



The Reachability Problem for Two-Dimensional Vector Addition Systems with States

Michael Blondin, Matthias Englert, Alain Finkel, Stefan Göller, Christoph Haase,
Ranko Lazić, Pierre McKenzie, Patrick Totzke

► To cite this version:

Michael Blondin, Matthias Englert, Alain Finkel, Stefan Göller, Christoph Haase, et al.. The Reachability Problem for Two-Dimensional Vector Addition Systems with States. Journal of the ACM (JACM), 2021, 68 (5), pp.1-43. <10.1145/3464794>. <hal-04270941>

HAL Id: hal-04270941

<https://hal.science/hal-04270941v1>

Submitted on 13 Nov 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

The Reachability Problem for Two-Dimensional Vector Addition Systems with States

MICHAEL BLONDIN, Université de Sherbrooke, Canada

MATTHIAS ENGLERT, University of Warwick, United Kingdom

ALAIN FINKEL, Laboratoire LMF, Université Paris-Saclay, CNRS, ENS Paris-Saclay, IUF, France

STEFAN GÖLLER, Universität Kassel, Germany

CHRISTOPH HAASE, University of Oxford, United Kingdom

RANKO LAZIĆ, University of Warwick, United Kingdom

PIERRE MCKENZIE, Université de Montréal, Canada

PATRICK TOTZKE, University of Liverpool, United Kingdom

We prove that the reachability problem for two-dimensional vector addition systems with states is NL-complete or PSPACE-complete, depending on whether the numbers in the input are encoded in unary or binary. As a key underlying technical result, we show that if a configuration is reachable then there exists a witnessing path whose sequence of transitions is contained in a bounded language defined by a regular expression of pseudo-polynomially bounded length. This in turn enables us to prove that the lengths of minimal reachability witnesses are pseudo-polynomially bounded.

CCS Concepts: • **Theory of computation** → **Automata over infinite objects; Complexity classes; Logic and verification**;

Additional Key Words and Phrases: vector addition systems with states, Petri nets, semi-linear sets, linear path schemes, reachability

ACM Reference Format:

Michael Blondin, Matthias Englert, Alain Finkel, Stefan Göller, Christoph Haase, Ranko Lazić, Pierre McKenzie, and Patrick Totzke. 2021. The Reachability Problem for Two-Dimensional Vector Addition Systems with States. *J. ACM* 0, 0, Article 0 (2021), 41 pages. <https://doi.org/0000001.0000001>

1 INTRODUCTION

Petri nets form a fundamental model of computation with a long history. Since their introduction by Carl Adam Petri in 1962 [37], tens of thousands of papers on Petri nets have been published. Due to their generic nature, Petri nets have found a variety of applications, ranging, for instance, from

Authors' addresses: Michael Blondin, Université de Sherbrooke, 2500 boulevard de l'Université, Sherbrooke, J1K 2R1, QC, Canada, michael.blondin@usherbrooke.ca; Matthias Englert, University of Warwick, Coventry, CV4 7AL, United Kingdom, m.englert@warwick.ac.uk; Alain Finkel, Laboratoire LMF, Université Paris-Saclay, CNRS, ENS Paris-Saclay, IUF, 4 avenue des sciences, CS 30008, Gif/Yvette, 91192, France, Alain.Finkel@ens-paris-saclay.fr; Stefan Göller, Universität Kassel, Wilhelmshöher Allee 73, Kassel, 34121, Germany, stefan.goeller@uni-kassel.de; Christoph Haase, University of Oxford, Parks Rd, Oxford, OX1 3QD, United Kingdom, christoph.haase@cs.ox.ac.uk; Ranko Lazić, University of Warwick, Coventry, CV4 7AL, United Kingdom, r.s.lazic@warwick.ac.uk; Pierre McKenzie, Université de Montréal, CP 6128 succ. Centre-ville, Montréal, H3C 3J7, QC, Canada, mckenzie@iro.umontreal.ca; Patrick Totzke, University of Liverpool, Department of Computer Science, Ashton Building, Ashton Street, Liverpool, L69 3BX, United Kingdom, totzke@liverpool.ac.uk.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

