Reachability in Vector Addition System with States Parameterized by Geometric Dimension

Yangluo Zheng
BASICS, Shanghai Jiao Tong University, China

Abstract

The geometric dimension of a vector addition system with states (VASS), emerged in Leroux and Schmitz (2019) and formalized by Fu, Yang, and Zheng (2024), quantifies the dimension of the vector space spanned by cycle effects in the system. This paper examines the VASS reachability problem through the lens of geometric dimension, revealing key differences from the traditional dimensional parameterization. Notably, we establish that the reachability problem for both geometrically 1-dimensional and 2-dimensional VASS is PSPACE-complete, achieved by extending the pumping technique initially proposed by Czerwiński et al. (2019).

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

Keywords and phrases Petri net, vector addition system, reachability, geometric dimension, pumping

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2025.38

Related Version *Full Version*: https://arxiv.org/abs/2412.14608

Funding The support from the National Science Foundation of China (62072299), Science and Technology Commission of Shanghai Municipality (24BC3200500) are acknowledged.

Acknowledgements The author would like to thank Wojciech Czerwiński, Łukasz Orlikowski and Qizhe Yang for a fruitful discussion which motivates Section 5. The author also thanks the members of BASICS for their interest. The efforts of anonymous reviewers are appreciated.

1 Introduction

Vector addition systems with states (VASSes), equivalent to Petri nets, serve as a fundamental model for concurrency [10]. A VASS extends finite automata with integer counters that cannot be zero-tested but must be kept non-negative. Central to the algorithmic theory of VASS is the *reachability problem*: determining if a run exists from one configuration to another. Due to its generic nature, numerous practical problems can be modeled via the reachability problem [19]. After decades of study, the computational complexity of the VASS reachability problem was settled to be Ackermann-complete [15, 7, 14]. However, when fixing the *dimension* – the number of counters, a gap remains in complexity bounds. For d-dimensional VASS where $d > 2$, reachability lies in F_d [12], the dth level of the Grzegorczyk hierarchy of complexity classes [17] (for $d = 3$ the upper bound has been improved to 2-EXPSPACE recently [8]), while F_d-hardness is achieved with $(2d + 3)$-dimensional VASS [5]. This gap is known to be closed only in low dimensions. PSPACE-completeness holds for 2-dimensional VASS under binary encoding [3, 6], and NL-completeness under unary encoding [9]. NP-completeness holds for 1-dimensional VASS under binary encoding [13], and again NL-completeness under unary encoding [9]. Dimension has traditionally served as the standard parameterization for reachability.

On the other hand, the structure of *cycles* was identified as a pivotal factor controlling the complexity of the VASS reachability problem [15]. The notion of *geometric dimension*, formalized in [12], measures the dimension of the *cycle space* – vector space spanned by all cycle effects. Insights from [15, 12] suggest that the F_d upper bound of the famous KMLST

36th International Conference on Concurrency Theory (CONCUR 2025).
Editors: Patricia Bouyer and Jaco van de Pol; Article No. 38; pp. 38:1–38:18
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

algorithm applies to geometric dimension d, not only to dimension d. In this work we propose geometric dimension as an alternative parameterization for the VASS reachability problem. In theory the geometric dimension is closer to the nature of the reachability problem, as suggested by [15] and [12], and the cycles provides a rich structure for analysis. Recently a work on 3-dimensional VASS also made use of the geometric dimension [8]. Moreover, the structure of cycles played an important role in a study on continuous VASSes [1] where the authors extended the technique of linear path schemes to study the reachability problem. In practice, as the system parameters might not necessarily be independent of each other, fixing geometric dimension rather than dimension allows one to introduce certain types of interconnections in system parameters for free, which could possibly make the model more expressive. For example, a bounded counter can be simulated by a pair of complementary counters such that transitions increasing one of them must simultaneously decrease the other by the same amount. One easily verifies that these two counters contribute at most 1 to the geometric dimension.

Our contribution

In this paper, we study the reachability problem in VASSes with fixed geometric dimensions. As mentioned before, the F_d upper bound in [12] directly applies to geometric dimension $d \geq 3$. Thus, the primary focus is on VASSes with geometric dimension ≤ 2. Our main contributions are the following theorems.

▶ **Theorem 1.1.** *Reachability in VASS of geometric dimension 2 is* PSPACE*-complete under binary encoding.*

Previous work [12] showing the semi-linearity of VASS reachability set in geometric dimension 2 utilized the technique of linear path schemes, which yielded merely an EXPSPACE upper bound (see [18] for an explicit statement). This is largely due to an exponential blow-up in the number of control states introduced by a dimension-reduction argument (see, e.g. [12, Lemma A.18]). In contrast, our proof relies on a pumping technique for 2-dimensional VASSes [6]. To apply this technique, we make use of the *sign-reflecting projection* proposed in [12], with some further properties developed in Subsection 4.1. Another tool called the *support projection* is introduced in Subsection 4.2. Combining these projection tools we establish a suitable coordinate system (Lemma 4.6) within the 2-dimensinoal cycle space of the VASS. This enables us to apply the arguments in [6] to obtain the PSPACE upper bound. Together with the PSPACE-hardness inherited from 2-dimensinoal VASS [3, 11], we conclude PSPACE-completeness. We mention that the projection tools in this paper do not provide a straightforward reduction from d-dimensional VASS of geometric dimension 2 to 2-dimensinoal VASS. But we will show such a reduction exists for $d = 3$ in Section 5.

Geometric dimensions lower than 2 are also studied in this paper:

▶ **Theorem 1.2.** *Reachability in VASS of geometric dimension 1 is* PSPACE*-complete, and that of geometric dimension 0 is* NP*-complete under binary encoding.*

Results in VASS of geometric dimension 1 and 0 are obtained by a re-examination of known results for VASS of dimension 1 and 2. These results show an interesting distinction in complexity of VASS reachability parameterized by dimension and by geometric dimension, as compared in Table 1.

In addition to the above complexity results, we also give an efficient (polynomial time) algorithm computing the geometric dimension of a VASS in Subsection 3.1.

Table 1 Complexity of VASS reachability parameterized by dimension and by geometric dimension.

	dimension d	geometric dimension d
$d = 0$	NL-complete (folklore)	**NP-complete**
$d = 1$	NP-complete [13]	**PSPACE-complete**
$d = 2$	PSPACE-complete [3]	**PSPACE-complete**
$d = 3$	PSPACE-hard [3], in 2-EXPSPACE[8]	PSPACE-hard [3], in TOWER = F_3 [12]
$d \geq 4$	no known distinctions	

Organization

Section 2 fixes notations and definitions. Section 3 introduces the geometric dimension of a VASS and discusses some useful properties. Section 4 proves the PSPACE-completeness of geometrically 2-dimensional VASS. Section 5 gives a straightforward reduction from geometrically 2-dimensional 3-VASS to 2-VASS. Section 6 considers geometric dimensions lower than 2. Section 7 concludes the paper. Omitted proofs can be found in the full version.

2 Preliminaries

We use $\mathbb{N}, \mathbb{Z}, \mathbb{Q}$ to denote the set of natural numbers (non-negative integers), integers, and rational numbers respectively. Let $m \leq n$ be integers, we use $[m, n]$ to denote the set $\{m, m+1, \ldots, n\}$. And we abbreviate $[n]$ for $[1, n]$. For a d-dimensional vector $\boldsymbol{v} \in \mathbb{Q}^d$, we write $\boldsymbol{v}(i)$ for its ith component, and we use its maximum norm $\|\boldsymbol{v}\| := \max_{i\in[d]} |\boldsymbol{v}(i)|$. The order $\leq$ is extended component-wise to vectors: we write $\boldsymbol{u} \leq \boldsymbol{v}$ if $\boldsymbol{u}(i) \leq \boldsymbol{v}(i)$ for all $i \in [d]$. Similarly we define the component-wise strict order $<$ for vectors. We write $\langle \boldsymbol{u}, \boldsymbol{v}\rangle = \sum_{i\in[d]} \boldsymbol{u}(i)\boldsymbol{v}(i)$ for their inner product. The *support* of a vector $\boldsymbol{v}$ is $\text{supp}(\boldsymbol{v}) := \{i \in [d], \boldsymbol{v}(i) \neq 0\}$. The support of a set S of vectors is $\text{supp}(S) := \bigcup_{\boldsymbol{v}\in S} \text{supp}(\boldsymbol{v})$. A vector $\boldsymbol{v}$ is *positive* if $\boldsymbol{v} > \boldsymbol{0}$, it is *semi-positive* if $\boldsymbol{v} \geq \boldsymbol{0}$ and $\boldsymbol{v} \neq \boldsymbol{0}$. For a string $s = a_1 a_2 \ldots a_n \in \Sigma^*$ over an alphabet Σ, we write $s[i..j]$ for the substring $a_i a_{i+1} \ldots a_j$ of s.

2.1 Vector Addition System with States

Let $d \geq 0$ be an integer. A *d-dimensional vector addition system with states (d-VASS)* is a pair $G = (Q, T)$ where Q is a finite set of *states* and $T \subseteq Q \times \mathbb{Z}^d \times Q$ is a finite set of *transitions*. Clearly a VASS can also be viewed as a directed graph with edges labelled by integer vectors. Given a word $\pi = (p_1, \boldsymbol{a}_1, q_1)(p_2, \boldsymbol{a}_2, q_2) \ldots (p_n, \boldsymbol{a}_n, q_n) \in T^*$ over transitions, we say that π is a *path from p_1 to q_n* if $q_i = p_{i+1}$ for all $i = 1, \ldots, n-1$. It is a *cycle* if we further have $p_1 = q_n$. Such a path is usually presented in the following form:

$$\pi = p_1 \xrightarrow{t_1} q_1 \xrightarrow{t_2} q_2 \xrightarrow{t_3} \cdots \xrightarrow{t_n} q_n \tag{1}$$

where $t_i = (p_i, \boldsymbol{a}_i, q_i)$. The *effect* of π is defined to be $\Delta(\pi) := \sum_{i=1}^n \boldsymbol{a}_i$.

Size, traversal number and characteristic

The norm of a transition $t = (p, \boldsymbol{a}, q)$ is defined by $\|t\| := \|\boldsymbol{a}\|$. For a d-VASS $G = (Q, T)$ we write $\|T\| := \max\{\|t\| : t \in T\}$. We shall mainly consider VASS under binary encoding, so the *size* of G is given by $|G| := |Q| + d \cdot |T| \cdot \lceil \log(\|T\| + 1) \rceil + 1$.

We define the *traversal number* of G to be the maximal number of distinct states that can be visited (traversed) by a path in G, denoted by $\varsigma(G)$. We remark that $\varsigma(G)$ is an upper bound of (i) the length of any simple path/cycle, and (ii) the number of connected components visited by any path. Note also the trivial fact $\varsigma(G) \leq |Q|$.

