Improving Reachability in Vector Addition Systems Through Pumpability

Weijun Chen
BASICS, Shanghai Jiao Tong University, China

Yuxi Fu
BASICS, Shanghai Jiao Tong University, China

Yangluo Zheng
BASICS, Shanghai Jiao Tong University, China

Abstract

Vector addition systems (VAS) constitute an important model of computation and concurrency that is equally expressive as the Petri net model. Recently, a lot of research has been conducted on vector addition systems with states (VASS), which are VASes equipped with a finite state control. Results on VASS naturally carry over to VAS, but no straightforward improvement is available. In this paper, we investigate the reachability problem in VAS in fixed dimensions. Based on a pumpability analysis of VAS that refines Rackoff's extraction for VASS, we obtain an F_{d-2} upper bound for the d-dimensional VAS reachability problem, improving the F_d upper bound inherited from the d-dimensional VASS reachability problem. Low-dimensional VASes are also considered. In particular, we establish a PSPACE upper bound for reachability in 4-dimensional VAS and an ELEMENTARY upper bound for 5-dimensional VAS, while the same upper bounds were known only for 2-VASS and 3-VASS, respectively. The result for 4-VAS particularly hinges on a simplified projection technique developed for geometrically 2-dimensional VASSes, whose reachability problem is shown to be equivalent to 2-VASS.

2012 ACM Subject Classification Theory of computation → Logic and verification; Theory of computation → Concurrency

Keywords and phrases vector addition system, reachability, pumpability

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2026.24

Related Version *Full Version*: https://arxiv.org/abs/2604.24095

Funding All the authors are supported by the NSFC (National Natural Science Foundation of China) grant No. 62572319.

Acknowledgements The authors would like to thank the members of BASICS for their interest in this work. The authors are also grateful to the anonymous reviewers for their valuable comments and efforts.

1 Introduction

Vector addition systems (VAS), being equivalent to Petri nets, are a well-established model of concurrency. A VAS is simply defined by a finite set of integer vectors called actions. Configurations, which are non-negative integer vectors, may move along these actions. The reachability problem, asking whether there is a run from one configuration to another, is the central algorithmic problem in the study of VAS and Petri nets that has found various applications [22]. Decidability of the reachability problem in VAS has been known for over 40 years [15, 9, 11]. Its complexity was settled to be Ackermann-complete only recently [14, 13, 5]. Zooming in to the d-dimensional VAS reachability problem with d fixed, a significant gap still remains between the complexity lower bound $\mathsf{F}_{d/2+O(1)}$ [3] and upper bound F_d [14], where F_d is the d-th level of the fast-growing complexity hierarchy [18].

37th International Conference on Concurrency Theory (CONCUR 2026).
Editors: Ana Sokolova and Patrick Totzke; Article No. 24; pp. 24:1–24:17
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In the literature, the enhanced model "vector addition system with states (VASS)" turns out to be more popular. VASS augments VAS with a finite state control, so that each state has its own set of actions. Indeed, the Ackermann-completeness result mentioned above was proved directly for VASS. Considering the expressive power, Hopcroft and Pansiot [8] showed that the finite state control can be simulated by a VAS with three additional dimensions. In general, this implies problems in VASS are inter-reducible with those in VAS if the dimension is not a fixed parameter. In fixed dimensions, one might naturally expect better complexity bounds for VAS than for VASS. However, such results are not known for most problems. For the d-dimensional VAS reachability problem, the best upper bound we have is F_d inherited from d-dimensional VASS [6]. Gaps between VAS and VASS are clearer in low dimensions. For example, it was shown in [8] that reachability relations in both 2-dimensional VASS and 5-dimensional VAS are effectively semilinear. Inspired by this semilinearity, PSPACE-completeness is established for reachability in 2-VASS [1]. However, for 5-VAS, currently we do not even have an elementary upper bound. This paper is devoted to narrowing the gap between VAS and VASS.

Contribution

In this paper, we mainly focus on the reachability problem in fixed-dimensional VAS. Our first result shows that in the same fixed dimension d, the complexity of reachability in VAS could be lower than that in VASS.

▶ **Theorem 1.1.** *Reachability in $(d+2)$-VAS is in F_d for every $d \geq 3$.*

This result is based on a sharper analysis of pumpability in VAS. A d-dimensional configuration $\boldsymbol{s}$ is said to be pumpable if $\boldsymbol{s}$ can reach some $\boldsymbol{s}' \geq \boldsymbol{s} + \mathbf{1}$. Intuitively, when both the source and the target configurations are pumpable, one can bound the length of the shortest runs between them by a polynomial. On the other hand, when pumpability fails, a known technique called Rackoff's extraction [17] exhibits one bounded coordinate on any run. The dimension can then be reduced by one. In Section 3, we improve this technique and show that for a specific type of VAS, the failure of pumpability yields two bounded coordinates. This allows us to reduce the $(d+2)$-VAS reachability problem to "almost" d-VASS reachability: the reduced VASS might have dimension more than d, but the *geometric dimension* of each of its strongly connected components is at most d. Geometric dimension, originating from [14], is defined as the dimension of the vector space spanned by effects of simple cycles. It is a metric on VASS gaining importance in recent research [6, 4, 21, 7]. A closer inspection of [6] shows that the F_d upper bound for d-VASS indeed holds for VASSes produced by our reduction where each strongly connected component has geometric dimension at most d.

Our second result focuses on low-dimensional VAS, especially 4-VAS and 5-VAS.

▶ **Theorem 1.2.** *Reachability in 5-VAS is in* ELEMENTARY*. Reachability in 4-VAS is in* PSPACE *under binary encoding and is* NL*-complete under unary encoding.*

To the best of our knowledge, this provides the first elementary upper bound for VAS of dimension greater than 3 (we remark that a PSPACE upper bound for 3-VAS could be derived from [4, Lemma 2]). For low-dimensional VAS, we need a fine-grained characterization of shortest runs. We show that the shortest runs in d-dimensional VAS are captured either by $(d-2)$-dimensional VASS, or by a VASS consisting of two sub-VASSes whose geometric dimensions are at most $d-3$. For 4- and 5-VAS, this allows us to use as building blocks the existing results on 2-VASS [1], 3-VASS [4], and geometrically 2-dimensional VASS [6, 21]. We remark that a projection technique converting geometrically 2-dimensional VASS to 2-VASS is also provided in this paper, simplifying that of [21]. The technique produces tight complexity bounds for reachability in geometrically 2-dimensional VASS under unary encoding, which is not answered in [21]. To be specific, we show that this problem is NL-complete when the actual dimension is fixed, and is NP-complete otherwise.

Organization

In Section 2, we fix notation and introduce important concepts in VAS. Section 3 exhibits our main technical contribution – a weaker pumpability condition based on fine-tuning Rackoff's extraction. In Section 4 we apply the pumpability condition to show F_{d-2} upper bound for d-dimensional VAS where $d > 4$. In Section 5 we show a fine-grained analysis of VAS runs using similar techniques. This will be useful in Section 6, where we derive better bounds for 4- and 5-VAS. Finally, Section 7 proposes possible future directions. Proofs omitted from the text can be found in the full version.

2 Preliminaries

Let $\mathbb{N}, \mathbb{Z}, \mathbb{Q}$ denote the sets of natural numbers, integers, and rational numbers, respectively. We use, for example, $\mathbb{Q}_{\geq 0}$ for the set of rationals restricted to non-negative values. Let $m \leq n$ be two integers, we write $[m, n]$ for the set $\{m, m+1, \ldots, n\}$. The interval $[1, n]$ is abbreviated as $[n]$. Vectors are written in boldface $\boldsymbol{u}, \boldsymbol{v}, \boldsymbol{x}, \boldsymbol{y}$, etc. For a vector $\boldsymbol{x} \in \mathbb{Q}^d$, we write $\boldsymbol{x}(i)$ for its i-th component, where $i \in [d]$. We compare vectors component-wise, so that $\boldsymbol{x} \leq \boldsymbol{y}$ means $\boldsymbol{x}(i) \leq \boldsymbol{y}(i)$ for all $i \in [d]$. We use the max norm for vectors, denoted by $\|\boldsymbol{x}\| := \max_{i \in [d]} |\boldsymbol{x}(i)|$. For a finite set $X \subseteq \mathbb{Q}^d$ we write $\|X\| := \max_{\boldsymbol{x} \in X} \|\boldsymbol{x}\|$. The *inner product* of two vectors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{Q}^d$ is written as $\langle \boldsymbol{x}, \boldsymbol{y} \rangle := \sum_{i \in [d]} \boldsymbol{x}(i) \cdot \boldsymbol{y}(i)$. When the dimension d is clear from the context, we write $\boldsymbol{e}_i \in \mathbb{Q}^d$ for the unit vector with a 1 in its i-th coordinate. We write $\mathbf{1} := (1, 1, \ldots, 1)$ and $\mathbf{0} := (0, 0, \ldots, 0)$. For a finite set $X \subseteq \mathbb{Q}^d$ of vectors where $X = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n\}$, the *cone* generated by X is $\mathrm{cone}(X) := \{\lambda_1 \boldsymbol{x}_1 + \cdots + \lambda_n \boldsymbol{x}_n \mid \lambda_1, \ldots, \lambda_n \in \mathbb{Q}_{\geq 0}\}$.