0004-5411/2021/0-ART0 \$15.00

<https://doi.org/0000001.0000001>

modeling of biological, chemical and business processes to the formal verification of concurrent programs, see e.g. [1, 2, 21, 24, 32]. For the analysis of algorithmic properties of Petri nets, in the contemporary literature they are often equivalently viewed as *vector addition systems with states* (VASS), and we will adopt this view in this article. A VASS comprises a finite-state controller with a finite number of counters ranging over the natural numbers. The number of counters is usually referred to as the dimension of the VASS, and we write d -VASS when we talk about VASS in dimension d . When taking a transition, a VASS can add or subtract an integer from a counter, provided that the resulting counter values are greater than or equal to zero; otherwise the transition is blocked. Consequently, configurations of VASS are tuples consisting of a control state and an assignment of the counters to natural numbers, and every VASS gives rise to an infinite directed graph whose vertices are the set of configurations. The central decision problem for VASS is *reachability*: given two configurations, is there a path connecting them in the infinite graph induced by the VASS?

Even clarifying the decidability status of the reachability problem for VASS required tremendous efforts, and it actually took until 1981 for it to be shown decidable. This was achieved by Mayr [34], who built upon an earlier partial proof by Sacerdote and Tenney [40]. Mayr's argument was then polished and simplified by Kosaraju [22] in 1982, and Kosaraju's argument was in turn simplified ten years later by Lambert [23]. Only in the past decade, Leroux developed, in a series of papers, a substantially different approach to the decidability of the reachability problem [25–27]. Regarding the computational complexity of the general reachability problem, the best known bounds have been obtained very recently: an ACKERMANN upper bound by Leroux and Schmitz [28], and a TOWER lower bound by Czerwiński et al. [10]. Prior to the latter work, the state of the art lower bound was Lipton's EXPSPACE [33].

Motivated both by the unsettled status of the complexity of the general problem, and by the wide interest in classes of one-counter and two-counter automata (see e.g. [4, 7, 14]), the reachability problem for VASS of small fixed dimensions has attracted considerable attention. Deciding reachability of 1-VASS assuming unary encoding of numbers is NL-complete: the lower bound is inherited from directed graph reachability [36, Theorem 16.2] and the upper bound follows from a straightforward depumping argument [42]. When numbers are encoded in binary, reachability in 1-VASS is known to be NP-complete [18]. A substantial contribution towards showing the decidability of the general reachability problem was made by Hopcroft and Pansiot in 1979, who showed that reachability in 2-VASS is decidable [19]. To this end, they developed an intricate algorithm that implicitly exploits the fact that the reachability set of a 2-VASS is semi-linear. Moreover, they could show that their method breaks down for VASS in any greater dimension, as the authors exhibited a 3-VASS with a reachability set that is not semi-linear. Yet, aspects of computational complexity were left completely unanswered in [19]. In 1986, Howell, Rosier, Huynh and Yen [20] analyzed Hopcroft and Pansiot's algorithm and showed that it runs in nondeterministic doubly-exponential time, independently of whether numbers are presented in unary or binary. They could improve this nondeterministic doubly-exponential time upper bound to a deterministic doubly-exponential one and also identify a family of 2-VASS on which Hopcroft and Pansiot's algorithm requires doubly-exponential time. In summary, since 1986 it has been the state of the art that reachability in 2-VASS is in 2-EXPTIME, and NL-hard and NP-hard, depending on whether numbers are encoded in unary or binary.

Contributions

In this article, in terms of the broad brush-strokes of complexity theory, we settle the status of the reachability problem for 2-VASS and show that it is:

- NL-complete assuming unary encoding of numbers; and
- PSPACE-complete assuming binary encoding of numbers.