The *characteristic* of G, denoted by $\chi(G)$, is defined to be $\chi(G) := \varsigma(G) \cdot \|T\|$. So $\chi(G)$ upper bounds the norm of effect of any simple path/cycle in G.

Semantics of VASSes

Let $G = (Q, T)$ be a d-VASS. A *configuration* c of G is a pair of a state $p \in Q$ and a vector $\boldsymbol{v} \in \mathbb{N}^d$, written as $c = p(\boldsymbol{v})$. We will often confuse a configuration with its vector. So we shall write, for example, $\|c\|$ for $\|\boldsymbol{v}\|$, and $c(i)$ for $\boldsymbol{v}(i)$. The semantics of G is defined as follows. For each transition $t = (p, \boldsymbol{a}, q) \in T$, the one-step transition relation $\xrightarrow{t}$ relates all pairs of configurations of the form $(p(\boldsymbol{u}), q(\boldsymbol{v}))$ where $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{N}^d$ and $\boldsymbol{v} = \boldsymbol{u} + \boldsymbol{a}$. Then for a word $\pi = t_1 t_2 \dots t_n \in T^*$, the relation $\xrightarrow{\pi}$ is the composition $\xrightarrow{\pi} := \xrightarrow{t_1} \circ \cdots \circ \xrightarrow{t_n}$. So $p(\boldsymbol{u}) \xrightarrow{\pi} q(\boldsymbol{v})$ if and only if there are configurations $p_0(\boldsymbol{u}_0), \dots, p_n(\boldsymbol{u}_n) \in Q \times \mathbb{N}^d$ such that

$$p(\boldsymbol{u}) = p_0(\boldsymbol{u}_0) \xrightarrow{t_1} p_1(\boldsymbol{u}_1) \xrightarrow{t_2} \cdots \xrightarrow{t_n} p_n(\boldsymbol{u}_n) = q(\boldsymbol{v}). \tag{2}$$

Also, when $\pi = \epsilon$ is the empty word, the relation $\xrightarrow{\epsilon}$ is the identity relation over $Q \times \mathbb{N}^d$. Note that $\xrightarrow{\pi}$ is non-empty only if π is a path. When $p(\boldsymbol{u}) \xrightarrow{\pi} q(\boldsymbol{v})$ we also say that π induces (or is) a *run* from $p(\boldsymbol{u})$ to $q(\boldsymbol{v})$, and we write $\mathrm{src}(\pi) := p(\boldsymbol{u})$ and $\mathrm{trg}(\pi) := q(\boldsymbol{v})$ for the *source* and *target* of π. We emphasize that all configurations on this run lie in $\mathbb{N}^d$, and that they are uniquely determined by $p(\boldsymbol{u})$ and π. Finally, the *reachability relation* of G is defined to be $\xrightarrow{*} := \bigcup_{\pi \in T^*} \xrightarrow{\pi}$.

Reachability problem

The general reachability problem in VASS is formulated as follows:

<u>REACHABILITY IN VECTOR ADDITION SYSTEM WITH STATES</u>
Input: A VASS $G = (Q, T)$, two configurations $p(\boldsymbol{u}), q(\boldsymbol{v})$ of G.
Question: Does $p(\boldsymbol{u}) \xrightarrow{*} q(\boldsymbol{v})$ hold in G?

It is a folklore that this problem can be reduced to the following one in polynomial time without affecting the dimension or the geometric dimension:

<u>**0**-REACHABILITY IN VECTOR ADDITION SYSTEM WITH STATES</u>
Input: A $G = (Q, T)$, two states $p, q \in Q$.
Question: Does $p(\mathbf{0}) \xrightarrow{*} q(\mathbf{0})$ hold in G?

Thus, in this paper we shall mainly care about runs starting from $\mathbf{0}$ and ending at $\mathbf{0}$. Such a run is called a $\mathbf{0}$-*run* in the following.

Reverse of VASS

For a VASS $G = (Q, T)$ we define its *reverse* as the VASS $G^{\mathrm{rev}} = (Q, T^{\mathrm{rev}})$ where $T^{\mathrm{rev}} := \{(q, -\boldsymbol{a}, p) : (p, \boldsymbol{a}, q) \in T\}$. The *reverse* $\mathrm{rev}(\pi)$ of a path (or a run) π is defined naturally by reversing the order of transitions in π, switching the source and target states of those transitions, and negating their effects. We note that $p(\boldsymbol{u}) \xrightarrow{\pi} q(\boldsymbol{v})$ in G if and only if $q(\boldsymbol{v}) \xrightarrow{\mathrm{rev}(\pi)} p(\boldsymbol{u})$ in G^{rev}.

3 Geometric Dimension

▶ **Definition 3.1.** *Let G be a d-VASS. The* cycle space *of G is the vector space* $\mathrm{Cyc}(G) \subseteq \mathbb{Q}^d$ *spanned by the effects of all cycles in G, that is:* $\mathrm{Cyc}(G) := \mathrm{span}\{\Delta(\theta) : \theta \text{ is a cycle in } G\}$.

The dimension of the cycle space of G is called the geometric dimension *of G, denoted by* $\mathrm{gdim}(G) := \dim(\mathrm{Cyc}(G))$. *We say G is* geometrically k-dimensional *if* $\mathrm{gdim}(G) \le k$.

It should be noticed that the cycle space of a VASS is indeed spanned by the effects of all simple cycles in it.

▶ **Lemma 3.2.** *Let G be a d-VASS, then $\mathrm{Cyc}(G)$ equals to the vector space spanned by the effects of all simple cycles in G, that is,* $\mathrm{Cyc}(G) = \mathrm{span}\{\Delta(\beta) : \beta \text{ is a simple cycle in } G\}$.

Proof. Just note that the effect of every cycle is a finite sum of effects of simple cycles. ◀

A naive algorithm that computes the geometric dimension of a VASS by enumerating all simple cycles in it requires PSPACE. Indeed, we show here a more efficient algorithm which computes $\mathrm{gdim}(G)$ in polynomial time.

3.1 Computing Geometric Dimension

We present a stronger algorithm that given a VASS G as input, computes a basis for $\mathrm{Cyc}(G)$. Observe that every cycle lies within some maximal strongly connected component (SCC) of G. Once we have an algorithm that computes a basis for the cycle spaces of every SCC of G, a basis for $\mathrm{Cyc}(G)$ is just a maximal linearly independent subset of the union of these bases, which can be computed by Gaussian elimination in polynomial time. So the problem reduces to computing a basis for $\mathrm{Cyc}(G)$ in case G is strongly connected.

Fix a strongly connected d-VASS $G = (Q, T)$. We introduce an operation called *cycle shrinking*. Let θ be a simple cycle in G that is not a self-loop and has the form

$$\theta = p_0 \xrightarrow{t_1} p_1 \xrightarrow{t_2} \cdots \xrightarrow{t_n} p_n = p_0. \tag{3}$$

First we define a "shift function" $s : Q \to \mathbb{Z}^d$ as follows. If $q = p_k$ for some $k \in [n]$, we set $s(q) := \Delta(\theta[1..k]) = \Delta(t_1) + \ldots + \Delta(t_k)$; otherwise we set $s(q) := \mathbf{0}$. Note that $s(p_0) = s(p_n) = \Delta(\theta)$. Let $P = \{p_1, \ldots, p_n\}$ be the set of states that occurs on θ. Now we define a new VASS $G/\theta = (Q^\theta, T^\theta)$ that "shrinks" θ into a single states as follows.

- $Q^\theta := (Q \setminus P) \cup \{\theta\}$.
- Let $h : Q \to Q^\theta$ be defined as $h(q) = \theta$ if $q \in P$ and $h(q) = q$ otherwise. The transitions are given by $T^\theta := \{(h(p), s(p) + \boldsymbol{a} - s(q), h(q)) : (p, \boldsymbol{a}, q) \in T\}$.

It should be clear that $|Q^\theta| < |Q|$ as θ is not a self-loop, and that G/θ can be constructed in polynomial time given G and θ. Note that $\|T^\theta\| \le (2|Q| + 1)\,\|T\|$, so the size of G/θ is bounded by $|G/\theta| \le |G| + d \cdot |T| \cdot \lceil \log(2|Q| + 1) \rceil$. Observe that h is a graph homomorphism from G onto G/θ. Thus G/θ is strongly connected as long as G is. Besides, we can show that G/θ preserves the cycle space of G.

▶ **Proposition 3.3.** $\mathrm{Cyc}(G/\theta) = \mathrm{Cyc}(G)$.

Now the algorithm for computing a basis of $\mathrm{Cyc}(G)$ where G is strongly connected should be clear. As listed in Algorithm 1, we repeatedly shrink a cycle in G until there remains only one state. Then its cycle space is the span of effects of self-loops in it. A basis of $\mathrm{Cyc}(G)$ can be computed using Gaussian elimination. Since a cycle shrinking reduces the number of states by at least one, after at most $|Q|$ iterations we must stop with a single state remained. Note that the size of VASS is always bounded by $|G| + d \cdot |Q| \cdot |T| \lceil \log(2|Q| + 1) \rceil \le |G|^2$ in each iteration. Thus the algorithm runs in polynomial time.

Algorithm 1 CycleSpaceBasis.

input : a VASS G which is strongly connected
output : basis of $\mathrm{Cyc}(G)$

```
while G contains more than one state do
    θ ← a simple cycle in G that is not a self-loop
    G ← G/θ
end
U ← the set of effects of all self-loops in G
return a basis of U found by Gaussian elimination
```

3.2 Geometry of Reachability Sets and Runs

Given a VASS $G = (Q, T)$ and a configuration $p(\boldsymbol{u}) \in Q \times \mathbb{N}^d$, we write $\mathrm{Reach}_G(p(\boldsymbol{u}))$ for all configurations that is reachable from $p(\boldsymbol{u})$: $\mathrm{Reach}_G(p(\boldsymbol{u})) := \{q(\boldsymbol{v}) \in Q \times \mathbb{N}^d : p(\boldsymbol{u}) \xrightarrow{*} q(\boldsymbol{v})\}$. The next lemma shows that the "dimension" of any reachable set is bounded by $\mathrm{gdim}(G)$, in the sense that it is contained in a finite union of affine copies of $\mathrm{Cyc}(G)$. Here the sum of a vector $\boldsymbol{v} \in \mathbb{Q}^d$ and a set $S \subseteq \mathbb{Q}^d$ is defined as $\boldsymbol{v} + S := \{\boldsymbol{v} + \boldsymbol{s} : \boldsymbol{s} \in S\}$.