VASS, VAS, and Sequential VAS

A d-dimensional *vector addition system with states* (d-VASS) is a pair $V = (Q, T)$ where Q is a finite set of *states* and $T \subseteq Q \times \mathbb{Z}^d \times Q$ is a finite set of *transitions*. By a *configuration* of V we mean a pair $(p, \boldsymbol{x})$, often denoted by $p(\boldsymbol{x})$, of a state $p \in Q$ and a non-negative vector $\boldsymbol{x} \in \mathbb{N}^d$. A transition $t = (p, \boldsymbol{a}, q)$ induces one-step runs $p(\boldsymbol{x}) \xrightarrow{t} q(\boldsymbol{y})$ where $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{N}^d$ satisfy $\boldsymbol{x} + \boldsymbol{a} = \boldsymbol{y}$. Paths are words over transitions $\pi = (p_1, \boldsymbol{a}_1, q_1) \ldots (p_k, \boldsymbol{a}_k, q_k) =: t_1 \ldots t_k \in T^*$ such that $q_i = p_{i+1}$ for $i \in [k-1]$. If further $q_k = p_1$, we say π is a *cycle* in V. A cycle consisting of a single transition is called a *self-loop*. The effect of such a path π is $\Delta(\pi) := \sum_{i=1}^{k} \boldsymbol{a}_k$. The length of π is denoted by $|\pi| := k$. Runs induced by paths are concatenations of the one-step runs induced by the transitions, e.g.,

$$\pi : p_1(\boldsymbol{x}_1) \xrightarrow{t_1} q_1(\boldsymbol{y_1}) \xrightarrow{t_2} \cdots \xrightarrow{t_k} q_k(\boldsymbol{y}_k). \tag{1}$$

We emphasize that $\boldsymbol{x}_1, \boldsymbol{y}_1, \ldots, \boldsymbol{y}_k \geq \mathbf{0}$ must be non-negative. We write $p(\boldsymbol{x}) \xrightarrow{\pi} q(\boldsymbol{y})$ if there is a run induced by π from $p(\boldsymbol{x})$ to $q(\boldsymbol{y})$. We also write $p(\boldsymbol{x}) \xrightarrow{*} q(\boldsymbol{y})$, quantifying the path existentially. Two paths π_1, π_2 can be concatenated into $\pi_1\pi_2$ if the target state of π_1 matches the source state of π_2. If π is a cycle, we write π^m for the m-fold concatenation of π. Sometimes we also consider $\mathbb{Z}$-configurations $p(\boldsymbol{z})$, where the vector $\boldsymbol{z} \in \mathbb{Z}^d$ can take negative values. We define $\mathbb{Z}$-runs similarly over $\mathbb{Z}$-configurations, denoted by $p(\boldsymbol{z}) \xrightarrow{\pi}_{\mathbb{Z}} q(\boldsymbol{w})$.

Consider a run $\pi = t_1 \dots t_k$ from state p to q. We define the *drop* of π to be the vector $\text{DROP}(\pi) \in \mathbb{N}^d$ where for each $i \in [d]$, $\text{DROP}(\pi)(i)$ is the maximum decrease of the i-th component along this path, i.e.,

$$\text{DROP}(\pi)(i) := -\min(0, \min_{j \in [k]} \Delta(t_1 \dots t_j)(i)). \tag{2}$$

It is straightforward that for any $\boldsymbol{z} \geq \text{DROP}(\pi)$, $p(\boldsymbol{z}) \xrightarrow{\pi} q(\boldsymbol{w})$ holds for some $\boldsymbol{w} \in \mathbb{N}^d$.

The (unary-encoded) *size* of V is defined by $\text{SIZE}(V) := |Q| + d \cdot |T| \cdot (\|T\| + 1)$, where $\|T\| := \max_{(p,\boldsymbol{a},q) \in T} \|\boldsymbol{a}\|$. The *reverse* of V is given by $V^{\text{rev}} = (Q, T^{\text{rev}})$ where $T^{\text{rev}} = \{(q, -\boldsymbol{a}, p) \mid (p, \boldsymbol{a}, q) \in T\}$. The norm of a ($\mathbb{Z}$-)configuration $p(\boldsymbol{x})$ is defined as $\|p(\boldsymbol{x})\| := \|\boldsymbol{x}\|$.

Vector addition systems (VAS) are the main focus of this paper. A d-VAS is a d-VASS $V = (Q, T)$ where Q contains a single state only. Transitions in T are determined by their effects, and we can view T as a finite subset of $\mathbb{Z}^d$. The unique state is often omitted. So a configuration is regarded as a vector in $\mathbb{N}^d$.

In the analysis of VAS runs, we may identify some transitions that can be fired only a bounded number of times. In order to distinguish these transitions from others, we introduce the model of *sequential VAS*. Let V be a d-VAS whose transition set is $A \subseteq \mathbb{Z}^d$, and let $\boldsymbol{a}_1, \dots, \boldsymbol{a}_k \in \mathbb{Z}^d$ be vectors. The sequential VAS $S = V[\boldsymbol{a}_1, \dots, \boldsymbol{a}_k]$ is defined to be the VASS (Q, T) where $Q := \{q_0, q_1, \dots, q_k\}$ and

$$T := \{(q_j, \boldsymbol{a}, q_j) \mid j \in [0, k], \boldsymbol{a} \in A\} \cup \{(q_{j-1}, \boldsymbol{a}_j, q_j) \mid j \in [k]\}. \tag{3}$$

We call q_0 the *source state* and q_k the *target state* in S. Vectors $\boldsymbol{a}_1, \dots, \boldsymbol{a}_k$ are called the *bridges* in S. By a source/target configuration, we mean a configuration whose state is the source/target state in S. Intuitively, $V[\boldsymbol{a}_1, \dots, \boldsymbol{a}_k]$ can be viewed as the VAS V enhanced with k special transitions which can be fired exactly once in the order they are specified.

Geometric Dimension and Reachability Problems

Let V be a d-VASS. Its *cycle space* $\text{CYC}(V) \subseteq \mathbb{Q}^d$ is the vector space spanned by the effects of simple cycles in V, i.e., $\text{CYC}(V) := \text{span}\{\Delta(\theta) \mid \theta \text{ is a simple cycle in } V\}$. Geometric dimension was introduced recently as an important parameter of VASSes. The concept originated from [14]. Here, we consider two types of geometric dimension: The *cycle-dimension* $\dim_{\text{cyc}}(V) := \dim(\text{CYC}(V))$ is the dimension of the whole cycle space. This was also known as *geometric dimension* in [6] and [21]. The *component-dimension* $\dim_{\text{com}}(V) := \max_{\text{SCC } V' \text{ in } V} \dim(\text{CYC}(V'))$ is the maximum value among the dimensions of the cycle spaces of strongly connected components in V. This coincides with the dimension introduced in [7]. By convention, we say that a VASS V is geometrically d-dimensional if $\dim_{\text{cyc}}(V) \leq d$. Notice that for a (sequential) VAS V, both $\dim_{\text{cyc}}(V)$ and $\dim_{\text{com}}(V)$ are equal to the dimension of the vector space spanned by the effects of transitions of V, as every transition is a self-loop. We refer the readers to [2] for further research about these two parameters.

The reachability problem in VASSes can be formulated as follows:

Input: A VASS $V = (Q, T)$ with configurations $p(\boldsymbol{x}), q(\boldsymbol{y})$.
Question: Does it hold $p(\boldsymbol{x}) \xrightarrow{*} q(\boldsymbol{y})$?

Complexity bounds for reachability problems are usually obtained from length bounds on runs witnessing the reachability. Let V be a VASS and s, t be two configurations of V. We also call the triple (V, s, t) a VASS, with s and t understood as the source and target configurations. We write $\text{LEN}(V, s, t)$ for the set of lengths of runs from s to t, and $\text{SIZE}(V, s, t)$ for $\text{SIZE}(V) + d \cdot (\|s\| + \|t\| + 1)$. The following bound can be obtained from the famous KLMST decomposition algorithm.

▶ **Theorem 2.1** ([6, Lemma 6.2]). *Let (V, s, t) be a VASS of component-dimension $m \geq 3$. Then* $\min \text{LEN}(V, s, t) \leq F_m(g(\text{SIZE}(V, s, t)))$ *if* $\text{LEN}(V, s, t) \neq \emptyset$*, where g is some elementary function independent of m, and (V, s, t).*

The function F_m is the m-th fast-growing function, which will be introduced below. We remark that the statement of this theorem in [6] is for fixed dimension m rather than fixed component-dimension m. The version we use here can be easily obtained from the ranking function defined in [6].

Fast-Growing Complexity Hierarchy

The fast-growing functions $(F_k)_{k\in\mathbb{N}}$ are defined as follows: $F_0(x) := x + 1$, $\quad F_{k+1}(x) := F_k^{x+1}(x)$, where F_k^{x+1} means the $(x + 1)$-fold iteration of F_k. For example, $F_2(x) = 2^{x+1}(x + 1) - x$ is exponential in x, and $F_3(x)$ grows even faster than a tower of exponents. We say a function is *elementary* if it is bounded by F_2^c for some constant c. Based on these functions we may define a hierarchy $(\mathcal{F}_k)_{k\in\mathbb{N}}$ of computable functions: $\mathcal{F}_k := \bigcup_{c\in\mathbb{N}} \mathsf{FDTIME}(F_k^c(n))$. The fast-growing complexity hierarchy is defined as follows: $\mathsf{F}_k := \bigcup_{p\in\mathcal{F}_{k-1}} \mathsf{DTIME}(F_k(p(n)))$. For more details of this complexity hierarchy – which we will not use for this paper – we refer the readers to [18].

Wideness, Pumpability and Fixed Coordinates

We have introduced the cycle space of a VASS, which is the vector space spanned by the effects of simple cycles. However, it is the *cone* generated by the effects of simple cycles that better describes the behavior of the VASS. Let V be a VASS and E be the set of effects of simple cycles in V. We define the cone of V as $\text{cone}(V) := \text{cone}(E)$. And V is said to be *wide* if $\text{cone}(V) = \text{CYC}(V)$, so any combination of cycle effects can be expressed using non-negative coefficients. We remark that the notion of wideness was proposed in [4] in a slightly different manner.