The lower bounds are straightforward: NL-hardness holds already in dimension one as we remarked, and PSPACE-hardness follows as an easy consequence of a result by Fearnley and Jurdziński who showed PSPACE-completeness of reachability in bounded one-counter automata (i.e., 1-VASS with zero tests in which the counter is bounded below and above) [14]. Most of the work is therefore to establish the upper bounds, and it is organized in two stages:

- (1) Our starting point is a careful analysis of an argument developed by Leroux and Sutre in [29] for the purpose of showing that reachability relations of VASS in dimension two can be captured by the language of regular expression of star height one, i.e., speaking in the terminology of [29], 2-VASS can be *flattened*. More precisely, this means that for any 2-VASS there is a finite set S of *linear path schemes*, which are regular expressions of the form $\alpha_0 \beta_1^* \alpha_1 \cdots \beta_k^* \alpha_k$ over the set of transitions viewed as an alphabet, such that for any two reachable configurations there exists a witnessing run in the language defined by S . The paper of Leroux and Sutre reports that from any 2-VASS it is possible to construct such a regular language. This immediately implies that the reachability relation of 2-VASS is semi-linear. In dimension three, the reachability relation is no longer semi-linear and hence such linear path schemes cannot exist; the classical example by Hopcroft and Pansiot [19, proof of Lemma 2.8] gives a 3-VASS that does not possess a semi-linear reachability set. The paper [29] has not appeared as a fully refereed publication and some proof details are omitted in it. Thus, while in the first stage we mostly follow the proof strategy presented in [29], we provide a complete proof that 2-VASS can be flattened, and doing so requires the development of new arguments in order to enable a tight analysis which establishes a pseudo-polynomial bound on the lengths of the linear path schemes.¹
- (2) We next prove that, for linear path schemes obtained from 2-VASS, the lengths of runs witnessing reachability can be pseudo-polynomially bounded. The technique we have developed for showing this bound is surprisingly involved and can be seen as an extension to two dimensions of the classical depumping argument known as *hill cutting* due to Valiant and Paterson [42]. Combined with the pseudo-polynomially bounded flattening of any 2-VASS, this result enables us to conclude that 2-VASS, depending on whether transition updates are given with numbers encoded in unary or binary, have respectively polynomially or exponentially bounded reachability witnesses. The NL and PSPACE upper bounds then immediately follow from considering a simple nondeterministic algorithm that guesses and checks a reachability witness of bounded length by storing at most one configuration at a time (cf. [39, proof of Theorem 3.5]).

To the best of our knowledge, this is the first time that the complexity of an interesting restriction of the VASS reachability problem has broken ‘the size of the reachability set barrier’. Namely, it is well-known that general VASS may have reachability sets which are finite but Ackermannianly large [6], and in spite of the latest results [28], the Ackermann barrier remains. When the dimension is fixed to two and updates in the transitions are given in unary, it is not difficult to construct examples of 2-VASS with exponentially large reachability sets (by employing weak doubling a number of times proportional to the number of states—this uses integers only up to absolute value 2), but we prove that polynomially long reachability witnesses always exist.

This article has been produced by completing, integrating and streamlining the conference papers [3] and [13]. Ever since their publication, the main results as well as technical results that

¹Recall that a pseudo-polynomial bound is polynomial when the numbers in the input are given in unary, and in general is exponential when they are given in binary.

we present in this paper have found applications for a variety of problems. Czerwinski et al. [11] have found an alternative proof of our PSPACE upper bound, whereas Leroux and Sutre [30] have built on our work to show PSPACE membership even with one of the two counters being zero testable. Czerwiński and Lasota employ our PSPACE upper bound and the bounds on the flattening of 2-VASS in order to give a PSPACE algorithm for the regular separability problem for languages of one-counter automata [9]. Given two languages L and M accepted by one-counter automata, this problem is to decide whether there exists a regular language R such that $L \subseteq R$ and $R \cap M = \emptyset$. Michaliszyn, Otop and Wieczorek have used an intermediate technical result of [3] on the length of shortest paths in \mathbb{Z} -VASS (VASS whose counters are allowed to take negative values) in any fixed dimension in order to derive optimal algorithms for querying graph databases [35]. Finally, Brázdil et al. have studied the long-run average behaviour of probabilistic extensions of 1-VASS and 2-VASS [5]. They did not give complexity upper bounds for the case of probabilistic 2-VASS, the main reason being that bounds on shortest paths in unary 2-VASS were not known at the time. Moreover, their results crucially rely on the flatness property of 2-VASS for which no bounds existed at the time either. It seems conceivable that the results in this article should, with some efforts, enable obtaining (tight) upper bounds for the problems studied in [5].