▶ **Lemma 3.4.** *Let $G = (Q, T)$ be a d-VASS, $p(\boldsymbol{u}) \in Q \times \mathbb{N}^d$ be a configuration of G. Then*

$$\mathrm{Reach}_G(p(\boldsymbol{u})) \subseteq Q \times \bigcup_{\substack{\boldsymbol{z} \in \mathbb{Z}^d \\ \|\boldsymbol{z}\| \le \chi(G)}} \boldsymbol{u} + \mathrm{Cyc}(G) + \boldsymbol{z}. \tag{4}$$

In other words, for any configuration $q(\boldsymbol{v}) \in Q \times \mathbb{N}^d$ with $p(\boldsymbol{u}) \xrightarrow{*} q(\boldsymbol{v})$, we have $\boldsymbol{v} = \boldsymbol{u} + \boldsymbol{c} + \boldsymbol{z}$ for some $\boldsymbol{c} \in \mathrm{Cyc}(G)$ and $\boldsymbol{z} \in \mathbb{Z}^d$, where $\|\boldsymbol{z}\| \le \chi(G)$.

Note that one may need exponentially many (roughly $O(\chi(G)^d)$) affine copies of $\mathrm{Cyc}(G)$ to cover $\mathrm{Reach}_G(p(\boldsymbol{u}))$. The next lemma shows that any fixed run from $p(\boldsymbol{u})$ is confined in, however, at most $|Q|$ affine copies of $\mathrm{Cyc}(G)$.

▶ **Lemma 3.5.** *Let $G = (Q, T)$ be a d-VASS. For any run π in G with source $p(\boldsymbol{u}) \in Q \times \mathbb{N}^d$, there is a function $f_\pi : Q \to \mathbb{Z}^d$ such that for every configuration $q(\boldsymbol{v})$ occurring on π, we have $\boldsymbol{v} \in \boldsymbol{u} + \mathrm{Cyc}(G) + f_\pi(q)$. Moreover, $\|f_\pi(q)\| \le \chi(G)$ for every $q \in Q$.*

These two lemmas follow easily from the fact that we can view a run as a simple path interleaved with cycles. Their proofs can be found in the full version of this paper.

4 Geometrically 2-Dimensional VASS

In this section we focus exclusively on geometrically 2-dimensional VASSes. We prove that reachability in geometrically 2-dimensional VASSes is PSPACE-complete. The lower bound is a simple corollary of [3, Lemma 20], so most effort will be devoted to the upper bound. Our proof is based on the pumping technique proposed in [6] for 2-VASSes, where they showed that every run in a 2-VASS is either *thin* – confined in some belt-shaped regions, or *thick* – enjoying good pumping properties that can be exploited to shrink long runs. We extend this technique and prove a similar thin-thick classification for runs in geometrically 2-dimensional VASSes. This will enable us to obtain the following exponential length bound for reachability witnesses.

▶ **Theorem 4.1.** *For any* $\mathbf{0}$*-run* τ *in a geometrically 2-dimensional* d*-VASS* G*, there is a* $\mathbf{0}$*-run* ρ *in* G *with the same source and target states as* τ *and* $|\rho| \leq \chi(G)^{O(\varsigma(G)\cdot d^4)}$*.*

From the above theorem the PSPACE-completeness follows immediately.

▶ **Theorem 4.2.** *Reachability in geometrically 2-dimensional VASS is* PSPACE*-complete.*

Proof. The lower bound is inherited from the PSPACE-hardness of reachability in 2-VASS [3, Lemma 20]. For the upper bound, an algorithm only need to search a run of length up to $\chi(G)^{O(\varsigma(G)\cdot d^4)}$ after reducing to the $\mathbf{0}$-reachability problem, for which PSPACE is enough. ◀

In order to establish the thin-thick classification of runs in a geometrically 2-dimensinoal VASS, we need to create a suitable coordinate system within its cycle space. This is achieved using the projection tools developed in Subsection 4.1 and Subsection 4.2. Depending on whether these tools can be applied, we classify geometrically 2-dimensinoal VASSes into degenerate ones (where projection tools are not applicable) and proper ones (where projection tools are applicable). Degenerate VASSes are easier to handle as they only admit thin runs, see Subsection 4.3. For proper VASSes in Subsection 4.4 we adapt the argument of [6] within the coordinate system created by the projection tools and establish the thin-thick classification. Finally we recall in Subsection 4.5 that both thin runs and thick runs can be shrunk to exponential length, thus Theorem 4.1 follows.

4.1 Sign Reflecting Projection

An *orthant* is one of the 2^d regions in $\mathbb{Q}^d$ split by the d axes. Formally, given a vector $t \in \{+1, -1\}^d$, the orthant Z_t defined by t is the set $Z_t := \{\boldsymbol{u} \in \mathbb{Q}^d : \boldsymbol{u}(i) \cdot t(i) \geq 0 \text{ for all } i \in [d]\}$. The non-negative orthant $\mathbb{Q}^d_{\geq 0} = Z_{(+1,+1,\dots,+1)}$ is a major concern in this paper.

Let $I \subseteq [d]$ be a subset of indices. For a vector $\boldsymbol{u} \in \mathbb{Q}^d$, we define its *projection onto indices in* I as a function $\boldsymbol{u}|_I \in \mathbb{Q}^I$ given by $(\boldsymbol{u}|_I)(i) = \boldsymbol{u}(i)$ for all $i \in I$. We tacitly identify the function $\boldsymbol{u}|_I \in \mathbb{Q}^I$ as a vector in $\mathbb{Q}^{|I|}$. For a set of vectors $V \subseteq \mathbb{Q}^d$, we define $V|_I := \left\{\boldsymbol{v}|_I : \boldsymbol{v} \in V\right\}$. It should be clear that the projection of a vector space onto indices in I is again a vector space in $\mathbb{Q}^{|I|}$, and the projection of an orthant Z_t onto I is an orthant in $\mathbb{Q}^{|I|}$ defined by $t|_I$.

▶ **Definition 4.3** ([12, Definition A.7])**.** *Let* $P \subseteq \mathbb{Q}^d$ *be a vector space and* Z *be an orthant in* $\mathbb{Q}^d$*. A set of indices* $I \subseteq [d]$ *is called a* sign-reflecting projection for P with respect to Z *if for any* $\boldsymbol{v} \in P$*,* $\boldsymbol{v}|_I \in Z|_I$ *implies* $\boldsymbol{v} \in Z$*.*

Sign-reflecting projection helps us project the vectors in a vector space to some of its components so that the pre-image of a certain orthant still belongs to one orthant. Moreover, we have that such a projection is one-to-one.

▶ **Lemma 4.4** ([12, Lemma A.8])**.** *Let* $P \subseteq \mathbb{Q}^d$ *be a vector space and* Z *be an orthant in* $\mathbb{Q}^d$*. Let* $I \subseteq [d]$ *be a sign-reflecting projection for* P *w.r.t.* Z*. Then every vector* $\boldsymbol{v} \in P$ *is uniquely determined by* $\boldsymbol{v}|_I$*. In other words, for any* $\boldsymbol{v}, \boldsymbol{v}' \in P$*,* $\boldsymbol{v}|_I = \boldsymbol{v}'|_I$ *implies* $\boldsymbol{v} = \boldsymbol{v}'$*.*

We mainly care about sign-reflecting projections for a plane, i.e. a 2-dimensional subspace of $\mathbb{Q}^d$. In this case a good sign-reflecting projection is given by the following lemma.

▶ **Lemma 4.5** ([12, Theorem A.9])**.** *For* $d \geq 2$*, let* $P \subseteq \mathbb{Q}^d$ *be a plane (i.e. a 2-dimensional subspace), and* Z *be an orthant in* $\mathbb{Q}^d$ *such that* $P \cap Z$ *contains two linearly independent vectors. Then there is a sign-reflecting projection* I *for* P *w.r.t.* Z *such that* $|I| = 2$*.*

Intuitively Lemma 4.5 projects a plane onto an axis plane. The preimages of those two axes will play an important role in our work, as stated in the following lemma.

▶ **Lemma 4.6.** *Let $P \subseteq \mathbb{Q}^d$ be a plane. Let $I = \{i_1, i_2\} \subseteq [d]$ be a sign-reflecting projection for P w.r.t. $\mathbb{Q}^d_{\geq 0}$. Suppose $P = \mathrm{span}\{\boldsymbol{v}_1, \boldsymbol{v}_2\}$ for linearly independent vectors $\boldsymbol{v}_1, \boldsymbol{v}_2 \in \mathbb{Z}^d$ with norm $N := \max\{\|\boldsymbol{v}_1\|, \|\boldsymbol{v}_2\|\}$. Then there exists two non-zero vectors $\boldsymbol{u}_1, \boldsymbol{u}_2 \in P \cap \mathbb{N}^d$ such that $\boldsymbol{u}_1(i_2) = \boldsymbol{u}_2(i_1) = 0$ and $\boldsymbol{u}_1(i_1) = \boldsymbol{u}_2(i_2) > 0$ and that $\|\boldsymbol{u}_1\|, \|\boldsymbol{u}_2\| \leq 2N^2$.*

We call vectors $\boldsymbol{u}_1, \boldsymbol{u}_2$ given by this lemma the *canonical horizontal / vertical vector* derived from $\boldsymbol{v}_1, \boldsymbol{v}_2$ for P with respect to $\mathbb{Q}^d_{\geq 0}$. We remark that this notion can be generalized to orthants other than $\mathbb{Q}^d_{\geq 0}$. As any vector in P is uniquely represented as a linear combination of $\boldsymbol{u}_1$ and $\boldsymbol{u}_2$, we can obtain a bound for components of vectors in P in terms of their projections.

▶ **Lemma 4.7.** *Let $P \subseteq \mathbb{Q}^d$ be a plane. Let $I = \{i_1, i_2\} \subseteq [d]$ be a sign-reflecting projection for P w.r.t. $\mathbb{Q}^d_{\geq 0}$. Suppose $P = \mathrm{span}\{\boldsymbol{v}_1, \boldsymbol{v}_2\}$ for linearly independent vectors $\boldsymbol{v}_1, \boldsymbol{v}_2 \in \mathbb{Z}^d$ with norm $N := \max\{\|\boldsymbol{v}_1\|, \|\boldsymbol{v}_2\|\}$. For every $\boldsymbol{w} \in P \cap \mathbb{N}^d$ and every $i \in \mathrm{supp}(P)$, we have*

$$\frac{\min\{\boldsymbol{w}(i_1), \boldsymbol{w}(i_2)\}}{2N^2} \leq \boldsymbol{w}(i) \leq 2N^2 \cdot (\boldsymbol{w}(i_1) + \boldsymbol{w}(i_2)). \tag{5}$$

4.2 Support Projection

Lemma 4.7 gives bounds on the components in the support of the plane. So we would like the cycle space of a VASS G to have full support, i.e. $\mathrm{supp}(\mathrm{Cyc}(G)) = [d]$. In this section we develop a technique called *support projection* to transform an arbitrary geometrically 2-dimensional VASS to one with such good property, without increasing its traversal number and the characteristic.