Pumpability is another important notion that we will investigate. Let $p(\boldsymbol{x})$ be a configuration in V. We say $p(\boldsymbol{x})$ is *forward pumpable* (or simply *pumpable*) in V if there is a run $p(\boldsymbol{x}) \xrightarrow{*} p(\boldsymbol{x}')$ with $\boldsymbol{x}' \geq \boldsymbol{x} + \mathbf{1}$. We say $p(\boldsymbol{x})$ is *backward pumpable* in V if it is forward pumpable in V^{rev}, or equivalently, if there is a run $p(\boldsymbol{x}') \xrightarrow{*} p(\boldsymbol{x})$ in V with $\boldsymbol{x}' \geq \boldsymbol{x} + \mathbf{1}$.

Let $S = V[\boldsymbol{a}_1, \ldots, \boldsymbol{a}_k]$ be a sequential d-VAS. We say that a coordinate $i \in [d]$ is *fixed* in S if $\boldsymbol{u}(i) = 0$ holds for every transition $\boldsymbol{u}$ in V. Fixed coordinates are annoying from time to time. We may use the following easy lemma to get rid of them.

▶ **Lemma 2.2.** *Let (S, s, t) be a sequential d-VAS. Then there exists a sequential d'-VAS (S', s', t') containing no fixed coordinates such that $d' \leq d$,* $\text{SIZE}(S', s', t') \leq \text{SIZE}(S, s, t)$*, and* $\text{LEN}(S, s, t) = \text{LEN}(S', s', t')$.

Unless otherwise specified, we may assume without loss of generality that the sequential d-VAS in the subsequent discussion contains no fixed coordinates, i.e., there exists at least one transition $\boldsymbol{u}_i$ with $\boldsymbol{u}_i(i) \neq 0$ in V for each $i \in [d]$.

3 Pumpability from $d - 1$ Unbounded Counters

Rackoff's extraction [14, Lemma A.1] (see also [17], [4, Lemma 5], and [10]) is an important technique to extract pumpability or coverability from an "unbounded run". In the scenario of d-VASS runs, this technique states, intuitively, that if there is a run where each of the d coordinates has ever been raised above some sufficiently large value, then one can construct a run from the same source whose target has large values in all d coordinates. A consequence of this fact is that, within one strongly connected component of VASS, at least one coordinate on a run from some unpumpable configuration must remain bounded. Dimension reduction can then be conducted for reachability instances starting from an unpumpable configuration.

In this section, we improve Rackoff's extraction for wide sequential d-VASes. We show that coverability or pumpability is guaranteed if *all but one* coordinates can be raised above some large value in a run. This will give us at least two bounded coordinates in the unpumpable case. Let us now state the result formally. Let π be a run in some d-VASS V and $B \in \mathbb{N}$, and we introduce the notation $\textsc{ub}(\pi, B)$ for the set of indices $i \in [d]$ such that π contains a configuration $q_i(\boldsymbol{x}_i)$ with $\boldsymbol{x}_i(i) \geq B$. That is, $\textsc{ub}(\pi, B)$ are those coordinates that ever reached a value above B along π. Our improved Rackoff's extraction is stated as follows.

▶ **Theorem 3.1.** *There is a polynomial P such that, given a wide sequential d-VAS $S = V[\boldsymbol{a}_1, \ldots, \boldsymbol{a}_k]$ with source configuration $q_0(\boldsymbol{x}_0)$, if there is a run π from $q_0(\boldsymbol{x}_0)$ such that $|\textsc{ub}(\pi, U)| \geq d-1$ where $U := P(\|\boldsymbol{x}_0\|, \textsc{size}(S))^{(d+1)!}$, then there is a run ρ with $|\rho| \leq U$ such that $q_0(\boldsymbol{x}_0) \xrightarrow{\rho} q(\boldsymbol{x})$ for some configuration $q(\boldsymbol{x})$ pumpable in S. Moreover, ρ can be obtained from π by removing some self-loops.*

We remark that the "moreover" part of this theorem guarantees that the target of π is reachable from the pumpable configuration $q(\boldsymbol{x})$ by a $\mathbb{Z}$-run consisting of transitions in V.

First, we show that the pumpability of $q(\boldsymbol{x})$ is guaranteed if $d-1$ coordinates in $\boldsymbol{x}$ are larger than some polynomial bound. The trick is that wideness ensures the existence of a cycle with a positive effect in the remaining coordinate, and this cycle is actually a self-loop in the sequential VAS, which can be fired whenever other coordinates have large values.

▶ **Lemma 3.2.** *Let $S = V[\boldsymbol{a}_1, \ldots, \boldsymbol{a}_k]$ be a wide sequential d-VAS with source configuration $q_0(\boldsymbol{x}_0)$. If there is a run $q_0(\boldsymbol{x}_0) \xrightarrow{*} q(\boldsymbol{x})$ for some state q in S such that $\boldsymbol{x}(i) \geq B(\|\boldsymbol{x}_0\|, \textsc{size}(S))$ for all $i \in [d-1]$, then $q(\boldsymbol{x})$ is pumpable in S, where*

$$B(\|\boldsymbol{x}_0\|, \textsc{size}(S)) := (\|\boldsymbol{x}_0\| + 1 + \textsc{size}(S)) \cdot (1 + \textsc{size}(S)). \tag{4}$$

Proof. Let the states of S be $q_0, q_1, \ldots, q_k$. We expand the run $q_0(\boldsymbol{x}_0) \xrightarrow{*} q(\boldsymbol{x})$ as

$$\pi : q_0(\boldsymbol{x}_0) \xrightarrow{\pi_0} q_0(\boldsymbol{y}_0) \xrightarrow{\boldsymbol{a}_1} q_1(\boldsymbol{x}_1) \xrightarrow{\pi_1} q_1(\boldsymbol{y}_1) \xrightarrow{\boldsymbol{a}_2} \cdots \xrightarrow{\boldsymbol{a}_\ell} q_\ell(\boldsymbol{x}_\ell) \xrightarrow{\pi_\ell} \quad q_\ell(\boldsymbol{y}_\ell) = q(\boldsymbol{x}) \tag{5}$$

for some $\ell \in [0, k]$, where each π_j is a path in the VAS V.

We have assumed that S was wide and has no fixed coordinates (recall Lemma 2.2), and so is V. Therefore, there must be a transition (i.e. a self-loop) $\boldsymbol{t}$ in V such that $\boldsymbol{t}(d) > 0$. Let $H := \|\boldsymbol{x}_0\| + 1 + \textsc{size}(S)$. We claim that $q(\boldsymbol{x}) \xrightarrow{\boldsymbol{t}^H} q(\boldsymbol{x} + H \cdot \boldsymbol{t})$ is a legal run. To see this, one just needs to verify the target $\boldsymbol{x} + H \cdot \boldsymbol{t}$ is non-negative, as all the intermediate vectors on this path are a convex combination of $\boldsymbol{x}$ and $\boldsymbol{x} + H \cdot \boldsymbol{t}$. For $i \in [d-1]$, from $\boldsymbol{x}(i) \geq B(\|\boldsymbol{x}_0\|, \textsc{size}(S))$ we have

$$\boldsymbol{x}(i) + H \cdot \boldsymbol{t}(i) \geq H \cdot (1 + \textsc{size}(S)) - H \cdot \textsc{size}(S) \geq H \geq 0. \tag{6}$$

And for $i = d$, recall that $\boldsymbol{t}(d) > 0$, so $\boldsymbol{x}(d) + H \cdot \boldsymbol{t}(d) \geq H \geq 0$. Now we have justified that $q(\boldsymbol{x}) \xrightarrow{\boldsymbol{t}^H} q(\boldsymbol{x} + H \cdot \boldsymbol{t})$ is a legal run, and moreover that, the target $\boldsymbol{x}' := \boldsymbol{x} + H \cdot \boldsymbol{t}$ satisfies $\boldsymbol{x}'(i) \geq H$ for all $i \in [d]$. Regarding the effect of $q_0(\boldsymbol{x}_0) \xrightarrow{\pi} q(\boldsymbol{x}) \xrightarrow{\boldsymbol{t}^H} q(\boldsymbol{x}')$, we have

$$\Delta(\pi_0 \ldots \pi_\ell \boldsymbol{t}^H) + \sum_{i \in [\ell]} \boldsymbol{a}_i = \boldsymbol{x}' - \boldsymbol{x}_0 \geq (H - \|\boldsymbol{x}_0\|) \cdot \mathbf{1} \geq (\textsc{size}(S) + 1) \cdot \mathbf{1}. \tag{7}$$

Here we slightly abuse the notation, viewing $\pi_1, \dots, \pi_\ell$ as paths in the VAS V, so they can be concatenated. Notice that $\|\boldsymbol{a}_1 + \cdots + \boldsymbol{a}_\ell\| \le \textsc{size}(S)$. Let $\rho := \boldsymbol{t}^H \pi_0 \pi_1 \dots \pi_\ell$ be a cycle in V, we then have $\Delta(\rho) \ge \mathbf{1}$. Next, we show that ρ is fireable from $q(\boldsymbol{x})$.

Indeed, we have shown that the prefix $\boldsymbol{t}^H$ induces a legal run $q(\boldsymbol{x}) \xrightarrow{\boldsymbol{t}^H} q(\boldsymbol{x}')$. It remains to show that $\rho' := \pi_0\pi_1 \dots \pi_\ell$ is fireable from $\boldsymbol{x}'$ in V. Notice that ρ' is obtained from π by removing all the bridges $\boldsymbol{a}_1, \dots, \boldsymbol{a}_\ell$. Since π is fireable from $q_0(\boldsymbol{x}_0)$, we know that $\textsc{drop}(\pi) \le \|\boldsymbol{x}_0\| \cdot \mathbf{1}$. Using the fact that $\|\boldsymbol{a}_1 + \cdots + \boldsymbol{a}_j\| \le \textsc{size}(S)$ for any j, we deduce $\textsc{drop}(\rho') \le \textsc{drop}(\pi) + \textsc{size}(S) \cdot \mathbf{1} \le (\|\boldsymbol{x}_0\| + \textsc{size}(S)) \cdot \mathbf{1} \le H \cdot \mathbf{1} \le \boldsymbol{x}'$. This shows that ρ' is fireable from $\boldsymbol{x}'$. ◀

Lemma 3.2 reduces pumpability to a special kind of coverability, in the sense that only $d-1$ coordinates need to cover some large value simultaneously. The following lemma, which is similar to the standard Rackoff's extraction, further reduces the coverability condition to covering a larger value in each coordinate separately. We remark that the proof has a similar fashion to [4, Lemma 5]. All missing proofs can be found in the full version.

