is revisited, and hence we can "apply" σ to $\mathcal{A}$ by removing the other outgoing transitions. The resulting Markov chain is denoted by $\mathcal{A}_\sigma$. Note that every BSCC $\mathbb{B}$ of $\mathcal{A}_\sigma$ can also be seen as an end component of $\mathcal{A}$. For a MEC M of $\mathcal{A}$, we write $\mathbb{B} \subseteq M$ if all states and transitions of $\mathbb{B}$ are included in M.

For every BSCC $\mathbb{B}$ of $\mathcal{A}_\sigma$, let $p_\mathbb{B}$ be the least state of $\mathbb{B}$ with respect to $\sqsubseteq$. Let $\mathbb{U}_\mathbb{B}$ be a function assigning to every infinite path $\pi = p_0, \mathbf{u}_1, p_1, \mathbf{u}_2, \ldots$ the sum $\sum_{i=1}^{\ell} \mathbf{u}_i$ if $p_0 = p_\mathbb{B}$ and $\ell \geq 1$ is the least index such that $p_\ell = p_\mathbb{B}$, otherwise $\mathbb{U}_\mathbb{B}(\pi) = 0$. Hence, $\mathbb{U}_\mathbb{B}(\pi)$ is the change of the (only) counter c along π until $p_\mathbb{B}$ is revisited.

▶ **Definition 12.** *Let $\mathbb{B}$ be a BSCC of $\mathcal{A}_\sigma$. We say that $\mathbb{B}$ is*

- increasing *if $\mathbb{E}^\sigma_{p_\mathbb{B}}(\mathbb{U}_B) > 0$,*
- decreasing *if $\mathbb{E}^\sigma_{p_\mathbb{B}}(\mathbb{U}_B) < 0$,*
- bounded-zero *if $\mathbb{E}^\sigma_{p_\mathbb{B}}(\mathbb{U}_B) = 0$ and $\mathbb{P}^\sigma_{p_\mathbb{B}}[\mathbb{U}_\mathbb{B}{=}0] = 1$,*
- unbounded-zero *if $\mathbb{E}^\sigma_{p_\mathbb{B}}(\mathbb{U}_B) = 0$ and $\mathbb{P}^\sigma_{p_\mathbb{B}}[\mathbb{U}_\mathbb{B}{=}0] < 1$.*

Note that the above definition does not depend on the concrete choice of $\sqsubseteq$. We prove the following results relating the existence of upper/lower estimates of $\mathcal{L}$, $\mathcal{C}[c]$, and $\mathcal{T}[t]$ to the existence of BSCCs with certain properties. More concretely,

- for $\mathcal{C}[c]$, we show that
  - $\mathcal{C}[c]$ is unbounded for some type $\beta = M$ if there exists an increasing BSCC $\mathbb{B}$ of $\mathcal{A}_\sigma$ for some $\sigma \in \Sigma_{\mathrm{MD}}$ such that $\mathbb{B} \subseteq M$;
  - otherwise, n is a tight estimate of $\mathcal{C}[c]$ for every type.
- for $\mathcal{L}$, we show that
  - $\mathcal{L}$ is unbounded for some type $\beta = M$ if there exists an increasing or bounded-zero BSCC $\mathbb{B}$ of $\mathcal{A}_\sigma$ for some $\sigma \in \Sigma_{\mathrm{MD}}$ such that $\mathbb{B} \subseteq M$;
  - otherwise, n^2 is an upper estimate of $\mathcal{L}$ for every type β;
  - if there exists an unbounded-zero BSCC $\mathbb{B}$ of $\mathcal{A}_\sigma$ for some $\sigma \in \Sigma_{\mathrm{MD}}$, then n^2 is a lower estimate of $\mathcal{L}$ for $\beta = M$ where $\mathbb{B} \subseteq M$;
  - if every BSCC $\mathbb{B}$ of every $\mathcal{A}_\sigma$ is decreasing, then $\mathcal{L}_{\exp}(n) \in \Theta(n)$ (this follows from [3]), and hence n is a tight estimate of $\mathcal{L}$ for every type (Observation 5);
- for $\mathcal{T}[t]$, we distinguish two cases:
  - If t is not contained in any MEC of $\mathcal{A}$, then for every type β of length k, the transition t cannot be executed more than k times along a arbitrary computation π where $mecs(\pi) = \beta$.
  - If t is contained in a MEC M of $\mathcal{A}$, then
    * $\mathcal{T}[t]$ is unbounded for $\beta = M$ if there exist an increasing BSCC $\mathbb{B}$ of $\mathcal{A}_\sigma$ for some $\sigma \in \Sigma_{\mathrm{MD}}$ such that $\mathbb{B} \subseteq M$, or bounded-zero BSCC $\mathbb{B}$ of $\mathcal{A}_\sigma$ for some $\sigma \in \Sigma_{\mathrm{MD}}$ such that $\mathbb{B}$ contains t;
    * $\mathcal{T}[t]$ is unbounded for every $\beta = M_1, \ldots, M_k$ such that $M = M_i$ for some i and there exists an increasing BSCC $\mathbb{B}$ of $\mathcal{A}_\sigma$ for some $\sigma \in \Sigma_{\mathrm{MD}}$ such that $\mathbb{B} \subseteq M_j$ for some $j \leq i$;
    * otherwise, n^2 is an upper estimate of $\mathcal{T}[t]$ for every type;
    * if there is an unbounded-zero BSCC $\mathbb{B}$ of $\mathcal{A}_\sigma$ for some $\sigma \in \Sigma_{\mathrm{MD}}$ such that $\mathbb{B}$ contains t, then n^2 is a lower estimate of $\mathcal{T}[t]$ for $\beta = M$;
    * if every BSCC $\mathbb{B}$ of every $\mathcal{A}_\sigma$ is decreasing, then $\mathcal{T}[t]_{\exp}(n) \in \Theta(n)$ (this follows from [3]), and hence n is an upper estimate of $\mathcal{T}[t]$ for every type (Observation 5).