Let $G = (Q, T)$ be a geometrically 2-dimensional d-VASS. Let $S = \mathrm{supp}(\mathrm{Cyc}(G))$ and $\overline{S} = [d] \setminus S$. For vectors $\boldsymbol{v} \in \mathbb{Z}^S$ and $\overline{\boldsymbol{v}} \in \mathbb{Z}^{\overline{S}}$, we define their composition $\boldsymbol{v} \circ_S \overline{\boldsymbol{v}} \in \mathbb{Z}^d$ naturally by $(\boldsymbol{v} \circ_S \overline{\boldsymbol{v}})(i) = \boldsymbol{v}(i)$ if $i \in S$ and $(\boldsymbol{v} \circ_S \overline{\boldsymbol{v}})(i) = \overline{\boldsymbol{v}}(i)$ if $i \in \overline{S}$. Since S is always clear from the context, we will simply write $\boldsymbol{v} \circ \overline{\boldsymbol{v}}$ for this composition.

The *support projection* of G is the $|S|$-dimensional VASS $G^S = (Q^S, T^S)$ where

$$Q^S := \{(q, \boldsymbol{v}) \in Q \times \mathbb{N}^{\overline{S}} : \|\boldsymbol{v}\| \leq 2\chi(G)\}, \tag{6}$$

$$T^S := \{((p, \boldsymbol{u}), \boldsymbol{a}|_S, (q, \boldsymbol{v})) \in Q^S \times \mathbb{N}^S \times Q^S : (p, \boldsymbol{a}, q) \in T, \boldsymbol{u} + \boldsymbol{a}|_{\overline{S}} = \boldsymbol{v}\}. \tag{7}$$

A state of the form $(q, \boldsymbol{v})$ in G^S is denoted $q^{\boldsymbol{v}}$ for conciseness.

There is a huge expansion in the size of G^S, as $|Q^S| = |Q| \cdot (2\chi(G))^{|\overline{S}|}$. On the other hand, we can show that support projection does not increase traversal number and characteristic, and the projected VASS has full support as we expected.

▶ **Proposition 4.8.** $\varsigma(G^S) \leq \varsigma(G)$*,* $\chi(G^S) \leq \chi(G)$*, and* $\mathrm{supp}(\mathrm{Cyc}(G^S)) = S$*.*

Proof. Let's denote by $\varsigma(\pi)$ the number of distinct states in the path π. Consider any path π^S in G^S of the form $\pi^S = p_0^{\boldsymbol{v}_0} \xrightarrow{t_1^S} p_1^{\boldsymbol{v}_1} \xrightarrow{t_2^S} \cdots \xrightarrow{t_n^S} p_n^{\boldsymbol{v}_n}$ where $t_i^S \in T^S$. We define a corresponding path π in G as $\pi := p_0 \xrightarrow{t_1} p_1 \xrightarrow{t_2} \cdots \xrightarrow{t_n} p_n$ where $t_i = (p_{i-1}, \boldsymbol{a}_i, p_i)$ and $\boldsymbol{a}_i := \Delta(t_i^S) \circ (\boldsymbol{v}_i - \boldsymbol{v}_{i-1})$. Verify that t_i is indeed a transition in T by the definition of T^S. We claim that for any $i, j \in [0, n]$, $p_i = p_j$ implies $\boldsymbol{v}_i = \boldsymbol{v}_j$, then it follows that $\varsigma(\pi^S) = \varsigma(\pi)$. Indeed, suppose $p_i = p_j$, then the sub path in π from p_i to p_j is a cycle with effect $\Delta(t_{i+1}t_{i+2}\ldots t_j) \in \mathrm{Cyc}(G)$. So $\boldsymbol{v}_j = \boldsymbol{v}_i + \Delta(t_{i+1}t_{i+2}\ldots t_j)|_{\overline{S}} = \boldsymbol{v}_i$, which proves the claim. As the choice of π^S is arbitrary, for any path in G^S, there exists a path in G visits the same number of distinct states. This proves $\varsigma(G) \geq \varsigma(G^S)$.

Since it is clear that $\|T^S\| \leq \|T\|$, we immediately have $\chi(G^S) = \varsigma(G^S) \cdot \|T^S\| \leq \varsigma(G) \cdot \|T\| = \chi(G)$.

Finally, we show $\text{supp}(\text{Cyc}(G^S)) = S$. Observe that it suffices to prove $\text{Cyc}(G)|_S \subseteq \text{Cyc}(G^S)$. By Lemma 3.2, Cyc($G$) is spanned by effects of all simple cycles in G. So consider any cycle θ in G of the form $\theta := p_0 \xrightarrow{t_1} p_1 \xrightarrow{t_2} \cdots \xrightarrow{t_n} p_n = p_0$ where $t_1, \ldots, t_n \in T$. We define vectors $\boldsymbol{v}_0, \ldots, \boldsymbol{v}_n \in \mathbb{N}^{\overline{S}}$ by

$$\boldsymbol{v}_0 := \chi(G) \cdot \mathbf{1}, \qquad \boldsymbol{v}_{i+1} := \boldsymbol{v}_i + \Delta(t_{i+1})|_{\overline{S}}. \tag{8}$$

Since θ is simple, we have $n \leq \varsigma(G)$. It follows that $\mathbf{0} \leq \boldsymbol{v}_i \leq 2\chi(G) \cdot \mathbf{1}$ for all $i \in [0, n]$. So $p_i^{\boldsymbol{v}_i}$ is a state in G^S. Also note that $\boldsymbol{v}_n = \boldsymbol{v}_0 + \Delta(\theta)|_{\overline{S}} = \boldsymbol{v}_0$. We can define a corresponding cycle $\theta^S := p_0^{\boldsymbol{v}_0} \xrightarrow{t_1^S} p_1^{\boldsymbol{v}_1} \xrightarrow{t_2^S} \cdots \xrightarrow{t_n^S} p_n^{\boldsymbol{v}_n}$, where $t_i = (p_{i-1}^{\boldsymbol{v}_{i-1}}, \Delta(t_i)|_S, p_i^{\boldsymbol{v}_i})$. Verify that each t_i is a transition in T^S by definition. So θ^S is a cycle in G^S. In particular, we have $\Delta(\theta)|_S = \Delta(\theta^S) \in \text{Cyc}(G^S)$. As the choice of θ is arbitrary, we conclude that

$$\text{Cyc}(G)|_S = (\text{span}\{\Delta(\theta) : \theta \text{ is a simple cycle in } G\})|_S \subseteq \text{Cyc}(G^S), \tag{9}$$

which is the desired result. ◀

Using the support projection we can change our focus on VASSes with full support. We will prove the following lemma in the remaining of this section.

▶ **Lemma 4.9.** *For any* $\mathbf{0}$*-run* τ *in a geometrically 2-dimensional* d*-VASS* G *with the additional property that* $\text{supp}(\text{Cyc}(G)) = [d]$*, there is a* $\mathbf{0}$*-run* ρ *in* G *with the same source and target states as* τ *and* $|\rho| \leq \chi(G)^{O(\varsigma(G) \cdot d^4)}$*.*

Once Lemma 4.9 is established, Theorem 4.1 follows by plugging in the support projection of G. The rest is devoted to Lemma 4.9, so we can always assume that the VASS G has full support, i.e., $\text{supp}(\text{Cyc}(G)) = [d]$.

4.3 Degenerate VASS and Thin Runs

Depending on whether the cycle space of a VASS can be sign-reflectively projected onto an axes-plane with respect to the non-negative orthant $\mathbb{Q}_{\geq 0}^d$, we classify geometrically 2-dimensional VASSes into the following two classes.

▶ **Definition 4.10.** *A geometrically 2-dimensional VASS* G *is* proper *if* $\text{Cyc}(G) \cap \mathbb{Q}_{\geq 0}^d$ *contains two linearly independent vectors; it is* degenerate *otherwise.*

In this subsection we focus on degenerate VASSes. We show that every run from $\mathbf{0}$ in a degenerate VASS is *thin* in the sense of the following definitions. An illustration of thin runs can be found in Figure 1a.

▶ **Definition 4.11.** *Let* $\boldsymbol{v} \in \mathbb{N}^d$ *and* $W \in \mathbb{N}$*. The* beam $\mathcal{B}_{\boldsymbol{v},W}$ *is defined by*

$$\mathcal{B}_{\boldsymbol{v},W} := \left\{\boldsymbol{u} \in \mathbb{N}^d : \exists \alpha \in \mathbb{Q}_{\geq 0}, \|\boldsymbol{u} - \alpha \boldsymbol{v}\| \leq W\right\}. \tag{10}$$

The beam $\mathcal{B}_{\boldsymbol{v},W}$ *is said to be an* A-beam *where* $A \in \mathbb{N}$ *if* $\|\boldsymbol{v}\| \leq A$ *and* $W \leq A$*.*

▶ **Definition 4.12.** *Let* G *be a* d*-VASS. A run* π *in* G *is said to be* A-thin *if for every configuration* $p(\boldsymbol{u})$ *occurring in* π*, the vector* $\boldsymbol{u}$ *belongs to some* A*-beam.*

Indeed, we can relax the definition of beams by letting the direction $\boldsymbol{v}$ range over all integer vectors in $\mathbb{Z}^d$. Let $\boldsymbol{v} \in \mathbb{Z}^d$ and $W \in \mathbb{N}$. The *generalized beam* $\mathcal{B}^{\mathbb{Z}}_{\boldsymbol{v},W}$ is defined by

$$\mathcal{B}^{\mathbb{Z}}_{\boldsymbol{v},W} := \left\{\boldsymbol{u} \in \mathbb{N}^d : \exists \alpha \in \mathbb{Q}, \|\boldsymbol{u} - \alpha \boldsymbol{v}\| \leq W\right\}. \tag{11}$$

▶ **Lemma 4.13.** *For any $\boldsymbol{v} \in \mathbb{Z}^d$ and $W \in \mathbb{N}$, there exist $\boldsymbol{v}^+ \in \mathbb{N}^d$ and $\boldsymbol{v}^- \in \mathbb{N}^d$ such that $\mathcal{B}^{\mathbb{Z}}_{\boldsymbol{v},W} \subseteq \mathcal{B}_{\boldsymbol{v}^+,W} \cup \mathcal{B}_{\boldsymbol{v}^-,W}$ and that $\|\boldsymbol{v}^+\|, \|\boldsymbol{v}^-\| \leq \|\boldsymbol{v}\|$.*

With Lemma 4.13, Definition 4.12 is equivalent to stating that each configuration is located in some generalized A-beam. This makes it easier to argue if a run is thin, as demonstrated in the following lemma. We remark that the following result does not depend on whether the VASS has full support.