▶ **Lemma 3.3.** *There exists a polynomial R such that, given a d-VASS V and $U \in \mathbb{N}$, for any run π in V, there is a run ρ starting from the same source as π with $|\rho| \le R(U, \textsc{size}(V))^{(d+1)!}$ such that its target $q(\boldsymbol{x})$ satisfies $\boldsymbol{x}(i) \ge U$ for all $i \in \textsc{ub}(\pi, R(U, \textsc{size}(V))^{(d+1)!})$. Moreover, ρ can be obtained from π by removing some cycles.*

Now we can combine these two lemmas to establish Theorem 3.1.

Proof of Theorem 3.1. Let B be the polynomial in Lemma 3.2 and R be the polynomial defined in Lemma 3.3. We may take the polynomial P defined by

$$P(\|\boldsymbol{x}_0\|, \textsc{size}(S)) := R(B(\|\boldsymbol{x}_0\|, \textsc{size}(S)), \textsc{size}(S)). \tag{8}$$

Let π be a run from $q_0(\boldsymbol{x}_0)$ in S such that $|\textsc{ub}(\pi, U)| \ge d-1$. By Lemma 3.3 there is a run ρ with $|\rho| \le P(\|\boldsymbol{x}_0\|, \textsc{size}(S))^{(d+1)!}$ such that $q_0(\boldsymbol{x}_0) \xrightarrow{\rho} q(\boldsymbol{x})$ for some $q(\boldsymbol{x})$ with $\boldsymbol{x}(i) \ge B(\|\boldsymbol{x}_0\|, \textsc{size}(S))$ for all $i \in \textsc{ub}(\pi, U)$. We may assume that $[d-1] \subseteq \textsc{ub}(\pi, U)$ by renaming the coordinates. Therefore by Lemma 3.2 the configuration $q(\boldsymbol{x})$ is pumpable in S. Finally, since each cycle in S consists only of self-loops, ρ can be obtained from π by removing self-loops. ◀

4 Bounding Shortest Runs of d-VAS

In this section, we prove our first main result, that the length of the shortest runs in a $(d+2)$-VAS can be bounded by that in a VASS with component-dimension at most d.

▶ **Theorem 4.1.** *For every $d \in \mathbb{N}$, there is a polynomial P_d such that, for every $(d+2)$-VAS $(V, \boldsymbol{s}, \boldsymbol{t})$, there is a VASS (V', s', t') with $\dim_{\mathrm{com}}(V') \le d$ such that $\textsc{size}(V', s', t') \le P_d(\textsc{size}(V, \boldsymbol{s}, \boldsymbol{t}))$, and $\min \textsc{len}(V, \boldsymbol{s}, \boldsymbol{t}) \le \min \textsc{len}(V', s', t')$.*

If $d \ge 3$, by Theorem 2.1, we can further deduce that the length of shortest reachability witnesses for $(V, \boldsymbol{s}, \boldsymbol{t})$ must be bounded by $F_d(g(\textsc{size}(V, \boldsymbol{s}, \boldsymbol{t})))$, where g is an elementary function. All such runs can be enumerated and checked in $F_d(g(n))$-space. The following improved complexity upper bound follows immediately.

▶ **Theorem 1.1.** *Reachability in $(d+2)$-VAS is in F_d for every $d \ge 3$.*

The pattern of statement in Theorem 4.1 will appear frequently in the following sections. For ease of narration, we shall introduce the following definition.

▶ **Definition 4.2.** *Let $f : \mathbb{N} \to \mathbb{N}$ be a function.*

- *A VASS (V, s, t) is said to be* length-bounded *by a VASS (V', s', t')* with f-amplification*, if* $\min \textsc{len}(V, s, t) \le \min \textsc{len}(V', s', t')$ *and* $\textsc{size}(V', s', t') \le f(\textsc{size}(V, s, t))$*, and $s \xrightarrow{*} t$ in V if and only if $s' \xrightarrow{*} t'$ in V'.*
- *A VASS (V, s, t) is said to be* length-captured *by a family $\mathcal{V}$ of VASSes (V', s', t')* with f-amplification*, if* $\textsc{len}(V, s, t) = \bigcup_{(V',s',t') \in \mathcal{V}} \textsc{len}(V', s', t')$ *and* $\textsc{size}(V', s', t') \le f(\textsc{size}(V, s, t))$ *for all $(V', s', t') \in \mathcal{V}$.*

Hence, Theorem 4.1 could be reformulated as follows:

▶ **Theorem 4.3** (Theorem 4.1)**.** *For every $d \in \mathbb{N}$, there is a polynomial P_d such that, every $(d+2)$-VAS $(V, \boldsymbol{s}, \boldsymbol{t})$ is length-boounded by a VASS (V', s', t') with P_d-amplification, where* $\dim_{\text{com}}(V') \le d$.

Notice that we may always assume a run $s \xrightarrow{\pi} t$ exists in the VASS (V, s, t), as otherwise we can take (V', s', t') to be an unreachable instance with constant size. In the following, we tacitly adopt this assumption.

4.1 Non-Wide VAS to Wide Sequential VAS

We will actually prove Theorem 4.1 for wide sequential VASes. Here, we show how to convert a non-wide VAS to a family of wide sequential VASes.

▶ **Lemma 4.4.** *For every $d \in \mathbb{N}$, there is a polynomial W_d such that, every d-VAS $(V, \boldsymbol{s}, \boldsymbol{t})$ is length-captured by a family $\mathcal{S}$ of wide sequential d-VASes (S, s', t') with W_d-amplification. Moreover, if V is non-wide, then* $\dim_{\text{cyc}}(S) \le \dim_{\text{cyc}}(V) - 1$ *for all $S \in \mathcal{S}$.*

The proof idea is similar to [4, Sect. 4, Case 2]. If V is non-wide, then there exists a vector $\boldsymbol{n}$ such that every transition belongs to the half-space determined by $\boldsymbol{n}$. In particular, the inner product of $\boldsymbol{n}$ with the configuration cannot be decreased by any transition. Hence, every transition whose effect has a non-zero inner product with $\boldsymbol{n}$ can only be fired for a bounded number of times determined by $\langle \boldsymbol{t} - \boldsymbol{s}, \boldsymbol{n} \rangle$. We then enumerate every combination of these transitions and expand V as a sequential VAS correspondingly. What is crucial here is that we need to find a good $\boldsymbol{n}$ such that the VAS with transitions orthogonal to $\boldsymbol{n}$ is wide. For this purpose, we introduce the following lemma.

▶ **Lemma 4.5.** *Let $X \subseteq \mathbb{Z}^d$ be a finite set of vectors such that* $\operatorname{cone}(X) \neq \operatorname{span}(X)$*. There exists a non-zero vector $\boldsymbol{n} \in \mathbb{Z}^d$ such that*

1. $\langle \boldsymbol{n}, \boldsymbol{x} \rangle \ge 0$ *for all $\boldsymbol{x} \in X$;*
2. *let* $X_0 := \{\boldsymbol{x} \in X \mid \langle \boldsymbol{n}, \boldsymbol{x} \rangle = 0\}$*; then* $\operatorname{cone}(X_0) = \operatorname{span}(X_0)$*;*
3. $\|\boldsymbol{n}\| \le (d+1) \cdot (r \, \|X\|)^r$*, where* $r := \dim(\operatorname{span}(X))$.

Notice that such a vector $\boldsymbol{n}$ without the desired bound can be obtained by summing up the rows of the matrix A satisfying $\operatorname{cone}(X) = \{\boldsymbol{x} \in \mathbb{Q}^d \mid A\boldsymbol{x} \ge \boldsymbol{0}\}$. The size bound is obtained by applying the upper bound for solutions to linear systems established in [20].

With the help of Lemma 4.5, we are able to identify the set of bounded transitions. A VAS can then be turned into a wide sequential VAS by taking these transitions as bridges. The formal proof of Lemma 4.4 can be found in the full version.

We remark that after the decomposition presented in the proof of Lemma 4.4, we essentially obtain a family of wide sequential d-VASes, possibly with several fixed coordinates. However, by applying Lemma 2.2, we shall assume that these new VASSes contain no fixed coordinates, and the dimension might be less than d.

4.2 Boundedness of Runs

Now we focus on runs in wide sequential VAS. We identify two cases depending on whether Theorem 3.1 can be applied to extract forward and backward pumpable configurations.

Let π be a run in some d-VASS V and $B \in \mathbb{N}$. Recall that $\textsc{ub}(\pi, B)$ is the set of indices $i \in [d]$ such that π contains a configuration $q_i(\boldsymbol{x}_i)$ with $\boldsymbol{x}_i(i) \geq B$. We say π is *k-bounded* by B if $|\textsc{ub}(\pi_1, B)| \leq d - k$. In the literature of VASS, 1-boundedness is often used to perform dimension reduction. For wide sequential VAS, Theorem 3.1 allows us to shift the focus to 2-boundedness. For technical reasons, we shall make a slight generalization here. We say that π is *almost (2-)bounded by B* if there exists a partition $\pi = \pi_1\pi_2$ such that both π_1 and π_2 are 2-bounded by B. On the other hand, π is *almost unbounded by B* if there exists a partition $\pi = \pi_1\pi_2$ such that $|\textsc{ub}(\pi_1, B)|, |\textsc{ub}(\pi_2, B)| \geq d - 1$. Intuitively, almost bounded runs allow us to perform dimension reduction, while almost unbounded runs permit us to apply Theorem 3.1 to extract forward and backward pumpable configurations. Although these two conditions do not appear complementary, a dichotomy between them persists.