Organization

After fixing the notations in Section 2 and stating the main results in Section 3, we present the two main stages of the proof in Sections 4 and 5, and finish with remarks in Section 6.

Since in terms of computational complexity our focus is on a classification with respect to classes such as NL and PSPACE, most of the constants including exponents are omitted from the stated bounds to aid readability, but there are exceptions such as many of the technical lemmas in Section 5.

2 PRELIMINARIES

General Notation

By \mathbb{Q} , \mathbb{Z} and \mathbb{N} we denote the sets of *rational*s, *integer*s and *natural*s (i.e., non-negative integers). We shortly write $\mathbb{Q}_{\geq 0}$ for the set of non-negative rationals and use analogous notations for similarly restricted sets of numbers. For $i, j \in \mathbb{Z}$, we write $[i, j] \stackrel{\text{def}}{=} \{i, i+1, \dots, j\}$ for the set of all integers between i and j inclusively, and simply $[i]$ for $[1, i]$.

For scalars $\lambda \in \mathbb{Q}$ and d -dimensional vectors $\mathbf{v}, \mathbf{w} \in \mathbb{Q}^d$, we use standard notations: $\lambda \mathbf{v}$ for scalar multiplication, $\mathbf{v} \cdot \mathbf{w}$ for dot product, $\mathbf{v} + \mathbf{w}$ for component-wise sum, $\mathbf{v} \leq \mathbf{w}$ for component-wise non-strict order, $\mathbf{v} < \mathbf{w}$ for component-wise strict order (i.e., strict in at least one component), and $\|\mathbf{v}\|$ for the infinity norm (i.e., the maximum absolute value of the components of \mathbf{v}). The component-wise sum is extended to sets V and W of d -dimensional vectors as $V + W \stackrel{\text{def}}{=} \{\mathbf{v} + \mathbf{w} : \mathbf{v} \in V, \mathbf{w} \in W\}$, and the infinity norm is extended to finite sets of vectors V as $\|V\| \stackrel{\text{def}}{=} \max_{\mathbf{v} \in V} \|\mathbf{v}\|$, and to matrices $A \in \mathbb{Q}^{m \times n}$ as $\|A\| \stackrel{\text{def}}{=} \max_{i \in [m], j \in [n]} |a_{i,j}|$.

Semi-Linear Sets

For $\mathbb{D} \subseteq \mathbb{Q}$ and $P \subseteq \mathbb{Q}^d$, the \mathbb{D} -cone generated by P is the set of all non-empty linear combinations of elements of $P = \{\mathbf{p}_1, \dots, \mathbf{p}_n\}$ with coefficients from \mathbb{D} :

$$\text{cone}_{\mathbb{D}}(P) \stackrel{\text{def}}{=} \{\lambda_1 \mathbf{p}_1 + \dots + \lambda_n \mathbf{p}_n : \lambda_1, \dots, \lambda_n \in \mathbb{D} \text{ and } n \geq 1\}.$$

For $B \subseteq \mathbb{Q}^d$ and \mathbb{D}, P as above, the *hybrid-linear set* with coefficients \mathbb{D} , basis B and periods P is:

$$L_{\mathbb{D}}(B, P) \stackrel{\text{def}}{=} B + \text{cone}_{\mathbb{D}}(P).$$




Fig. 1. An example 2-VASS.

Linear sets $L_{\mathbb{D}}(\mathbf{b}, P)$ are hybrid-linear sets with singleton bases $B = \{\mathbf{b}\}$, and semi-linear sets are finite unions of linear sets. When the set of coefficients \mathbb{D} is the set of naturals \mathbb{N} , we may omit it from these notations.