▶ **Lemma 4.14.** *Let G be a geometrically 2-dimensional VASS that is degenerate. Then every $\mathbf{0}$-run in G is $\chi(G)^{O(d)}$-thin.*

Proof sketch. A degenerate VASS G falls into one of the following 3 cases: (i) $\mathrm{gdim}(G) < 2$; (ii) $\mathrm{gdim}(G) = 2$ and $\mathrm{Cyc}(G) \cap \mathbb{Q}^d_{\geq 0} = \{\mathbf{0}\}$, or (iii) $\mathrm{gdim}(G) = 2$ and $\mathrm{Cyc}(G) \cap \mathbb{Q}^d_{\geq 0} = \mathbb{Q}_{\geq 0} \cdot \boldsymbol{u}$ for some $\boldsymbol{u} \in \mathbb{N}^d \setminus \{\mathbf{0}\}$. We consider only the first case here. For the other cases we refer the readers to the full version of this paper.

Suppose $\mathrm{gdim}(G) < 2$, then $\mathrm{Cyc}(G) = \mathrm{span}\{\boldsymbol{c}\}$ where $\boldsymbol{c}$ is the effect of a (possibly empty) simple cycle in G. So $\|\boldsymbol{c}\| \leq \chi(G)$. By Lemma 3.4 every configuration $q(\boldsymbol{v})$ reachable from $p(\mathbf{0})$ satisfy $\boldsymbol{v} = \alpha\boldsymbol{c} + \boldsymbol{z}$ for some $\alpha \in \mathbb{Q}$ and $\|\boldsymbol{z}\| \leq \chi(G)$. This shows that $\boldsymbol{v} \in \mathcal{B}^{\mathbb{Z}}_{\boldsymbol{c},\chi(G)}$. As the choice of $q(\boldsymbol{v})$ is arbitrary, we deduce that every $\mathbf{0}$-run in G is confined in the generalized beam $\mathcal{B}^{\mathbb{Z}}_{\boldsymbol{c},\chi(G)}$, thus is $\chi(G)$-thin. ◀

4.4 Proper VASS and the Thin-Thick Classification

In this subsection we fix a geometrically 2-dimensional d-VASS $G = (Q, T)$ that is proper, and assume that $\mathrm{supp}(\mathrm{Cyc}(G)) = [d]$. So by Lemma 4.5, there exists $i_1 \neq i_2 \in [d]$ such that $I := \{i_1, i_2\}$ is a sign-reflecting projection of $\mathrm{Cyc}(G)$ with respect to $\mathbb{Q}^d_{\geq 0}$. Moreover, let $\boldsymbol{u}_1, \boldsymbol{u}_2 \in \mathbb{N}^d$ be the canonical horizontal and vertical vectors given by Lemma 4.6. We have $\boldsymbol{u}_1(i_2) = 0$ and $\boldsymbol{u}_2(i_1) = 0$. Observe that we can assume $\|\boldsymbol{u}_1\|, \|\boldsymbol{u}_2\| \leq 2\chi(G)^2$, as they can be derived from effects of simple cycles in G.

4.4.1 Thick Runs

We will show that every $\mathbf{0}$-run in G can be classified into thin runs and thick runs. Here we give the definition of thick runs.

Sequential cones

Recall that the *cone* generated by vectors $\boldsymbol{v}_1, \dots, \boldsymbol{v}_k \in \mathbb{Z}^d$ is the set $\mathrm{Cone}\{\boldsymbol{v}_1, \dots, \boldsymbol{v}_k\} := \{\sum_{j=1}^k a_j \boldsymbol{v}_j : a_1, \dots, a_k \in \mathbb{Q}_{\geq 0}\}$. The definition of cones is enhanced in [6] where every prefix sum is also required to be non-negative:

▶ **Definition 4.15** (sequential cones). *Let $\boldsymbol{v}_1, \dots, \boldsymbol{v}_{\boldsymbol{k}} \in \mathbb{Z}^d$ be a sequence of vectors, the* sequential cone *generated by these vectors is the following set:*

$$\mathrm{SeqCone}(\boldsymbol{v}_1, \dots, \boldsymbol{v}_{\boldsymbol{k}}) := \left\{\sum_{j=1}^k a_j \boldsymbol{v}_j : a_1, \dots, a_k \in \mathbb{Q}_{\geq 0}, \forall i. \sum_{j=1}^i a_j \boldsymbol{v}_j \geq \mathbf{0}\right\}. \tag{12}$$

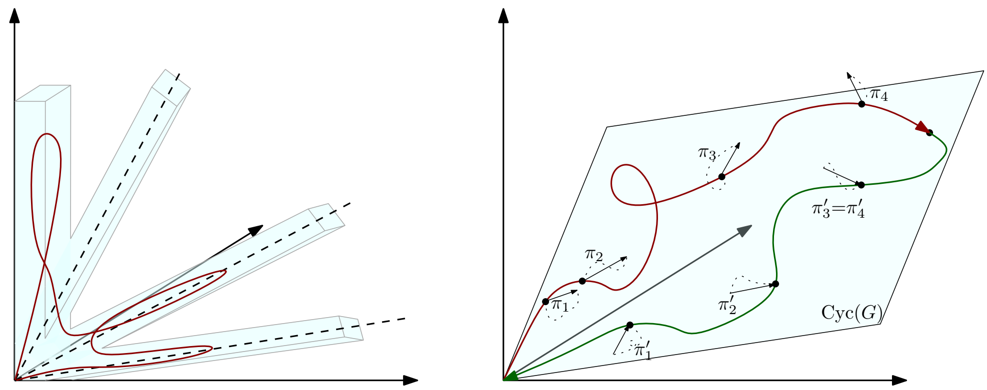

(a) Thin runs. **(b)** Thick runs.

Figure 1 Illustration of thin runs and thick runs.

In dimension 2 it was shown that a sequential cone is nothing but a cone generated by 2 vectors [6]. We generalize this result to sequential cones generated by vectors from the cycle space of a proper geometrically 2-dimensional VASS.

▶ **Lemma 4.16.** *Let* $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k \in \mathbb{Z}^d$ *be vectors in* $\mathrm{Cyc}(G)$*. Then* $\mathrm{SeqCone}(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k) = \mathrm{Cone}\{\boldsymbol{x}, \boldsymbol{y}\}$ *for some non-negative vectors* $\boldsymbol{x}, \boldsymbol{y}$ *where each of them is either* $\boldsymbol{v}_j$ *for some* $j \in [k]$*, or is the canonical horizontal / vertical vector* $\boldsymbol{u}_1$ *or* $\boldsymbol{u}_2$*.*

Indeed, this lemma reduces to [6, Lemma 2] easily by projecting the sequential cone onto coordinates in I.

Sequentially enabled cycles and thick runs

A path π is *enabled* at configuration c if there exists a configuration c' such that $c \xrightarrow{\pi} c'$ is a legal run. Let $S \subseteq [d]$, we say π is *S-enabled* at $c = p(\boldsymbol{u})$ if there exists a vector $\boldsymbol{z} \in \mathbb{N}^d$ with $\mathrm{supp}(\boldsymbol{z}) \subseteq [d] \setminus S$ such that π is enabled at $p(\boldsymbol{u} + \boldsymbol{z})$. In other words, π is S-enabled if it is enabled when ignoring the coordinates outside S.

▶ **Definition 4.17.** *Let* $A \in \mathbb{N}$*, and let* $\pi_1, \pi_2, \pi_3, \pi_4$ *be four cycles in* G*, we say these cycles are* A-sequentially enabled *in a run* ρ *in* G *if their lengths are at most* A*, and* ρ *can be factored into five parts* $\rho = \rho_1\rho_2\rho_3\rho_4\rho_5$ *such that*

- $\Delta(\pi_1)|_I$ *is semi-positive, and* π_1 *is enabled at* $\mathrm{trg}(\rho_1)$*. Moreover, both coordinates in* I *are bounded by* A *along* ρ_1*.*
- *If* $\Delta(\pi_1)|_I$ *is positive, then* π_2 *is* $\emptyset$*-enabled at* $\mathrm{trg}(\rho_2)$*. Otherwise,* π_2 *is* S*-enabled at* $\mathrm{trg}(\rho_2)$ *for* $S = [d] \setminus \mathrm{supp}(\Delta(\pi_1))$*, and, if* $\Delta(\pi_1)(i_b) = 0$*, then the* i_b*-th coordinate is bounded by* A *along* ρ_2 *where* $b = 1, 2$*.*
- $\mathrm{SeqCone}(\Delta(\pi_1), \Delta(\pi_2))$ *contains some positive vector. (We remark that this is possible only if* G *has full support.)*
- π_3, π_4 *are* $\emptyset$*-enabled at* $\mathrm{trg}(\rho_3), \mathrm{trg}(\rho_4)$ *respectively.*

▶ **Definition 4.18.** *Let* $A \in \mathbb{N}$*. A* $\mathbf{0}$*-run* τ *in* G *is* A-thick *if* τ *factors into* $\tau = \rho\rho'$ *such that*

- *some cycles* $\pi_1, \pi_2, \pi_3, \pi_4$ *in* G *are* A*-sequentially enabled in* ρ*,*
- *some cycles* $\pi'_1, \pi'_2, \pi'_3, \pi'_4$ *in* G^{rev} *are* A*-sequentially enabled in* $rev(\rho)$*,*
- $\mathrm{SeqCone}(\Delta(\pi_1), \Delta(\pi_2), \Delta(\pi_3), \Delta(\pi_4)) \cap \mathrm{SeqCone}(\Delta(\pi'_1), \Delta(\pi'_2), \Delta(\pi'_3), \Delta(\pi'_4))$ *is not trivial (i.e. it contains two linearly independent vectors).*

An illustration of thick runs is given in Figure 1b.

4.4.2 Thin-Thick Classification

In spirit of [6], the following classification lemma is of great significance.

▶ **Lemma 4.19.** *Let G be a proper geometrically 2-dimensional d-VASS with* $\mathrm{supp}(\mathrm{Cyc}(G)) = [d]$*, then there exists a number $A \leq \chi(G)^{O(\varsigma(G))}$ such that every $\mathbf{0}$-run in G is A-thick if it is not A-thin.*

An important technical lemma in [6] is the "non-negative cycle lemma" [6, Lemma 3], which states that a run in 2-VASS from $\mathbf{0}$ visiting a high configuration must contain a configuration enabling a semi-positive cycle. Here we need a similar lemma for geometrically 2-dimensional VASS.