▶ **Lemma 4.6.** *Suppose that (V, s, t) is a d-VASS with $s \xrightarrow{\pi} t$ in V and $B \in \mathbb{N}$. Then π is almost unbounded by B, or π is almost bounded by $B + \textsc{size}(V)$.*

Below, we handle almost unbounded runs and almost bounded runs separately.

4.2.1 Almost Unbounded Runs

In this part, we address almost unbounded runs in wide sequential VASes. We show that in this case, reachability is relatively easy to verify in the sense that a polynomial upper bound can be established on the length of shortest witnesses for such runs.

▶ **Lemma 4.7.** *For every $d \in \mathbb{N}$, there exists a polynomial U_d such that, given a wide sequential d-VAS (S, s, t) with $s \xrightarrow{\pi} t$ in S, if π is almost unbounded by $U_d(\textsc{size}(S, s, t))$, then there exists a run $s \xrightarrow{\rho} t$ in S of length $|\rho| \leq U_d(\textsc{size}(S, s, t))$.*

The proof idea is summarized as follows. If a run π is almost unbounded by a sufficiently large value, then π can be decomposed into two segments $\pi_1\pi_2$ such that, in both segments, at least $d - 1$ coordinates go beyond the given threshold. Recall that the pumping technique presented in Section 3 allows us to construct a $\mathbb{Z}$-run of the form

$$s \xrightarrow{\rho_1} q_1(\boldsymbol{x}_1) \rightarrow_{\mathbb{Z}} q_2(\boldsymbol{x}_2) \xrightarrow{\rho_2} t, \tag{9}$$

where $q_1(\boldsymbol{x}_1)$ and $q_2(\boldsymbol{x}_2)$ are forward pumpable and backward pumpable, respectively. Furthermore, the lengths of ρ_1 and ρ_2 are bounded by a polynomial of $\textsc{size}(S, s, t)$. It suffices to show that the second $\mathbb{Z}$-run can be replaced by a short run. A similar result has been established for strongly connected 3-VASS [4, Sect. 4, Case 1]. The proof consists of two steps: (i) replacing an arbitrary $\mathbb{Z}$-run by a short one coming from an integer linear program, and (ii) lifting the resulting $\mathbb{Z}$-run to sufficiently high via pumping cycles so that it becomes a run. Although the configurations $q_1(\boldsymbol{x}_1)$ and $q_2(\boldsymbol{x}_2)$ may be pumped in different directions, the wideness of S ensures that an appropriate combination of transitions can eventually compensate for this discrepancy. These two steps are realized in the following lemma.

▶ **Lemma 4.8.** *Let $S = V[\boldsymbol{a}_1, \ldots, \boldsymbol{a}_k]$ be a wide sequential d-VAS with source state p and target state q. Let $p(\boldsymbol{x}) \xrightarrow{*}_{\mathbb{Z}} q(\boldsymbol{y})$ be a $\mathbb{Z}$-run in S, where $p(\boldsymbol{x})$ is forward pumpable and $q(\boldsymbol{y})$ is backward pumpable. Then $p(\boldsymbol{x}) \xrightarrow{\rho} q(\boldsymbol{y})$ for some ρ in S with $|\rho| \leq L_d(\textsc{size}(V, p(\boldsymbol{x}), q(\boldsymbol{y})))$, where L_d is a polynomial independent of V, $p(\boldsymbol{x})$, and $q(\boldsymbol{y})$.*

We remark that the wideness is essential in the proof of Lemma 4.8. Moreover, the polynomial bound is obtained by combining the fundamental theorem of linear inequalities [19, Theorem 7.1] with Pottier's theorem [16, Theorem 1]. Next, we proceed to the proof of Lemma 4.7.

Proof of Lemma 4.7. Let P be the polynomial defined in Theorem 3.1 and L_d be the polynomial defined in Lemma 4.8. We take U_d as the following polynomial

$$U_d(x) := 2P(x,x)^{(d+1)!} + L_d\left(2x \cdot P(x,x)^{(d+1)!} + x\right). \tag{10}$$

Let $M := \textsc{size}(S,s,t)$, and $N := P(M,M)^{(d+1)!}$. For any run π almost unbounded by $U_d(M)$, by definition, there exists a configuration $q(\boldsymbol{w})$ such that π can be split into $s \xrightarrow{\pi_1} q(\boldsymbol{w}) \xrightarrow{\pi_2} t$, where $|\,\textsc{ub}(\pi_i, U_d(M))| \geq d-1$ holds for $i = 1, 2$. Note that $U_d(M) \geq N \geq P(\|s\|, \textsc{size}(S))^{(d+1)!}$. Applying Theorem 3.1 to π_1, one can obtain a run ρ_1 with $|\rho_1| \leq P(\|s\|, \textsc{size}(S))^{(d+1)!} \leq N$ such that $s \xrightarrow{\rho_1} q_1(\boldsymbol{x}_1)$ for some forward pumpable configuration $q_1(\boldsymbol{x})$ in S. Moreover, since ρ can be obtained from π by removing some self-loops, we have $q_1(\boldsymbol{x}_1) \xrightarrow{*}_{\mathbb{Z}} q(\boldsymbol{w})$. A symmetric argument to π^{rev} yields that $q(\boldsymbol{w}) \xrightarrow{*}_{\mathbb{Z}} q_2(\boldsymbol{x}_2) \xrightarrow{\rho_2} t$, where $|\rho_2| \leq N$ and $q_2(\boldsymbol{x}_2)$ is backward pumpable in S. Then we can merge these two $\mathbb{Z}$-runs as

$$s \xrightarrow{\rho_1} q_1(\boldsymbol{x}_1) \xrightarrow{*}_{\mathbb{Z}} q_2(\boldsymbol{x}_2) \xrightarrow{\rho_2} t. \tag{11}$$

The sequential VAS S' induced by $q_1(\boldsymbol{x}_1) \xrightarrow{*}_{\mathbb{Z}} q_2(\boldsymbol{x}_2)$ has size at most $M + 2NM$. The configuration $p(\boldsymbol{x})$ is still forward pumpable in S', and $q(\boldsymbol{y})$ is still backward pumpable in S'. Applying Lemma 4.8 to S' and this $\mathbb{Z}$-run, we have $q_1(\boldsymbol{x}_1) \xrightarrow{\rho'} q_2(\boldsymbol{x}_2)$ with $|\rho'| \leq L_d(M + 2NM)$. Let $\rho := \rho_1 \rho' \rho_2$. Then $s \xrightarrow{\rho} t$ with $|\rho| \leq 2N + L_d(M + 2NM) = U_d(M)$. ◀

4.2.2 Almost Bounded Runs

For k-bounded runs, it is standard to encode the k bounded coordinates into states so that they can be captured by a new VASS with lower (geometric) dimension.

▶ **Lemma 4.9.** *Let $B \in \mathbb{N}$ and (V,s,t) be a d-VASS with $s \xrightarrow{\pi} t$ in V, which is k-bounded by B. Then (V,s,t) is length-bounded by a d-VASS (V',s',t') with f-amplication, where* $\dim_{\text{com}}(V') \leq d-k$ *and $f(x) = B^k \cdot x$.*

However, for an almost bounded run π considered here, this dimension reduction technique cannot be applied directly, since π consists of two bounded segments π_1 and π_2. Instead, we construct two new VASSes according to π_1 and π_2, respectively, and then concatenate them. The resulting VASS captures π with a strictly smaller component-dimension.

▶ **Lemma 4.10.** *For each $d \in \mathbb{N}$, there is a polynomial A_d such that, for every $B \in \mathbb{N}$ and $(d+2)$-VASS (V,s,t) with $s \xrightarrow{\pi} t$ in V, if π is almost bounded by B, then (V,s,t) is length-bounded by a $(d+2)$-VASS (V',s',t') with A_d-amplification, where* $\dim_{\text{com}}(V') \leq d$.

The proofs of both lemmas are quite standard. We defer them to the full version for space limitations. So far, we have analyzed both cases of (almost) boundedness. We now return to the proof of Theorem 4.1.

Proof of Theorem 4.1. Let W_d, U_d and A_d be the polynomials defined in Lemma 4.4, Lemma 4.7, and Lemma 4.10, respectively. Note that these polynomials are also nondecreasing with respect to d. We can always adjust A_d such that $A_d(x,y) \geq x+1$ and then define

$$P_d(x) := A_d(U_{d+2}(W_{d+2}(x)) + W_{d+2}(x), W_{d+2}(x)). \tag{12}$$

Let $(V, \boldsymbol{s}, \boldsymbol{t})$ be a $(d+2)$-VASS with $\boldsymbol{s} \xrightarrow{*} \boldsymbol{t}$ and $M := \textsc{size}(V, \boldsymbol{s}, \boldsymbol{t})$. By Lemma 4.4, $(V, \boldsymbol{s}, \boldsymbol{t})$ is length-captured by a family $\mathcal{S}$ of wide sequential VASes with W_d-amplification. Since $\textsc{len}(V, \boldsymbol{s}, \boldsymbol{t}) \neq \emptyset$, there exists a wide sequential d_S-VAS $(S, s, t) \in \mathcal{S}$ with $d_S \leq d+2$ such that $\emptyset \neq \textsc{len}(S, s, t) \subseteq \textsc{len}(V, \boldsymbol{s}, \boldsymbol{t})$, and $\textsc{size}(S, s, t) \leq W := W_{d+2}(M)$. Assume that $s \xrightarrow{*} t$ in S is witnessed by π. Let $U := U_{d_S+2}(W)$. Applying Lemma 4.6 to π, there are two cases.