Vector Addition Systems with States

Even though this article is on 2-dimensional VASS, for completeness we introduce general d -dimensional VASS, or d -VASS for short. A d -VASS is given by a finite \mathbb{Z}^d -labelled directed graph (Q, T) , whose vertices $q \in Q$ are called *states*, and whose edges (p, z, q) are called *transitions* (here $p, q \in Q$ and $z \in \mathbb{Z}^d$). For a transition $t = (p, z, q)$, we define $\|t\| \stackrel{\text{def}}{=} \|z\|$, and $\|T\| \stackrel{\text{def}}{=} \max_{t \in T} \|t\|$.

Transitions of a d -VASS (Q, T) give rise to a number of relations between *configurations*, which consist of a state q and a vector $\mathbf{v} \in \mathbb{Z}^d$, and are written $q(\mathbf{v})$:

- for a single transition $t = (p, z, q)$, the relation $\xrightarrow{t}_{\mathbb{Z}^d}$ is between configurations of the forms $p(\mathbf{u})$ and $q(\mathbf{u} + z)$, where $\mathbf{u} \in \mathbb{Z}^d$, i.e. the transition specifies the source and target states, and its effect is to add the vector z ;
- for a word (i.e., a finite sequence) π of transitions, the relation $\xrightarrow{\pi}_{\mathbb{Z}^d}$ is obtained by composing in order the relations for the transitions in π , which yields a non-empty relation if and only if π is a path in the graph (Q, T) ;
- for a language (i.e., a set of words) L of transitions, the relation $\xrightarrow{L}_{\mathbb{Z}^d}$ is obtained as the union of the relations for the words in L ;
- we write $\xrightarrow{*}_{\mathbb{Z}^d}$ for the relation $\xrightarrow{T^*}_{\mathbb{Z}^d}$, so that $p(\mathbf{u}) \xrightarrow{*}_{\mathbb{Z}^d} q(\mathbf{v})$ if and only if there exists a path π from state p to state q and $\mathbf{v} = \mathbf{u} + \text{effect}(\pi)$, where $\text{effect}(\pi)$ is the sum of all the vectors that label the transitions in π .

We are typically interested in configurations whose vectors belong to some subset $\mathbb{D} \subseteq \mathbb{Z}^d$. To make that clear, we write \mathbb{D} as a subscript in the relational notations above, which requires that all configurations have vectors from \mathbb{D} . For example, $p(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{D}} q(\mathbf{v})$ means that the path π leads from configuration $p(\mathbf{u})$ to configuration $q(\mathbf{v})$, and that \mathbf{u} , the vector of every intermediate configuration and \mathbf{v} all belong to the set \mathbb{D} . In such cases, we say that π gives rise to (or shortly, is) a \mathbb{D} -run from \mathbf{u} to \mathbf{v} . We remark that \mathbf{u} and π determine q and all the remaining configurations in the \mathbb{D} -run, up to and including $q(\mathbf{v})$.

When the set \mathbb{D} is omitted, it is understood to be \mathbb{N}^d . That is the default restriction for d -VASS: their *runs* are \mathbb{N}^d -runs, which means that the vectors of all configurations (source, intermediate and target) have all components non-negative. We say that a path π is admissible from a vector $\mathbf{u} \in \mathbb{N}^d$ if and only if it gives rise to a run from \mathbf{u} , i.e. performing the transitions of π starting from \mathbf{u} does not cause any vector component to become negative.

For instance, in the 2-VASS depicted in Figure 1, we have both $p(0, 0) \xrightarrow{t_1 t_1 t_2 t_3}_{\mathbb{Z}^2} p(1, 1)$ and $p(0, 0) \xrightarrow{t_1 t_2 t_3 t_1}_{\mathbb{Z}^2} p(1, 1)$. However, the former expression is admissible from $(0, 0)$ but the latter is not, and so $p(0, 0) \xrightarrow{t_1 t_1 t_2 t_3}_{\mathbb{N}^2} p(1, 1)$ holds but $p(0, 0) \xrightarrow{t_1 t_2 t_3 t_1}_{\mathbb{N}^2} p(1, 1)$ does not.