▶ **Lemma 4.20.** *There exists a polynomial P such that every run ρ in G from $\mathbf{0}$ to $\boldsymbol{v}$ with $\|\boldsymbol{v}\| \geq P(\chi(G))^{\varsigma(G)}$ contains a configuration enabling a cycle θ of length at most $P(\chi(G))$ such that $\Delta(\theta)|_I$ is semi-positive.*

We also need some simple geometric facts.

▶ **Lemma 4.21.** *Let $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{Q}^d$. Let $X \subseteq \mathrm{span}\{\boldsymbol{u}, \boldsymbol{v}\}$ be a convex set such that $X \cap \mathbb{Q}_{\geq 0} \cdot \boldsymbol{u} = \emptyset$ and $X \cap \mathbb{Q}_{\geq 0} \cdot \boldsymbol{v} = \emptyset$. Then either $X \cap \mathrm{Cone}\{\boldsymbol{u}, \boldsymbol{v}\} = \emptyset$ or $X \subseteq \mathrm{Cone}\{\boldsymbol{u}, \boldsymbol{v}\}$.*

Given a 2-vector $\boldsymbol{v} \in \mathbb{Q}^2$, we define its *right rotation* $\boldsymbol{v}^R := (\boldsymbol{v}(2), -\boldsymbol{v}(1))$. For another vector $\boldsymbol{u} \in \mathbb{Q}^2$, we write $\boldsymbol{v} \curvearrowright \boldsymbol{u}$ if $\langle \boldsymbol{u}, \boldsymbol{v}^R \rangle > 0$, and write $\boldsymbol{v} \underline{\curvearrowright} \boldsymbol{u}$ if $\langle \boldsymbol{u}, \boldsymbol{v}^R \rangle \geq 0$. We generalize this notation to the 2-dimensional subspace $\mathrm{Cyc}(G)$: given two vectors $\boldsymbol{u}, \boldsymbol{v} \in \mathrm{Cyc}(G)$, we write $\boldsymbol{v} \curvearrowright \boldsymbol{u}$ if $\boldsymbol{v}|_I \curvearrowright \boldsymbol{u}|_I$, and $\boldsymbol{v} \underline{\curvearrowright} \boldsymbol{u}$ if $\boldsymbol{v}|_I \underline{\curvearrowright} \boldsymbol{u}|_I$.

▶ **Proposition 4.22.** *Let $\boldsymbol{u}, \boldsymbol{v} \in \mathrm{Cyc}(G)$ be such that $\boldsymbol{u} \curvearrowright \boldsymbol{v}$. Then for any $\boldsymbol{w} \in \mathrm{Cyc}(G)$, $\boldsymbol{w} \in \mathrm{Cone}\{\boldsymbol{u}, \boldsymbol{v}\}$ if and only if $\boldsymbol{u} \underline{\curvearrowright} \boldsymbol{w} \underline{\curvearrowright} \boldsymbol{v}$.*

The threshold A

Let P be the polynomial in Lemma 4.20. We define p to be the polynomial

$$p(x) = 4x^4 \cdot (2x^3 \cdot (P(x) + (x+1)^2 + 5) + 2x) + x. \tag{13}$$

Define $A := p(\chi(G))^{\varsigma(G)} \leq \chi(G)^{O(\varsigma(G))}$. Note that in particular we have

$$A \geq p(\chi(G))^{\varsigma(G)} \geq 4\chi(G)^4 \cdot (B + \chi(G)) + \chi(G) \tag{14}$$

where $B := 2\chi(G)^3 \cdot (P(\chi(G))^{\varsigma(G)} + (\chi(G)+1)^2 + 5) + \chi(G)$.

If a $\mathbf{0}$-run is not A-thin, then we can find a configuration that lies out of all A-beams. The property that G has full support helps to further show that each component of this configuration is high.

▶ **Lemma 4.23.** *Let ρ be a $\mathbf{0}$-run in G that is not A-thin. Then ρ contains a configuration $s(\boldsymbol{w})$ where $\boldsymbol{w}$ lies outside all A-beams, and such that $\boldsymbol{w}(i) \geq B$ for all $i \in [d]$.*

Main Lemma

▶ **Lemma 4.24.** *Let τ be a $\mathbf{0}$-run in G that is not A-thin. Let $s(\boldsymbol{w})$ be the configuration on τ given by Lemma 4.23. Then τ can be factored into two parts $\tau = \rho\rho'$ where* $\mathrm{trg}(\rho) = s(\boldsymbol{w}) = \mathrm{src}(\rho')$ *such that*

- *There are 4 cycles* $\pi_1, \pi_2, \pi_3, \pi_4$ *in* G *that are* B*-sequentially enabled in* ρ*, such that the set* $\mathrm{SeqCone}(\Delta(\pi_1), \Delta(\pi_2), \Delta(\pi_3), \Delta(\pi_4))$ *contains a vector* $\boldsymbol{x}$ *with* $\|\boldsymbol{x} - \boldsymbol{w}\| \leq \chi(G)$*.*
- *There are 4 cycles* $\pi_1', \pi_2', \pi_3', \pi_4'$ *in* G^{rev} *that are* B*-sequentially enabled in* $rev(\rho')$*, such that* $\mathrm{SeqCone}(\Delta(\pi_1'), \Delta(\pi_2'), \Delta(\pi_3'), \Delta(\pi_4'))$ *contains a vector* $\boldsymbol{x}'$ *with* $\|\boldsymbol{x}' - \boldsymbol{w}\| \leq \chi(G)$*.*

We show that Lemma 4.19 follows immediately from Lemma 4.24.

Proof of Lemma 4.19. If the run τ is not A-thin, then Lemma 4.24 applies. Regarding the definition of thick runs, we are only left to show $\mathrm{SeqCone}(\Delta(\pi_1), \Delta(\pi_2), \Delta(\pi_3), \Delta(\pi_4)) \cap \mathrm{SeqCone}(\Delta(\pi_1'), \Delta(\pi_2'), \Delta(\pi_3'), \Delta(\pi_4'))$ contains two linearly independent vectors. Define $U := \{\boldsymbol{u} \in \mathrm{Cyc}(G) : \|\boldsymbol{u} - \boldsymbol{w}\| \leq 2 \cdot \chi(G)\}$. By Lemma 4.16, $\mathrm{SeqCone}(\Delta(\pi_1), \Delta(\pi_2), \Delta(\pi_3), \Delta(\pi_4))$ is equal to some $\mathrm{Cone}\{\boldsymbol{v}_1, \boldsymbol{v}_2\}$ where $\|\boldsymbol{v}_1\|, \|\boldsymbol{v}_2\| \leq B \cdot \chi(G) \leq A$. As $\boldsymbol{w}$ lies out of all A-beams, we must have $U \cap \mathbb{Q}_{\geq 0}\boldsymbol{v}_1 = \emptyset$ and $U \cap \mathbb{Q}_{\geq 0}\boldsymbol{v}_2 = \emptyset$. On the other hand, Lemma 4.24 guarantees that $\mathrm{SeqCone}(\Delta(\pi_1), \Delta(\pi_2), \Delta(\pi_3), \Delta(\pi_4)) \cap U \neq \emptyset$. So by Lemma 4.21, we have $U \subseteq \mathrm{SeqCone}(\Delta(\pi_1), \Delta(\pi_2), \Delta(\pi_3), \Delta(\pi_4))$. A similar argument shows also that $U \subseteq \mathrm{SeqCone}(\Delta(\pi_1'), \Delta(\pi_2'), \Delta(\pi_3'), \Delta(\pi_4'))$. Finally, one easily verifies that U contains two linearly independent vectors. ◀

Proof sketch of Lemma 4.24. By symmetry we only need to prove the first item. Since ρ reaches the configuration $s(\boldsymbol{w})$ which is high enough, Lemma 4.20 shows that ρ contains a configuration c_1 enabling a semi-positive cycle π_1. For the second cycle, if π_1 is positive, then we simply let $\pi_2 := \pi_1$. Otherwise, $\Delta(\pi_1)$ is parallel to one of the canonical vectors $\boldsymbol{u}_1$ or $\boldsymbol{u}_2$. Say it is the latter case, so $\Delta(\pi_1)(i_1) = 0$. We need to find a cycle π_2 with $\Delta(\pi_2)(i_1) > 0$. First observe that Lemma 4.20 allows us to assume $\|c_1\| \leq P(\chi(G))^{\varsigma(G)} + \chi(G)$. As $\boldsymbol{w}(i_1) \geq B > c_i(i_1) + \chi(G)$, the run from c_1 to $s(\boldsymbol{w})$ must contain a cycle with positive effect in i_1-th coordinate. We would like to choose this cycle to be π_2, but its length may be unbounded. So we remove from it all sub-cycles whose effect is in parallel with $\Delta(\pi_1)$. This makes π_2 merely $([d] \setminus \mathrm{supp}(\Delta(\pi_1)))$-enabled, but that's enough. Finally, we exhibit cycles π_3, π_4 so that the sequential cone contains some vector $\boldsymbol{x}$ in the neighbor of $\boldsymbol{w}$. We would like $\boldsymbol{x}$ to be exactly $\boldsymbol{w}$, but this may be impossible as $\boldsymbol{w}$ might not even belong to $\mathrm{Cyc}(G)$. So we write $\boldsymbol{w} = \boldsymbol{c} + \boldsymbol{z}$ for $\boldsymbol{c} \in \mathrm{Cyc}(G)$ and $\|\boldsymbol{z}\| \leq \chi(G)$. The goal is now to contain $\boldsymbol{x} = \boldsymbol{c}$ in the sequential cone. Consider all simple cycles enabled at some configuration on the run ρ after the source of π_2. List them in the order they are enabled, and extract a subsequence $\pi_3, \pi_4, \ldots$ such that $\Delta(\pi_2) \curvearrowright \Delta(\pi_3) \curvearrowright \Delta(\pi_4) \curvearrowright \cdots$. So the effects of them sweep the cycle space clockwise. Then we are able to argue, as in [6], that there must be some π_i, π_{i+1} in this sequence such that $\Delta(\pi_i) \curvearrowright \boldsymbol{c} \curvearrowright \Delta(\pi_{i+1})$. We are done by choosing them as π_3 and π_4. ◀

4.5 Exponential Bounds of Reachability Witnesses

Lemma 4.19, together with Lemma 4.14, shows that every run in a geometrically 2-dimensional VASS G is either A-thin or A-thick for $A = \chi(G)^{O(d \cdot \varsigma(G))}$. For both cases we are able to exhibit a reachability witness of length exponential in A. The proofs are similar to that in [6]. So here we only state the results with basic sketches of the proof.