The run π is almost unbounded by U. Since S is wide and sequential, by Lemma 4.7, there exists a run $s \xrightarrow{\rho} t$ of length $|\rho| \leq U$. We can always construct a 0-VASS V' by concatenating $U+1$ states with source s' and target t'. Then $\textsc{size}(V', s', t') \leq U + 1 \leq A_d(U + \textsc{size}(S), W)$ and $\min \textsc{len}(V', s', t') = U \geq \min \textsc{len}(S, s, t) \geq \min \textsc{len}(V, \boldsymbol{s}, \boldsymbol{t})$.

The run π is almost bounded by. $U + \textsc{size}(S)$. By Lemma 4.10, there is a VASS (V', s', t') with $\dim_{\text{com}}(V') \leq d$ such that $s' \xrightarrow{*} t'$ in V', $\min \textsc{len}(V, \boldsymbol{s}, \boldsymbol{t}) \leq \min \textsc{len}(S, s, t) \leq \min \textsc{len}(V' s', t')$ and $\textsc{size}(V', s', t') \leq A_d(U + \textsc{size}(S), W)$.

In both cases, the size does not exceed $P_d(M)$, which justifies that $(V, \boldsymbol{s}, \boldsymbol{t})$ is length-bounded by a VASS of component dimension at most d with P_d-amplification. ◀

Notes on Geometrically d-Dimensional VASes

So far, we have established Theorem 4.1. A natural extension is to ask what the corresponding result would be if we fix the geometric dimension g of the original VAS rather than its (standard) dimension d. In fact, by taking a closer look at the above proofs, we can easily derive the following conclusion.

▶ **Lemma 4.11.** *For every $g \in \mathbb{N}$, there is a polynomial P_g such that, every d-VAS $(V, \boldsymbol{s}, \boldsymbol{t})$ with $\dim_{\text{cyc}}(V) = g+1$ is length-bounded by a d-VASS (V', s', t') with f-amplification, where $\dim_{\text{com}}(V') \leq g$ and $f(x) = P_g(x)^{(d+1)!}$.*

Proof Sketch. The analysis is analogous to that in Theorem 4.1. For the size bound of the new VASS, the $(d+1)!$ term in the exponent originates from Rackoff's extraction (Lemma 3.3). For the component-dimension of the new VASS, it suffices to show that the dimension reduction in Lemma 4.9 decreases the cycle-dimension of the sequential VAS S by at least one. Note that after encoding any $i \in [d]$ into the state, every cycle θ in the resulting VASS, denoted by V', satisfies $\Delta(\theta)(i) = 0$. In contrast, the original sequential VAS S contains no fixed coordinates, i.e., there exists a self-loop u with $\Delta(u)(i) \neq 0$. Consequently, we have $\textsc{cyc}(V') \subsetneq \textsc{cyc}(S)$, which implies that $\dim_{\text{cyc}}(V') \leq \dim_{\text{cyc}}(S) - 1 \leq g$. ◀

The existence of a 2-bounded run allows us to remove two coordinates of the VAS, which, however, might only decrease the geometric dimension by one. As an example, consider a 3-VAS $(V, \boldsymbol{s}, \boldsymbol{t})$ for which $\boldsymbol{u}(1) + \boldsymbol{u}(2) = 0$ holds for every transition vector $\boldsymbol{u}$ in V. It is straightforward that these two coordinates are always bounded by $\textsc{size}(V, \boldsymbol{s}, \boldsymbol{t})^{O(1)}$ in every run in V. However, removing these two coordinates can only reduce the component-dimension by one, as they are collinear.

Finally, we can combine Theorem 2.1 and Lemma 4.11 to derive the following corollary.

▶ **Corollary 4.12.** *Reachability in geometrically $(g+1)$-dimensional VAS is in F_g for $g \geq 3$.*

5 Fine-Grained Characterization of VAS Runs

Theorem 4.1 gives a nice characterization of shortest runs in d-dimensional VAS using VAS of component-dimension $d-2$. This suffices to prove the F_{d-2} upper-bound for fixed dimension $d \geq 5$. However, some good properties of this characterization are hidden beneath the technical proofs. In this section, we give a finer characterization of runs in d-VAS, relating them to $(d-2)$-VASS or VASS of component-dimension $d-3$, which is formulated as follows.

▶ **Proposition 5.1.** *For each $d \in \mathbb{N}$ where $d \geq 2$, there is a polynomial P_d such that every d-VAS $(V, \boldsymbol{s}, \boldsymbol{t})$ is length-bounded by a VASS (V', s', t') with P_d-amplification. Moreover,*

- *either V' is a $(d-2)$-VASS with $\dim_{\text{cyc}}(V') \leq \dim_{\text{cyc}}(V) - 1$,*
- *or V' is comprised of two d-VASSes V_1, V_2 connected by a single transition, where*

$$\dim_{\text{cyc}}(V_1), \dim_{\text{cyc}}(V_2) \leq \max\{0, \dim_{\text{cyc}}(V) - 3\}. \tag{13}$$

In particular, it holds that $\dim_{\text{com}}(V') \leq \max\{0, \dim_{\text{cyc}}(V) - 3\}$.

This proposition is particularly favorable for low-dimensional VAS compared to Theorem 4.1. For example, Theorem 4.1 bounds the length of shortest runs in 5-VAS by VASS of component-dimension 3, for which we know nothing better than Theorem 2.1. On the other hand, Proposition 5.1 allows us to rely on sharper results for 3-VASS [4] and VASS of component-dimension 2 [6].

Before proving Proposition 5.1, we introduce the notion of *collinear coordinates* in VASes. Let $V = (Q, T)$ be a d-VAS and $i, j \in [d]$. We say i and j are *collinear (in V)* if there exists $\alpha \in \mathbb{Q}$ such that for every transition $\boldsymbol{u} \in T$, we have $\boldsymbol{u}(i) = \alpha \cdot \boldsymbol{u}(j)$. Furthermore, i and j are *positively collinear* if $\alpha > 0$ and *negatively collinear* if $\alpha < 0$. We remark that by Lemma 2.2 we may assume all coordinates are not fixed, so we cannot have $\alpha = 0$ in the case where i and j are collinear. For a sequential VAS $S = V[\boldsymbol{a}_1, \ldots, \boldsymbol{a}_k]$, collinearity is defined with respect to V. We observe that negatively collinear coordinates are always bounded, and can be removed by encoding into states. Formally, we have the following lemma.

▶ **Lemma 5.2.** *There is a polynomial P such that every sequential d-VAS (S, s, t), which contains a pair of negatively collinear coordinates, is length-bounded by a $(d-2)$-VASS (V', s', t') with P-amplification, where $\dim_{\text{cyc}}(V') \leq \dim_{\text{cyc}}(S) - 1$.*

The core idea of proving Proposition 5.1 is to improve the pumping condition in Lemma 3.2. It states that in wide sequential VASes, pumpability always holds in a target configuration where at most one coordinate is bounded. This condition can be weakened to the point that all bounded coordinates are pairwise positively collinear.

▶ **Lemma 5.3.** *Let B be the polynomial given by Lemma 3.2, $S = V[\boldsymbol{a}_1, \ldots, \boldsymbol{a}_k]$ be a wide sequential d-VAS (containing no fixed coordinate) with source configuration $q_0(\boldsymbol{x}_0)$. If there is a run $q_0(\boldsymbol{x}_0) \xrightarrow{*} q(\boldsymbol{x})$ for some state q such that the coordinates in $\{i \in [d] \mid \boldsymbol{x}(i) < B(\|\boldsymbol{x}_0\|, \text{SIZE}(S))\}$ are pairwise positively collinear, then $q(\boldsymbol{x})$ is pumpable in S.*

Proof Sketch. Let $K := \{i \in [d] \mid \boldsymbol{x}(i) < B(\|\boldsymbol{x}_0\|, \text{SIZE}(S))\}$. The proof is almost the same as that of Lemma 3.2 with one difference: we need to find a transition $\boldsymbol{t}$ in V such that $\boldsymbol{t}(i) > 0$ for all $i \in K$ in order to pump up coordinates in K. Since V is wide, there must be a transition $\boldsymbol{t}$ with $\boldsymbol{t}(i) > 0$ for some $i \in K$. As every pair of coordinates in K is positively collinear, the effect of $\boldsymbol{t}$ on these coordinates must be positive simultaneously. ◀

In the following, we prove Proposition 5.1 depending on whether the VAS is wide or not.

5.1 Wide VAS

First, we deal with wide VASes. Recall that Theorem 4.1 was established for wide sequential VASes. There, we introduced the notions of almost boundedness and almost unboundedness to handle non-trivial bridges between VAS components. Since a wide VAS consists of only a single state, it suffices to distinguish runs based on whether they are 2-bounded. This simplification yields the result that the shortest run length in a wide d-VAS can be bounded by that of a $(d-2)$-VASS, rather than a VASS of component-dimension $d-2$.

▶ **Lemma 5.4.** *For each $d \in \mathbb{N}$ where $d \geq 2$, there is a polynomial M_d such that every wide d-VAS $(V, \boldsymbol{s}, \boldsymbol{t})$ with $\dim_{\text{cyc}}(V) \geq 1$ is length-bounded by a $(d-2)$-VASS (V', s', t') with M_d-amplification, where $\dim_{\text{cyc}}(V') \leq \dim_{\text{cyc}}(V) - 1$.*

Proof Sketch. Let P be the polynomial in Theorem 3.1, $M := \textsc{size}(V, \boldsymbol{s}, \boldsymbol{t})$. Pumpability is guaranteed if there is at most one coordinate bounded by $U := P(M, M)^{(d+1)!}$ along π. In this case, Lemma 4.8 bounds the length of shortest runs by a polynomial depending on d. On the other hand, if there are at least two coordinates bounded by U, we may encode these coordinates into states to obtain a $(d-2)$-VASS capturing this run. ◀

5.2 Non-Wide VAS

In the case where a given d-VAS V is non-wide, we still apply Lemma 4.4 to obtain a family of wide sequential VASes. We emphasize again the fact that the cycle-dimension of newly obtained sequential VASes is at most $\dim_{\text{cyc}}(V) - 1$. Compared with the result we want for Proposition 5.1, we still need to eliminate two coordinates or reduce the component-dimension by at least 2. Built upon Lemma 5.2, we focus on those sequential VASes containing no negatively collinear pair of coordinates here, which leads to the following lemma.