The polynomial time bound of Theorem 11 is then obtained by realizing the following: First, we need to decide the existence of an increasing BSCC of $\mathcal{A}_\sigma$ for some $\sigma \in \Sigma_{\mathrm{MD}}$. This can be done in polynomial time using the constraint system (I) of Figure 3. If no such increasing BSCC exists, we need to decide the existence of a bounded-zero BSCC, which can be achieved in polynomial time for a subclass of one-dimensional VASS MDPs where no increasing BSCC exists. Then, if no bounded-zero BSCC exists, we need to decide the existence of an unbounded-zero BSCC, which can again be done in polynomial time using the constraint system (I) of Figure 3 (realize that any solution $\mathbf{x}$ of (I) implies the existence of a BSCC that is either increasing, bounded-zero, or unbounded-zero).

Hence, the "algorithmic part" of Theorem 11 is an easy consequence of the above observations, but there is one remarkable subtlety. Note that we need to decide the existence of a bounded-zero BSCC only for a subclass of one-dimensional VASS MDPs where no increasing BSCCs exist. This is actually crucial, because deciding the existence of a bounded-zero BSCC in *general* one-dimensional VASS MDPs is **NP**-complete [2].

The main difficulties requiring novel insights are related to proving the observation about $\mathcal{C}[c]$, stating that if there is no increasing BSCC of $\mathcal{A}_\sigma$ for any $\sigma \in \Sigma_{\mathrm{MD}}$, then n is an upper estimate of $\mathcal{C}[c]$ for every type. A comparably difficult (and in fact closely related) task is to show that if there is no increasing or bounded-zero BSCC, then n^2 is an upper estimate of $\mathcal{L}$ for every type. Note that here we need to analyze the behaviour of $\mathcal{A}$ under *all* strategies (not just MD), and consider the notoriously difficult case when the long-run average change of the counter caused by applying the strategy is zero. Here we need to devise a suitable decomposition technique allowing for interpreting general strategies as "interleavings" of MD strategies and lifting the properties of MD strategies to general strategies. Furthermore, we need to devise techniques for reducing the problems of our interest to analyzing certain types of random walks that have already been studied in stochastic process theory. We discuss this more in the following subsection, and we refer to [2] for a complete exposition of these results.

5.1 MD decomposition

As we already noted, one crucial observation behind Theorem 11 is that if there is no increasing BSCC of $\mathcal{A}_\sigma$ for any $\sigma \in \Sigma_{\mathrm{MD}}$, then n is an upper estimate of $\mathcal{C}[c]$ for every type. In this section, we sketch the main steps towards this result.

First, we show that every path in $\mathcal{A}$ can be decomposed into "interweavings" of paths generated by MD strategies.