▶ **Lemma 4.25.** *For any* $\mathbf{0}$*-run* τ *that is* A*-thin in a* d*-VASS* G *where* $A \geq 2\chi(G)$*, there exists a* $\mathbf{0}$*-run* ρ *with the same source and target states as* τ *such that* $|\rho| \leq A^{O(d^2)}$*.*

Proof sketch. Recall that A-thin runs are confined in A-beams of finitely many possibilities. These beams become eventually disjoint in the region far away from the origin. If an A-thin $\mathbf{0}$-run τ without repeated configuration is long enough, it must contain a configuration c whose vector is far from origin. In this case we can find a sub-run π around c that stays within a unique A-beam. As τ is a $\mathbf{0}$-run, π first goes "up" and then "down" in this beam. From this observation we will extract from π two cycles of opposite effects by Pigeonhole Principle. Deleting these cycles shortens the length of τ. ◀

For thick runs, recall that they only occur in proper VASSes.

▶ **Lemma 4.26.** *Let G be a proper geometrically 2-dimensional d-VASS with* $\mathrm{supp}(\mathrm{Cyc}(G)) = [d]$*. For any $\mathbf{0}$-run τ in G that is A-thick for some $A \geq \chi(G)$, there is a $\mathbf{0}$-run ρ with the same source and target states as τ such that $|\rho| \leq A^{O(d^3)}$.*

Proof sketch. Consider a thick run τ. By definition, it can be decomposed into $\tau = \pi\pi'$ such that each of π and $\mathrm{rev}(\pi')$ contains four pumping cycles. For each path π and π', we remove any cycle whose effect can be realized by those pumping cycles to shorten their lengths. After that, we add some additional iterations to the pumping cycles to make the target configurations of π and π' meet in a common point. This is possible as we required that the sequential cones generated by the pumping cycles have a non-trivial intersection. The bound on the number of iterations of cycles can be obtained using standard tools in integer programming, e.g. [4] and [16]. ◀

With these bounds we can finish the proof of our major goal Lemma 4.9.

Proof of Lemma 4.9. By Lemma 4.14 and Lemma 4.19, any $\mathbf{0}$-run τ in G is either A-thin or A-thick for some $A = \chi(G)^{O(d \cdot \varsigma(G))}$. We can apply Lemma 4.25 or Lemma 4.26 to transform τ into a run ρ with the same source and target so that $|\rho| \leq A^{O(d^3)} \leq \chi(G)^{O(d^4 \cdot \varsigma(G))}$. ◀

5 A Note on Geometrically 2-Dimensional 3-VASS

The projection techniques used in Section 4 did not provide a straightforward reduction from geometrically 2-dimensional d-VASS to 2-VASS. In this section we will show that such a reduction is indeed possible for $d = 3$. We shall also mention some issues one will face when trying to generalize this reduction for $d > 3$.

In this section we care about the unary-encoding size of VASSes. Let $G = (Q, T)$ be a d-VASS, its *unary-encoding size* is defined as $|G|_1 := |Q| + d \cdot |T| \cdot \|T\| + 1$. The reduction is stated as the following lemma, from which one immediately gets the PSPACE upper bound for reachability in geometrically 2-dimensinoal 3-VASS even under binary encoding.

▶ **Lemma 5.1.** *Given a geometrically 2-dimensinoal 3-VASS G with 2 states p, q, one can compute in time polynomial in $|G|_1$ a 2-VASS $\overline{G}$ with 2 states $\overline{p}, \overline{q}$ satisfying $|\overline{G}|_1 \leq |G|_1^{O(1)}$, such that the following statements are equivalent:*

- *there exists a run from $p(\mathbf{0})$ to $q(\mathbf{0})$ of length ℓ in G;*
- *there exists a run from $\overline{p}(\mathbf{0})$ to $\overline{q}(\mathbf{0})$ of length 3ℓ in $\overline{G}$.*

▶ **Corollary 5.2.** *Reachability in geometrically 2-dimensinoal 3-VASS is in* PSPACE.

Proof. By [2, Theorem 3.2], in a 2-VASS $\overline{G}$, length of the shortest reachability witness is bounded by $|\overline{G}|_1^{O(1)}$. Together with Lemma 5.1 this shows that length of the shortest reachability witness in a geometrically 2-dimensinoal 3-VASS G is also bounded by $|G|_1^{O(1)}$, which is exponential in the binary-encoding size $|G|$. Polynomial space is enough for enumerating every run of length at most $|G|_1^{O(1)}$. ◀

The following is devoted to Lemma 5.1. Fix a geometrically 2-dimensinoal 3-VASS $G = (Q, T)$ with 2 states p, q. We can assume that $\mathrm{gdim}(G) = 2$, by adding isolated loops to G when necessary. Then there exists a normal vector $\boldsymbol{n} \in \mathbb{Z}^3$ such that $\mathrm{Cyc}(G) = \{\boldsymbol{c} \in \mathbb{Q}^3 : \langle \boldsymbol{n}, \boldsymbol{c} \rangle = 0\}$. As $\mathrm{Cyc}(G)$ is spanned by effects of two simple cycles, whose norms are bounded by $\chi(G) < |G|_1^{O(1)}$, applying Cramer's rule we can assume $\|\boldsymbol{n}\| \le |G|_1^{O(1)}$. Let $h \in \mathbb{Z}$, we define the set $C^h := \{\boldsymbol{u} \in \mathbb{N}^3 : \langle \boldsymbol{n}, \boldsymbol{u} \rangle = h\}$ as an affine copy of $\mathrm{Cyc}(G)$. The following proposition shows that any run from $p(\boldsymbol{0})$ moves within a limited number of C^h's.

▶ **Proposition 5.3.** *There exists a number $B \le |G|_1^{O(1)}$ such that for any configuration $s(\boldsymbol{w})$ reachable from $p(\boldsymbol{0})$, $\boldsymbol{w} \in C^h$ for some h with $|h| \le B$.*

Proof. By Lemma 3.4 we have $\boldsymbol{w} = \boldsymbol{c} + \boldsymbol{z}$ for some $\boldsymbol{c} \in \mathrm{Cyc}(G)$ and $\|\boldsymbol{z}\| \le \chi(G)$. So $\boldsymbol{w} \in C^h$ for $h = \langle \boldsymbol{n}, \boldsymbol{w} \rangle$. Note that $|h| = |\langle \boldsymbol{n}, \boldsymbol{w} \rangle| = |\langle \boldsymbol{n}, \boldsymbol{c} \rangle + \langle \boldsymbol{n}, \boldsymbol{z} \rangle| = |\langle \boldsymbol{n}, \boldsymbol{z} \rangle|$. We can pick $B := 3\chi(G) \cdot \|\boldsymbol{n}\| \le |G|_1^{O(1)}$. ◀

Proof of Lemma 5.1. Let B be the number given in Proposition 5.3. We consider 2 cases according to the signs of components in $\boldsymbol{n}$.

Case 1. $n \ge 0$ (or symmetrically $n \le 0$). By Proposition 5.3, every configuration $s(\boldsymbol{w})$ reachable from $p(\boldsymbol{0})$ satisfies $|\langle \boldsymbol{n}, \boldsymbol{w} \rangle| \le B$. We claim that $\left\| \boldsymbol{w}|_{\mathrm{supp}(\boldsymbol{n})} \right\| \le B$. Otherwise, as $\boldsymbol{n}$ is an integer vector, $\boldsymbol{n}(i) \ge 1$ for all $i \in \mathrm{supp}(\boldsymbol{n})$, then $\left\| \boldsymbol{w}|_{\mathrm{supp}(\boldsymbol{n})} \right\| > B$ implies $|\langle \boldsymbol{n}, \boldsymbol{w} \rangle| > B$. Then we can encode all possible values of the coordinates in $\mathrm{supp}(\boldsymbol{n})$ using the states of the VASS and thus reduce its dimension. As $\boldsymbol{n} \ne \boldsymbol{0}$, $\mathrm{supp}(\boldsymbol{n})$ contains at least one coordinate, so the resulting VASS is at most 2-dimensional. And the size is blown up by a factor of at most $(B+1)^2 \le |G|_1^{O(1)}$. Note that to validate the factor 3 in the amplification of lengths of runs as in the statements of Lemma 5.1, we also need to add two dummy transitions after each transition in the original VASS.

Case 2. n contains both positive and negative components. By negating $\boldsymbol{n}$ we can assume there is only one negative component in $\boldsymbol{n}$. Assume w.l.o.g. $\boldsymbol{n}(3) < 0$ and $\boldsymbol{n}(1), \boldsymbol{n}(2) \ge 0$. Let $\boldsymbol{n} =: (a, b, -c)$ for some $a, b, c \in \mathbb{N}$, then we have $C^h = \{(x, y, z) \in \mathbb{N}^3 : ax + by = cz + h\}$. We construct a 2-VASS $\overline{G}$ that projects G onto the first 2 coordinates. Let $\overline{G} = (\overline{Q}, \overline{T})$ be defined as follows. The state set $\overline{Q}$ contains $\{s^h : s \in Q, h \in \mathbb{Z}, |h| \le B\} \cup \{\overline{s^h} : s \in Q, h \in \mathbb{Z}, |h| \le B\}$. Intuitively a configuration $s^h(x, y)$ in $\overline{G}$ represents the configuration $s(x, y, z)$ in G such that $(x, y, z) \in C^h$, and the barred version marks an "unchecked" state, where $z = (ax + by - h)/c$ could be negative. So for each transition $t = (r, \boldsymbol{a}, s)$ of G, we add the transitions $t^h = (r^h, \boldsymbol{a}|_{\{1,2\}}, \overline{s^{h'}})$ to $\overline{G}$ for all $|h|, |h'| \le B$ and $h' = h + \langle \boldsymbol{n}, \boldsymbol{a} \rangle$. Indeed, if $r(\boldsymbol{u}) \xrightarrow{t} s(\boldsymbol{v})$ in G and $\boldsymbol{u} \in C^h$, $\boldsymbol{v} \in C^{h'}$, then $h' = \langle \boldsymbol{n}, \boldsymbol{v} \rangle = \langle \boldsymbol{n}, \boldsymbol{u} + \Delta(t) \rangle = h + \langle \boldsymbol{n}, \Delta(t) \rangle$. This justifies the correctness of the transitions we added to $\overline{G}$.