▶ **Lemma 5.5.** *For each $d \in \mathbb{N}$, there is a polynomial N_d such that every wide sequential d-VAS (S, s, t) containing no negatively collinear pairs is length-bounded by a d-VASS (V', s', t') with N_d-amplification, where V' is comprised of two d-VASSes V_1, V_2 connected by a single transition, where $\dim_{\text{cyc}}(V_1), \dim_{\text{cyc}}(V_2) \leq \max\{0, \dim_{\text{cyc}}(S) - 2\}$.*

The proof resembles that of Theorem 4.1 with one major difference: we need to make sure that the bounded coordinates in a bounded run are not collinear so that eliminating them can reduce the component-dimension by at least 2. To begin with, we introduce two new notions. Let $B \in \mathbb{N}$. A run π is said to be *quasi-unbounded by B* if it can be split into $\pi = \pi_1\pi_2$, where the coordinates in $[d] \setminus \textsc{ub}(\pi_i, B)$ are pairwise positively collinear for each $i = 1, 2$, and π is *quasi-bounded by B* if $\pi = \pi_1\pi_2$, where $[d] \setminus \textsc{ub}(\pi_i, B)$ contains at least one pair of non-collinear coordinates for each $i = 1, 2$. We have the following dichotomy lemma.

▶ **Lemma 5.6.** *Let (V, s, t) be a VASS with $s \xrightarrow{\pi} t$ and $B \in \mathbb{N}$. If V contains no negatively collinear pairs, then π is quasi-unbounded by B, or quasi-bounded by $B + \textsc{size}(V)$.*

With the help of Lemma 5.3, we can extract pumpability from quasi-unbounded runs. Hence, the following polynomial length bound on quasi-unbounded runs can be obtained. Its proof is just a replay of that of Lemma 4.7.

▶ **Lemma 5.7.** *For every $d \in \mathbb{N}$, there exists a polynomial U_d such that, given a wide sequential d-VAS (S, s, t) with $s \xrightarrow{\pi} t$ in S, if π is quasi-unbounded by $U_d(\textsc{size}(S, s, t))$, then there exists a run $s \xrightarrow{\rho} t$ of length $|\rho| \leq U_d(\textsc{size}(S, s, t))$.*

As for the quasi-bounded case, one can repeat the dimension reduction technique in Lemma 4.10, which captures the two parts $\pi = \pi_1\pi_2$ separately by encoding into states the two bounded coordinates. Quasi-boundedness ensures the existence of two bounded coordinates that are not collinear. Therefore, the cycle-dimension drops by 2 after the encoding. A complete proof of Lemma 5.5 can be found in the full version.

Now we can establish the finer characterization in Proposition 5.1.

Proof of Proposition 5.1. Consider a d-VAS $(V, \boldsymbol{s}, \boldsymbol{t})$ with $\boldsymbol{s} \xrightarrow{*} \boldsymbol{t}$ and let $M := \textsc{size}(V, \boldsymbol{s}, \boldsymbol{t})$. We may abbreviate the two cases in the lemma as the "first characterization" ($(V, \boldsymbol{s}, \boldsymbol{t})$ is length-bounded by a $(d-2)$-VASS (V', s', t') with $\dim_{\mathrm{cyc}}(V') \leq \dim_{\mathrm{cyc}}(V) - 1$) and the "second characterization" ($(V, \boldsymbol{s}, \boldsymbol{t})$ is length-bounded by a d-VASS (V', s', t') such that V' is comprised of two d-VASSes V_1, V_2 connected by a single transition, where $\dim_{\mathrm{cyc}}(V_1), \dim_{\mathrm{cyc}}(V_2) \leq \max\{0, \dim_{\mathrm{cyc}}(V) - 3\}$), respectively.

The case where $\dim_{\mathrm{cyc}}(V) = 0$ falls naturally into the second characterization. Now assume that $\dim_{\mathrm{cyc}}(V) \geq 1$. If V is wide itself, then by Lemma 5.4, we have the first characterization. Otherwise, applying Lemma 4.4, we obtain a wide sequential d'-VAS (S, s, t) where $d' \leq d$ and $\dim_{\mathrm{cyc}}(S) \leq \dim_{\mathrm{cyc}}(V) - 1$, such that $\textsc{size}(S, s, t) = W_d(M)$, $s \xrightarrow{*} t$ in S, and $\min \textsc{len}(V, \boldsymbol{s}, \boldsymbol{t}) \leq \min \textsc{len}(S, s, t)$. If S contains a pair of negatively collinear coordinates, the first characterization follows from Lemma 5.2. Now we may assume that S contains no negatively collinear pairs. Then by Lemma 5.5, there exists a d'-VASS (V', s', t') with $\textsc{size}(V', s', t') \leq N_d(W_d(M))$ such that $s' \xrightarrow{*} t'$ and $\min \textsc{len}(V, \boldsymbol{s}, \boldsymbol{t}) \leq \min \textsc{len}(V, \boldsymbol{s}, \boldsymbol{t})$, where N_d is given in Lemma 5.5. Moreover, V' is comprised of two d-VASSes V_1, V_2 connected by a single transition, where $\dim_{\mathrm{cyc}}(V_1), \dim_{\mathrm{cyc}}(V_2) \leq \max\{0, \dim_{\mathrm{cyc}}(S) - 2\} \leq \max\{0, \dim_{\mathrm{cyc}}(V) - 3\}$. Thus, the second characterization holds. The polynomial P_d is determined implicitly in the above argument. ◀

6 Low Dimensional VAS

In this section, we apply the fine-grained characterization of VAS runs (Proposition 5.1) to prove better complexity bounds for 4- and 5-VAS:

▶ **Theorem 1.2.** *Reachability in 5-VAS is in* ELEMENTARY*. Reachability in 4-VAS is in* PSPACE *under binary encoding and is* NL*-complete under unary encoding.*

The ELEMENTARY upper bound of 5-VAS is almost immediate. With the help of Proposition 5.1, the length of the shortest runs in a 5-VAS is bounded by that in a 3-VASS or a sequence of two geometrically 2-dimensional VASSes of polynomial size amplification. For the 3-VASS, we have a recent triply-exponential bound [4, Lemma 2]; while for the sequence of two geometrically 2-dimensional VASSes, we can convert each one into linear-path schemes [6, Theorem 3.4]. The shortest runs in linear-path schemes are characterized by systems of linear inequalities and thus enjoy an exponential length bound (in terms of unary-encoded size). These analyses sum up to the following lemma.

▶ **Lemma 6.1.** *In a 5-VAS $(V, \boldsymbol{s}, \boldsymbol{t})$, if $\boldsymbol{s} \xrightarrow{*} \boldsymbol{t}$, then $\boldsymbol{s} \xrightarrow{\rho} \boldsymbol{t}$ for some ρ with $|\rho| \leq 2^{2^{2^{\textsc{size}(V,\boldsymbol{s},\boldsymbol{t})^{O(1)}}}}$.*

Proof of Theorem 1.2, for 5-VAS. Enumerating all paths of triply-exponential length (in terms of unary-encoded size) can be done in 3-EXPSPACE (in terms of binary-encoded size). Therefore, we obtain the ELEMENTARY upper bound for 5-VAS. ◀

As for the reachability problem in 4-VAS, we can also apply Lemma 5.5 to capture the shortest runs by either a 2-VAS or a sequence of two geometrically 1-dimensional VASSes. On the one hand, 2-VAS inherits the PSPACE upper bound from 2-VASS [1]. On the other hand, we are not able to apply the linear-path schemes characterization to the sequence of geometrically 1-dimensional VASSes, as it yields merely a doubly-exponential bound in terms of binary-encoded size. Fortunately, the sequence of two geometrically 1-dimensional VASSes form a geometrically 2-dimensional one as a whole. We will develop tools to convert a geometrically 2-dimensional VASS into a real 2-VASS:

▶ **Lemma 6.2.** *There is a polynomial H such that, for every geometrically 2-dimensional d-VASS (V, s, t), there exists a family $\mathcal{V}$ of 2-VASSes (V', s', t') such that $(2d+1) \cdot \text{LEN}(V, s, t) = \bigcup_{(V',s',t') \in \mathcal{V}} \text{LEN}(V', s', t')$, and $\text{SIZE}(V', s', t') \leq H(\text{SIZE}(V, s, t))$.*

Hence, for both the case of a 2-VAS and the case of a geometrically 2-VASS, we are able to bound the length of their shortest runs using known results for 2-VASS [1]. We then establish the length bound for 4-VAS in the following lemma.

▶ **Lemma 6.3.** *For any 4-VAS $(V, \boldsymbol{s}, \boldsymbol{t})$ with $\boldsymbol{s} \xrightarrow{*} \boldsymbol{t}$, it holds that $\min \text{LEN}(V, \boldsymbol{s}, \boldsymbol{t}) \leq \text{SIZE}(V, \boldsymbol{s}, \boldsymbol{t})^{O(1)}$.*

Proof. By Proposition 5.1, the 4-VAS $(V, \boldsymbol{s}, \boldsymbol{t})$ is length-bounded by a VASS (V', s', t') with polynomial amplification, where, V' is either a 2-VASS or a 4-VASS of component-dimension at most 1. For the latter case, it should be noted that V' consists of two geometrically 1-dimensional sub-VASSes, and hence, it is geometrically 2-dimensional as a whole. Applying Lemma 6.2, V' can be converted into a 2-VASS of polynomial size. We conclude the proof as the shortest runs in a 2-VASS have a polynomial length bound [1, Theorem 3.2]. ◀

For any 4-VAS $(V, \boldsymbol{s}, \boldsymbol{t})$ with $\boldsymbol{s} \xrightarrow{*} \boldsymbol{t}$ in V, we deduce that there is a short witness whose length is bounded by $\text{SIZE}(V, \boldsymbol{s}, \boldsymbol{t})^{O(1)}$. A brute-force algorithm searching all such runs can be completed in polynomial space when the input is encoded in binary, and in logarithmic space when the input is encoded in unary. Finally, we are able to finish the proof of Theorem 1.2.