Let $\alpha = p_0, \mathbf{v}_1, \ldots, p_k$ be a path. For every $i \leq k$, we use $\alpha_{..i} = p_0, \mathbf{v}_1, \ldots, p_i$ to denote the prefix of α of length i. We say that α is *compatible* with a MD strategy σ if $\sigma(\alpha_{..i}) = (p_i, \mathbf{v}_{i+1}, p_{i+1})$ for all $i < k$ such that $p_i \in Q_n$. Furthermore, for every path $\beta = q_0, \mathbf{u}_1, q_1, \ldots, q_\ell$ such that $p_k = q_0$, we define a path $\alpha \circ \beta = p_0, \mathbf{v}_1, p_1, \ldots, p_k, \mathbf{u}_1, q_1, \ldots, q_\ell$.

▶ **Definition 13.** *Let $\mathcal{A}$ be a VASS MDP, $\pi_1, \ldots, \pi_k \in \Sigma_{\mathrm{MD}}$, and $p_1, \ldots, p_k \in Q$. An* MD-decomposition *of a path $\alpha = s_1, \ldots, s_m$ under $\pi_1, \ldots, \pi_k$ and $p_1, \ldots, p_k$ is a decomposition of α into finitely many paths $\alpha = \gamma_1^1 \circ \cdots \circ \gamma_1^k \circ \gamma_2^1 \circ \cdots \circ \gamma_2^k \circ \cdots \circ \gamma_\ell^1 \circ \cdots \circ \gamma_\ell^k$ satisfying the following conditions:*

- *for all $i < \ell$ and $j \leq k$, the last state of γ_i^j is the same as the first state of γ_{i+1}^j;*
- *for every $j \leq k$, $\gamma_1^j \circ \cdots \circ \gamma_\ell^j$ is a path that begins with p_j and is compatible with π_j.*

Note that $\pi_1, \ldots, \pi_k$ and $p_1, \ldots, p_k$ are not necessarily pairwise different, and the length of γ_i^j can be zero. Also note that the same α may have several MD-decompositions.

Intuitively, an MD decomposition of α shows how to obtain α by repeatedly selecting zero or more transitions by $\pi_1, \ldots, \pi_k$. The next lemma shows that for every VASS MDP $\mathcal{A}$, one can *fix* MD strategies $\pi_1, \ldots, \pi_k$ and states $p_1, \ldots, p_k$ such that *every* path α in $\mathcal{A}$ has an MD-decomposition under $\pi_1, \ldots, \pi_k$ and $p_1, \ldots, p_k$. Furthermore, such a decomposition is constructible *online* as α is read from left to right.

▶ **Lemma 14.** *For every VASS MDP $\mathcal{A}$, there exist $\pi_1, \ldots, \pi_k \in \Sigma_{\mathrm{MD}}$, $p_1, \ldots, p_k \in Q$, and a function $Decomp_{\mathcal{A}}$ such that the following conditions are satisfied for every finite path α:*

- *$Decomp_{\mathcal{A}}(\alpha)$ returns an MD-decomposition of α under $\pi_1, \ldots, \pi_k$ and $p_1, \ldots, p_k$.*
- *$Decomp_{\mathcal{A}}(\alpha) = Decomp_{\mathcal{A}}(\alpha_{..len(\alpha)-1}) \circ \gamma^1 \circ \cdots \circ \gamma^k$, where exactly one of γ^i has positive length (the i is called the* mode *of α).*
- *If the last state of $\alpha_{..len(\alpha)-1}$ is probabilistic, then the mode of α does not depend on the last transition of α.*