Next, we add to $\overline{G}$ transitions that checks, at each state $\overline{s^h}$, if the configuration $\overline{s^h}(x, y)$ represents a legal configuration $s(x, y, z)$ in G, where $(x, y, z) \in C^h$, so $cz = ax + by - h$. As $c > 0$, this is equivalent to checking if $ax + by - h \ge 0$. Recall that $a, b \ge 0$, one can verify that for $(x, y) \ge \boldsymbol{0}$, we have $ax + by - h \ge 0$ if and only if $(x, y) \ge \boldsymbol{m}$ for some $\boldsymbol{m} \in \mathcal{M}^h$, where $\mathcal{M}^h$ is the set of minimal solutions given by

$$\mathcal{M}^h := \begin{cases} \{(0,0)\} & \text{if } h \le 0; \\ \{(0, \lceil h/b \rceil)\} & \text{if } a = 0, b > 0, \text{ and } h > 0; \\ \{(\lceil h/a \rceil, 0)\} & \text{if } a > 0, b = 0, \text{ and } h > 0; \\ \{(x, \lceil (h - ax)/b \rceil) : x = 0, 1, \ldots, \lceil h/a \rceil\} & \text{if } a > 0, b > 0, \text{ and } h > 0. \end{cases} \tag{15}$$

Now for each number h with $|h| \leq B$, we add the path $\overline{s^h} \xrightarrow{-\boldsymbol{m}} \bullet \xrightarrow{\boldsymbol{m}} s^h$ to $\overline{G}$ for each $\boldsymbol{m} \in \mathcal{M}^h$, where $\bullet$ is an arbitrary but unique state added to $\overline{G}$. One can easily verify that $\overline{s^h}(x,y) \xrightarrow{*} s^h(x,y)$ if and only if $z = (ax+by-h)/c \geq 0$. Therefore, $p(\mathbf{0}) \xrightarrow{\pi} q(\mathbf{0})$ in G for a path π of length ℓ if and only if $p^0(\mathbf{0}) \xrightarrow{\overline{\pi}} q^0(\mathbf{0})$ in $\overline{G}$ for a path $\overline{\pi}$ of length 3ℓ. Also note that $|\mathcal{M}^h| \leq B \leq |G|_1^{O(1)}$, so we have $|\overline{G}|_1 \leq |G|_1^{O(1)}$. This completes the proof. ◀

Finally, we mention some issues when trying to generalize for higher dimensions.

1. If the dimension d is not fixed as a parameter, one can easily construct an acyclic (geometrically 0-dimensinoal) VASS so that each point in the Hamming cube $\{0,1\}^d$ is reachable from $\mathbf{0}$. Then the reduced VASS must handle an exponential number of affine copies of Cyc(G), which suggests an exponential blow-up in size.
2. For $d = 3$ the reduced 2-VASS only need to check inequalities of the form $\langle \boldsymbol{a}, \boldsymbol{x} \rangle \geq b$ for $\boldsymbol{a} \geq \mathbf{0}$. For higher dimensions we no longer have the assumption $\boldsymbol{a} \geq \mathbf{0}$. And we have no idea how to make a VASS check inequalities such as $x - y \geq 3$.

6 Geometrically 1-Dimensional and 0-Dimensional VASS

We have shown that the reachability problem in geometrically 2-dimensional VASS is PSPACE-complete, which is same as the complexity of reachability in 2-VASS. The situation becomes different for lower geometric dimensions. In this section we study the reachability problems for geometrically 1-dimensional and 0-dimensional VASSes. Indeed, these results turn out to be just rephrase of existing results for 2-VASS and 1-VASS.

▶ **Theorem 6.1.** *Reachability in geometrically 1-dimensional VASS is* PSPACE*-complete.*

Proof. The upper bound is implied by Theorem 4.2. For the lower bound, we refer the readers to the reduction from bounded one-counter automata to 2-VASS [3, Lemma 20]. We remark that as the reduced VASS uses 2 counters to simulate a bounded counter, every transition have effect of the form $(z, -z)$ for some $z \in \mathbb{Z}$. Thus, the cycle space of the reduced VASS is contained in span$\{(1,-1)\}$. So this reduction indeed establishes the PSPACE-hardness of reachability in geometrically 1-dimensional VASS. ◀

▶ **Theorem 6.2.** *Reachability in geometrically 0-dimensional VASS is* NP*-complete.*

Proof. We first show the upper bound. Let $G = (Q,T)$ be a geometrically 0-dimensional VASS. Then any cycle in G has effect $\mathbf{0}$. So for any run τ in G we can safely remove all cycles from τ and obtain a run $\overline{\tau}$ whose length is bounded by $|Q|$. Now a nondeterministic algorithm can decide reachability in G by simply guessing a run of length at most $|Q|$.

For the lower bound, we recall the folklore reduction from Subset-Sum to the reachability problem in 1-VASS. Given an instance $\langle S = \{a_1, \ldots, a_n\}, s\rangle$ of Subset-Sum, where $S \subseteq \mathbb{N}$ and the goal is to find a subset of S whose sum is s, we construct a 1-VASS $G = (Q,T)$ as follows. The states are $Q := \{q_0, \ldots, q_n\}$. For each $i = 1, \ldots, n$ we add two transitions $q_{i-1} \xrightarrow{a_i} q_i$ and $q_{i-1} \xrightarrow{0} q_i$. One can easily observe that there exists a subset of S with sum s if and only if $q_0(0) \to^* q_n(s)$ in G. Note that there are no cycles in G, thus G is indeed geometrically 0-dimensional. ◀

7 Concluding Remarks

In this paper, we have studied the reachability problem in vector addition systems with states (VASS) parameterized by geometric dimension. We introduced an efficient algorithm for computing the geometric dimension of a VASS and demonstrated some simple geometrical properties of reachable sets and runs. The primary focus is on VASS with low geometric dimensions, particularly those that are geometrically 2-dimensional. By generalizing existing pumping techniques for 2-VASS, we have shown that the reachability problem in geometrically 2-dimensional VASS is PSPACE-complete. The techniques of sign-reflecting projection and support projection also provide insights into how results for d-VASS can be adapted to geometrically d-dimensional VASS.

The results for geometrically 1- and 0-dimensional VASS are both interesting and technically straightforward. By re-examining existing results, we have shown that reachability in geometrically 1-dimensional VASS is PSPACE-complete, while in geometrically 0-dimensional VASS it is NP-complete. It is worth noting that reachability is known to be NP-complete in 1-VASS and NL-complete in 0-VASS. Our findings highlight a distinction in expressiveness and computational power between geometrically d-dimensional VASS and d-VASS. We suggest that comparing these two models could be a possible direction for future research.

In this paper our main focus is on binary encoded VASSes of low geometric dimensions. The situation would become different if unary encoding is considered. It is known that in fixed-dimensional settings the problem is NL-complete for dimensions 0, 1, and 2 [2]. However, if only the geometric dimension is fixed, we note that logarithmic space is not enough to store a single configuration containing d counters (recall that the size of $G = (Q, T)$ under unary encoding is $|G|_1 := |Q| + d \cdot |T| \cdot \|T\| + 1$). The complexity of reachability in low geometric dimensions under unary encoding is left for future work.

References

1. Shaull Almagor, Arka Ghosh, Tim Leys, and Guillermo A. Pérez. The geometry of reachability in continuous vector addition systems with states. *Inf. Comput.*, 304:105298, 2025. doi:10.1016/J.IC.2025.105298.
2. 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.
3. Michael Blondin, Alain Finkel, Stefan Göller, Christoph Haase, and Pierre McKenzie. Reachability in two-dimensional vector addition systems with states is pspace-complete. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2015, Kyoto, Japan, July 6-10, 2015*, pages 32–43. IEEE Computer Society, 2015. doi:10.1109/LICS.2015.14.
4. Dmitry Chistikov and Christoph Haase. The taming of the semi-linear set. In Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi, editors, *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*, volume 55 of *LIPIcs*, pages 128:1–128:13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPICS.ICALP.2016.128.
5. 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, December 18-20, 2023, IIIT Hyderabad, Telangana, India*, volume 284 of *LIPIcs*, pages 35:1–35:22. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPICS.FSTTCS.2023.35.
6. Wojciech Czerwinski, Slawomir Lasota, Christof Löding, and Radoslaw Piórkowski. New pumping technique for 2-dimensional VASS. In Peter Rossmanith, Pinar Heggernes, and Joost-Pieter Katoen, editors, *44th International Symposium on Mathematical Foundations of Computer Science, MFCS 2019, August 26-30, 2019, Aachen, Germany*, volume 138 of *LIPIcs*, pages 62:1–62:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICS.MFCS.2019.62.

7 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.

8 Wojciech Czerwiński, Ismaël Jecker, Sławomir Lasota, and Łukasz Orlikowski. Reachability in 3-vass is elementary. *CoRR*, abs/2502.13916, 2025. doi:10.48550/arXiv.2502.13916.

9 Matthias Englert, Ranko Lazic, and Patrick Totzke. Reachability in two-dimensional unary vector addition systems with states is nl-complete. In Martin Grohe, Eric Koskinen, and Natarajan Shankar, editors, *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '16, New York, NY, USA, July 5-8, 2016*, pages 477–484. ACM, 2016. doi:10.1145/2933575.2933577.

10 Javier Esparza and Mogens Nielsen. Decidability issues for petri nets - a survey. *J. Inf. Process. Cybern.*, 30(3):143–160, 1994.

11 John Fearnley and Marcin Jurdzinski. Reachability in two-clock timed automata is pspace-complete. *Inf. Comput.*, 243:26–36, 2015. doi:10.1016/J.IC.2014.12.004.

12 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, July 8-12, 2024, Tallinn, Estonia*, volume 297 of *LIPIcs*, pages 136:1–136:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICS.ICALP.2024.136.

13 Christoph Haase, Stephan Kreutzer, Joël Ouaknine, and James Worrell. Reachability in succinct and parametric one-counter automata. In Mario Bravetti and Gianluigi Zavattaro, editors, *CONCUR 2009 - Concurrency Theory, 20th International Conference, CONCUR 2009, Bologna, Italy, September 1-4, 2009. Proceedings*, volume 5710 of *Lecture Notes in Computer Science*, pages 369–383. Springer, 2009. doi:10.1007/978-3-642-04081-8_25.

14 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.

15 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.

16 Loic Pottier. Minimal solutions of linear diophantine systems: Bounds and algorithms. In Ronald V. Book, editor, *Rewriting Techniques and Applications, 4th International Conference, RTA-91, Como, Italy, April 10-12, 1991, Proceedings*, volume 488 of *Lecture Notes in Computer Science*, pages 162–173. Springer, 1991. doi:10.1007/3-540-53904-2_94.

17 Sylvain Schmitz. Complexity hierarchies beyond elementary. *CoRR*, abs/1312.5686, 2013. arXiv:1312.5686.

18 Qizhe Yang and Yuxi Fu. Reachability in 3-vass is in tower. *CoRR*, abs/2306.05710, 2023. doi:10.48550/arXiv.2306.05710.

19 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.