Proof of Theorem 1.2, for 4-VAS. The upper bounds follow from Lemma 6.3. It is shown in [8] that a 0-VASS can be simulated by a 3-VAS with a polynomial overhead. Moreover, the conversion from 0-VASS to unary 3-VAS can be done in logarithmic space. Since the reachability in directed graphs is NL-complete, we conclude that reachability in unary 3- and 4-VAS is NL-complete. ◀

We are left to prove Lemma 6.2. Due to space limitations, this part is postponed to the full version. There, we also study the reachability problem of geometrically 2-dimensional VASS on its own interest and prove the following characterization:

▶ **Theorem 6.4.** *Reachability in geometrically 2-dimensional d-VASS is*

1. PSPACE-*complete under binary encoding;*
2. NP-*complete under unary encoding, with d not fixed; and*
3. NL-*complete under unary encoding, with d fixed.*

We remark that the PSPACE-completeness under binary encoding was already proved in [21], while results for unary encoding are novel and cannot be derived by techniques in [21]. For NP-hardness under unary encoding, we use an idea similar to [12, Lemma 2] to encode binary representation in counters. Notice that for geometrically 2-dimensional VASS under unary encoding with dimension not fixed, we do not have an NL upper bound as in the case for 2-VASS. Indeed, in this case, enumerating configurations cannot fit in logarithmic space (there is an $O(d)$ factor). On the other hand, if the VASS is a VAS, we can actually retain the NL upper bound.

▶ **Theorem 6.5.** *Reachability in geometrically 2-dimensional VAS is in* NL *under unary encoding.*

7 Future Work

Several intriguing questions remain at the boundary of our current understanding. A fundamental open problem is whether the length of runs in d-VAS can be effectively captured by those in $(d-3)$-VASS. Regarding VAS in fixed dimensions, numerous challenges remain unresolved. For binary 2-VAS and 3-VASS, a complexity gap exists as the problems currently reside between an NP-hard lower bound and a PSPACE upper bound. Furthermore, the lower bound for the unary reachability problem in 2-VAS, and more generally, in geometrically 2-dimensional systems where d is not fixed, remains elusive; specifically, it is yet to be determined whether these problems are NL-hard. The case of 5-VAS presents another significant challenge. While the reachability sets of 5-VAS are known to be effectively semilinear, the current 3-EXPSPACE upper bound, derived from the length-boundedness of 3-VASS, appears to be loose. We suggest that a more refined analysis of the geometric constraints in 5-dimensional systems might yield tighter complexity bounds, making the reachability problem in 5-VAS a compelling subject for further exploration.

References

1. Michael Blondin, Matthias Englert, Alain Finkel, Stefan Göller, Christoph Haase, Ranko Lazic, Pierre McKenzie, and Patrick Totzke. The Reachability Problem for Two-Dimensional Vector Addition Systems with States. *J. ACM*, 68(5):34:1–34:43, 2021. doi:10.1145/3464794.
2. Wojciech Czerwinski, Roland Guttenberg, Lukasz Orlikowski, Henry Sinclair-Banks, and Yangluo Zheng. Exploring VASS Parameterised by Geometric Dimension. *CoRR*, abs/2602.15483, 2026. doi:10.48550/arXiv.2602.15483.
3. Wojciech Czerwinski, Ismaël Jecker, Slawomir Lasota, Jérôme Leroux, and Lukasz Orlikowski. New Lower Bounds for Reachability in Vector Addition Systems. In Patricia Bouyer and Srikanth Srinivasan, editors, *43rd IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2023, IIIT Hyderabad, Telangana, India, December 18-20, 2023*, volume 284 of *LIPIcs*, pages 35:1–35:22. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.FSTTCS.2023.35.
4. Wojciech Czerwinski, Ismaël Jecker, Slawomir Lasota, and Lukasz Orlikowski. Reachability in 3-VASS Is Elementary. In Keren Censor-Hillel, Fabrizio Grandoni, Joël Ouaknine, and Gabriele Puppis, editors, *52nd International Colloquium on Automata, Languages, and Programming, ICALP 2025, Aarhus, Denmark, July 8-11, 2025*, volume 334 of *LIPIcs*, pages 153:1–153:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPIcs.ICALP.2025.153.
5. Wojciech Czerwinski and Lukasz Orlikowski. Reachability in Vector Addition Systems is Ackermann-complete. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 1229–1240. IEEE, 2021. doi:10.1109/FOCS52979.2021.00120.
6. Yuxi Fu, Qizhe Yang, and Yangluo Zheng. Improved Algorithm for Reachability in d-VASS. In Karl Bringmann, Martin Grohe, Gabriele Puppis, and Ola Svensson, editors, *51st International Colloquium on Automata, Languages, and Programming, ICALP 2024, Tallinn, Estonia, July 8-12, 2024*, volume 297 of *LIPIcs*, pages 136:1–136:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.ICALP.2024.136.
7. Roland Guttenberg, Wojciech Czerwinski, and Slawomir Lasota. Reachability and Related Problems in Vector Addition Systems with Nested Zero Tests. In *40th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2025, Singapore, June 23-26, 2025*, pages 581–593. IEEE, 2025. doi:10.1109/LICS65433.2025.00050.
8. John E. Hopcroft and Jean-Jacques Pansiot. On the Reachability Problem for 5-Dimensional Vector Addition Systems. *Theor. Comput. Sci.*, 8:135–159, 1979. doi:10.1016/0304-3975(79)90041-0.

9 S. Rao Kosaraju. Decidability of Reachability in Vector Addition Systems (Preliminary Version). In Harry R. Lewis, Barbara B. Simons, Walter A. Burkhard, and Lawrence H. Landweber, editors, *Proceedings of the 14th Annual ACM Symposium on Theory of Computing, May 5-7, 1982, San Francisco, California, USA*, pages 267–281. ACM, 1982. `doi:10.1145/800070.802201`.

10 Marvin Künnemann, Filip Mazowiecki, Lia Schütze, Henry Sinclair-Banks, and Karol Wegrzycki. Coverability in VASS Revisited: Improving Rackoff's Bounds to Obtain Conditional Optimality. *J. ACM*, 72(5):33:1–33:27, 2025. `doi:10.1145/3762178`.

11 Jean-Luc Lambert. A Structure to Decide Reachability in Petri Nets. *Theor. Comput. Sci.*, 99(1):79–104, 1992. `doi:10.1016/0304-3975(92)90173-D`.

12 Jérôme Leroux. Flat Petri Nets (Invited Talk). In Didier Buchs and Josep Carmona, editors, *Application and Theory of Petri Nets and Concurrency - 42nd International Conference, PETRI NETS 2021, Virtual Event, June 23-25, 2021, Proceedings*, volume 12734 of *Lecture Notes in Computer Science*, pages 17–30. Springer, 2021. `doi:10.1007/978-3-030-76983-3_2`.

13 Jérôme Leroux. The Reachability Problem for Petri Nets is Not Primitive Recursive. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 1241–1252. IEEE, 2021. `doi:10.1109/FOCS52979.2021.00121`.

14 Jérôme Leroux and Sylvain Schmitz. Reachability in Vector Addition Systems is Primitive-Recursive in Fixed Dimension. In *34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019, Vancouver, BC, Canada, June 24-27, 2019*, pages 1–13. IEEE, 2019. `doi:10.1109/LICS.2019.8785796`.

15 Ernst W. Mayr. An Algorithm for the General Petri Net Reachability Problem. *SIAM J. Comput.*, 13(3):441–460, 1984. `doi:10.1137/0213029`.

16 Loic Pottier. Minimal Solutions of Linear Diophantine Systems: Bounds and Algorithms. In *International Conference on Rewriting Techniques and Applications*, pages 162–173. Springer, 1991. `doi:10.1007/3-540-53904-2_94`.

17 Charles Rackoff. The Covering and Boundedness Problems for Vector Addition Systems. *Theor. Comput. Sci.*, 6:223–231, 1978. `doi:10.1016/0304-3975(78)90036-1`.

18 Sylvain Schmitz. Complexity Hierarchies beyond Elementary. *ACM Trans. Comput. Theory*, 8(1):3:1–3:36, 2016. `doi:10.1145/2858784`.

19 Alexander Schrijver. *Theory of Linear and Integer Programming*. Wiley-Interscience series in discrete mathematics and optimization. Wiley, 1999.

20 Joachim von zur Gathen and Malte Sieveking. A Bound on Solutions of Linear Integer Equalities and Inequalities. *Proceedings of the American Mathematical Society*, 72(1):155–158, October 1978. `doi:10.1090/s0002-9939-1978-0500555-0`.

21 Yangluo Zheng. Reachability in Vector Addition System with States Parameterized by Geometric Dimension. In Patricia Bouyer and Jaco van de Pol, editors, *36th International Conference on Concurrency Theory, CONCUR 2025, Aarhus, Denmark, August 26-29, 2025*, volume 348 of *LIPIcs*, pages 38:1–38:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. `doi:10.4230/LIPIcs.CONCUR.2025.38`.

22 Richard Zurawski and MengChu Zhou. Petri Nets and Industrial Applications: A Tutorial. *IEEE Trans. Ind. Electron.*, 41(6):567–583, 1994. `doi:10.1109/41.334574`.