According to Lemma 14, *every* strategy σ for $\mathcal{A}$ just performs a certain "interleaving" of the MD strategies $\pi_1, \ldots, \pi_k$ initiated in the states $p_1, \ldots, p_k$. We aim to show that if every BSCC of every $\mathcal{A}_{\pi_j}$ is non-increasing, then n is an upper estimate of $\mathcal{C}[c]$ for every type. Since we do not have any control over the length of the individual γ_i^j occurring in MD-decompositions, we need to introduce another concept of *extended VASS MDPs* where the strategies $\pi_1, \ldots, \pi_k$ can be interleaved in "longer chunks". Intuitively, an extended VASS MDP is obtained from $\mathcal{A}$ by taking k copies of $\mathcal{A}$ sharing the same counter. The j-th copy selects transitions according to π_j. At each round, only one π_j makes a move, where the j is selected by a special type of "pointing" strategy defined especially for extended MDPs. Note that σ can be faithfully simulated in the extended VASS MDP by a pointing strategy that selects the indexes consistently with $Decomp_{\mathcal{A}}$. However, we can also construct another pointing strategy that simulates each π_j longer (i.e., "precomputes" the steps executed by π_j in the future) and thus "close cycles" in the BSCC visited by π_j. This computation can be seen as an interleaving of a finite number of independent random walks with non-positive expectations. Then, we use the optional stopping theorem to get an upper bound on the total expected number of "cycles", which can then be used to obtain the desired upper estimate. We refer to [2] for details.

5.2 A Note about Energy Games

One-dimensional VASS MDPs are closely related to energy games/MDPs [6, 7, 8, 13]. An important open problem for energy games is the complexity of deciding the existence of a *safe* configuration where, for a sufficiently high energy amount, the responsible player can avoid decreasing the energy resource (counter) below zero. This problem is known to be in $\mathbf{NP} \cap \mathbf{coNP}$, and a pseudopolynomial algorithm for the problem exists; however, it is still open whether the problem is in $\mathbf{P}$ when the counter updates are encoded in binary. Our analysis shows that this problem is solvable in polynomial time for energy (i.e., one-dimensional VASS) MDPs $\mathcal{A}$ such that there is no increasing SCC of $\mathcal{A}_\sigma$ for any $\sigma \in \Sigma_{\mathrm{MD}}$.

We say that a SCC $\mathbb{B}$ of $\mathcal{A}_\sigma$ is *non-decreasing* if $\mathbb{B}$ does not contain any negative cycles. Note that every bounded-zero SCC is non-decreasing, and a increasing SCC may but does not have to be non-decreasing.

▶ **Lemma 15.** *An energy MDP has a safe configuration iff there exists a non-decreasing SCC $\mathbb{B}$ of $\mathcal{A}_\sigma$ for some $\sigma \in \Sigma_{\mathrm{MD}}$.*

The "$\Leftarrow$" direction of Lemma 15 is immediate, and the other direction can be proven using our MD decomposition technique, see [2].

Note that if there is no increasing SCC $\mathbb{B}$ of $\mathcal{A}_\sigma$ for any $\sigma \in \Sigma_{\mathrm{MD}}$, then the existence of a non-decreasing SCC is equivalent to the existence of a bounded-zero SCC, and hence it can be decided in polynomial time (see the results presented above). However, for general energy MDPs, the best upper complexity bound for the existence of a non-decreasing SCC is $\mathbf{NP} \cap \mathbf{coNP}$. Interestingly, a small modification of this problem already leads to $\mathbf{NP}$-completeness, as demonstrated by the following lemma.

▶ **Lemma 16.** *The problem whether there exists a non-decreasing SCC $\mathbb{B}$ of $\mathcal{A}_\sigma$ for some $\sigma \in \Sigma_{\mathrm{MD}}$ such that $\mathbb{B}$ contains a given state $p \in Q$ is* **NP***-complete.*

6 Conclusions

We introduced new estimates for measuring the asymptotic complexity of probabilistic programs and their VASS abstractions. We demonstrated the advantages of these measures over the asymptotic analysis of expected values, and we have also shown that tight complexity estimates can be computed efficiently for certain subclasses of VASS MDPs.

A natural continuation of our work is extending the results achieved for one-dimensional VASS MDPs to the multi-dimensional case. In particular, an interesting open question is whether the polynomial asymptotic analysis for non-deterministic VASS presented in [23] can be generalized to VASS MDPs. Since the study of multi-dimensional VASS MDPs is notoriously difficult, a good starting point would be a complete understanding of VASS MDPs with two counters.