Linear Path Schemes

Suppose (Q, T) is a d -VASS. The notions of path and cycle have standard meanings in the context of the directed graph (Q, T) . More precisely, a *path* is a sequence $\pi = t_1 \cdots t_n$, where each $t_i = (p_i, z_i, q_i)$ is a transition from T , that satisfies $q_i = p_{i+1}$ for every $i \in [n-1]$. Such a path π is a *cycle* if $q_n = p_1$. By a *linear path scheme* (for short, *LPS*), we mean a regular expression of the form

$$\Lambda = \alpha_0 \beta_1^* \alpha_1 \cdots \beta_k^* \alpha_k,$$

where each α_i is a path, each β_j is a cycle, and $\alpha_0 \beta_1 \alpha_1 \cdots \beta_k \alpha_k$ is a path. A consequence of those restrictions is that each word in the language Λ is also a path; they differ in the numbers of times (possibly 0) that the cycles β_j are performed.

We write $\text{cycles}(\Lambda) \stackrel{\text{def}}{=} \{\beta_1, \dots, \beta_k\}$ for the set of cycles of Λ , $|\Lambda|_* \stackrel{\text{def}}{=} k$ for their number, $|\Lambda| \stackrel{\text{def}}{=} |\alpha_0 \beta_1 \alpha_1 \cdots \beta_k \alpha_k|$ for the length of Λ which is the length of the underlying path, and $\|\Lambda\|$ for the maximum of the norms of the vectors that occur in the underlying path.

The Reachability Problem

For a d -VASS $V = (Q, T)$, we call $\xrightarrow{*}_{\mathbb{Z}^d}$ the \mathbb{Z} -reachability relation (where there is no restriction on the vectors), and $\xrightarrow{*}$ (i.e., $\xrightarrow{*}_{\mathbb{N}^d}$) the reachability relation. The reachability problem is then:

Given a d -VASS \mathcal{V} and two configurations $p(\mathbf{u})$ and $q(\mathbf{v})$, is $q(\mathbf{v})$ reachable from $p(\mathbf{u})$ (i.e., is it the case that $p(\mathbf{u}) \xrightarrow{*} q(\mathbf{v})$)?

When considering complexity, we distinguish between two cases depending on whether the integers in the transition labels of \mathcal{V} and in \mathbf{u} and \mathbf{v} are given in unary or binary.

Unary: The size of \mathcal{V} is $|\mathcal{V}|_1 \stackrel{\text{def}}{=} |Q| + d \cdot |T| \cdot \|T\| + 1$, the size of \mathbf{u} is $|\mathbf{u}|_1 \stackrel{\text{def}}{=} d \cdot \|\mathbf{u}\|$, and analogously for \mathbf{v} .

Binary: The size of \mathcal{V} is $|\mathcal{V}|_2 \stackrel{\text{def}}{=} |Q| + d \cdot |T| \cdot \lceil \log(\|T\| + 1) \rceil + 1$, the size of \mathbf{u} is $|\mathbf{u}|_2 \stackrel{\text{def}}{=} d \cdot \lceil \log(\|\mathbf{u}\| + 1) \rceil$, and analogously for \mathbf{v} .

3 OVERVIEW OF THE MAIN RESULTS

To obtain the crowning achievements of this article, namely the NL or PSPACE upper bounds on the complexity of the reachability problem for 2-dimensional VASS depending on whether the numbers in the input are encoded in unary or binary (respectively), we prove two main theorems. The first one says that 2-VASS can be flattened with small linear path schemes, and the second one that short reachability witnesses can be obtained from such a flattening. Before stating these two contributions formally, let us briefly explain what we mean by ‘flattening’ and ‘reachability witness’.