References

1. M. Ajdarów and A. Kučera. Deciding polynomial termination complexity for VASS programs. In *Proceedings of CONCUR 2021*, volume 203 of *Leibniz International Proceedings in Informatics*, pages 30:1–30:15. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2021.
2. M Ajdarów and A. Kučera. Asymptotic complexity estimates for probabilistic programs and their vass abstractions. *arXiv*, 2307.04707 [cs.FL], 2023.
3. T. Brázdil, K. Chatterjee, A. Kučera, P. Novotný, and D. Velan. Deciding fast termination for probabilistic VASS with nondeterminism. In *Proceedings of ATVA 2019*, volume 11781 of *Lecture Notes in Computer Science*, pages 462–478. Springer, 2019.
4. T. Brázdil, K. Chatterjee, A. Kučera, P. Novotný, D. Velan, and F. Zuleger. Efficient algorithms for asymptotic bounds on termination time in VASS. In *Proceedings of LICS 2018*, pages 185–194. ACM Press, 2018.
5. M. Broy and M. Wirsing. On the algebraic specification of nondeterministic programming languages. In *Proceedings of CAAP'81*, volume 112 of *Lecture Notes in Computer Science*, pages 162–179. Springer, 1981.
6. K. Chatterjee and L. Doyen. Energy parity games. In *Proceedings of ICALP 2010, Part II*, volume 6199 of *Lecture Notes in Computer Science*, pages 599–610. Springer, 2010.
7. K. Chatterjee and L. Doyen. Energy and mean-payoff parity Markov decision processes. In *Proceedings of MFCS 2011*, volume 6907 of *Lecture Notes in Computer Science*, pages 206–218. Springer, 2011.
8. K. Chatterjee, M. Henzinger, S. Krinninger, and D. Nanongkai. Polynomial-time algorithms for energy games with special weight structures. In *Proceedings of ESA 2012*, volume 7501 of *Lecture Notes in Computer Science*, pages 301–312. Springer, 2012.
9. W. Czerwiński, S. Lasota, R. Lazić, J. Leroux, and F. Mazowiecki. The reachability problem for Petri nets is not elementary. In *Proceedings of STOC 2019*, pages 24–33. ACM Press, 2019.
10. J. Esparza. Decidability of model checking for infinite-state concurrent systems. *Acta Informatica*, 34:85–107, 1997.
11. J.E. Hopcroft and J.-J. Pansiot. On the reachability problem for 5-dimensional vector addition systems. *Theoretical Computer Science*, 8:135–159, 1979.
12. P. Jančar. Undecidability of bisimilarity for Petri nets and some related problems. *Theoretical Computer Science*, 148(2):281–301, 1995.
13. M. Jurdziński, R. Lazić, and S. Schmitz. Fixed-dimensional energy games are in pseudo-polynomial time. In *Proceedings of ICALP 2015*, volume 9135 of *Lecture Notes in Computer Science*, pages 260–272. Springer, 2015.

14 A. Kučera. Algorithmic analysis of termination and counter complexity in vector addition systems with states: A survey of recent results. *ACM SIGLOG News*, 8(4):4–21, 2021.

15 A. Kučera, J. Leroux, and D. Velan. Efficient analysis of VASS termination complexity. In *Proceedings of LICS 2020*, pages 676–688. ACM Press, 2020.

16 J. Leroux. Polynomial vector addition systems with states. In *Proceedings of ICALP 2018*, volume 107 of *Leibniz International Proceedings in Informatics*, pages 134:1–134:13. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2018.

17 J. Leroux and Ph. Schnoebelen. On functions weakly computable by Petri nets and vector addition systems. In *Reachability Problems*, volume 8762 of *Lecture Notes in Computer Science*, pages 190–202. Springer, 2014.

18 R. Lipton. The reachability problem requires exponential space. Technical report 62, Yale University, 1976.

19 E.W. Mayr and A.R. Meyer. The complexity of the finite containment problem for Petri nets. *Journal of the Association for Computing Machinery*, 28(3):561–576, 1981.

20 C.A. Petri. Kommunikation mit automaten. *Schriften des Institutes für Instrumentelle Mathematik*, 3, 1962.

21 M. Sinn, F. Zuleger, and H. Veith. A simple and scalable static analysis for bound analysis and amortized complexity analysis. In *Proceedings of CAV 2014*, volume 8559 of *Lecture Notes in Computer Science*, pages 745–761. Springer, 2013.

22 M. Sinn, F. Zuleger, and H. Veith. Complexity and resource bound analysis of imperative programs using difference constraints. *Journal of Automated Reasoning*, 59(1):3–45, 2017.

23 F. Zuleger. The polynomial complexity of vector addition systems with states. In *Proceedings of FoSSaCS 2020*, volume 12077 of *Lecture Notes in Computer Science*, pages 622–641. Springer, 2020.