Consider the 2-VASS depicted in Figure 1. The set of all paths from state p to state q is described by the regular expression $r = t_1^* t_2 (t_3 t_1^* t_2)^*$. Such an expression can be valuable to determine whether $q(\mathbf{v})$ is reachable from $p(\mathbf{u})$ for some given $\mathbf{u}, \mathbf{v} \in \mathbb{N}^2$. However, the nested Kleene stars in r complicate this task. Fortunately, r can be *flattened* due to a simple observation: if a path allows us to reach $q(\mathbf{v})$ from $p(\mathbf{u})$, then it can be reordered into a path where all occurrences of t_1 are at the beginning. This holds because $\text{effect}(t_1) \geq 0$ and because executing non-negative transitions earlier can only help. Therefore, if we are interested in reachability from state p to state q , then it suffices to examine the runs induced by paths from the linear path scheme $\Lambda = t_1^* t_2 (t_3 t_2)^*$. In other words,

$$p(\mathbf{u}) \xrightarrow{*} q(\mathbf{v}) \quad \text{if, and only if,} \quad p(\mathbf{u}) \xrightarrow{t_1^x t_2 (t_3 t_2)^y} q(\mathbf{v}) \quad \text{for some } x, y \in \mathbb{N}. \quad (1)$$

Determining whether the right-hand side of (1) holds is not immediate, but it is a simpler task, e.g. it can be done by solving a well-chosen system of linear Diophantine inequalities. Moreover,

a careful analysis shows that if a solution (x, y) exists, then there exists a solution whose norm is polynomial in terms of $\|\mathbf{u}\|$ and $\|\mathbf{v}\|$. This implies that reachability from $p(\mathbf{u})$ to $q(\mathbf{v})$ can be *witnessed* by a short path.

In general, the structure of a VASS can be much more complicated than the one depicted in Figure 1, e.g., there can be nested cycles where no cycle is non-negative. Thus, it is not even clear whether reachability can always be captured by (small) linear path schemes. As we shall see, it is in fact always the case for 2-VASS. More precisely, our first main result will be:

THEOREM 3.1. *For every 2-VASS $\mathcal{V} = (Q, T)$, there exists a set S of LPSs such that*

- (1) $\rightarrow^* = \bigcup S$, and
- (2) *every linear path scheme in S has $O(|Q|^2)$ many cycles and length $|\mathcal{V}|_1^{O(1)}$.*

We shall then bound the number of times each cycle of a linear path scheme must be traversed to witness reachability, which will allow us to derive the second main result:

THEOREM 3.2. *For every 2-VASS \mathcal{V} , if $p(\mathbf{u}) \rightarrow^* q(\mathbf{v})$, then $p(\mathbf{u}) \xrightarrow{\pi} q(\mathbf{v})$ for a path π of length $(|\mathcal{V}|_1 + \|\mathbf{u}\| + \|\mathbf{v}\|)^{O(1)}$.*

Two remarks are in order:

- In Theorem 3.1, the set S is necessarily finite since the lengths of its members are bounded.
- Like in other statements in this article where $O(1)$ replaces explicit exponents, they are constant and do not depend on the 2-VASS under consideration.

COROLLARY 3.3. *The 2-VASS reachability problem is NL-complete under unary encoding and PSPACE-complete under binary encoding.*

PROOF. Under unary encoding, we have by Theorem 3.2 that polynomially long reachability witnesses always exist, so logarithmic space suffices for a nondeterministic algorithm that guesses and checks such a witness by storing at most one transition at a time (cf. [39, proof of Theorem 3.5]). The problem is NL-hard by an immediate reduction from reachability in directed graphs [36, Theorem 16.2].

Under binary encoding, since $2^{|\mathcal{V}|_2 + \|\mathbf{u}\|_2 + \|\mathbf{v}\|_2} \geq |\mathcal{V}|_1 + \|\mathbf{u}\| + \|\mathbf{v}\|$, we have by Theorem 3.2 that exponentially long reachability witnesses always exist, so polynomial space suffices for a nondeterministic algorithm as before. The problem is PSPACE-hard by a straightforward logarithmic-space reduction from the reachability problem for bounded one-counter automata [14]: the latter are essentially 1-VASS in which the counter x is restricted to a given range $[0, B]$, and can be simulated by 2-VASS that have a second counter that maintains the value $B - x$. \square

In the remainder, we focus on proving Theorem 3.1 in Section 4, and Theorem 3.2 in Section 5. An initial depiction of the structures of those proofs can be found in Figure 2. Generally speaking, our results are obtained in a bottom-up fashion. At the basis, we analyze very subtle yet crucial properties of semi-linear sets. We also develop bounds on linear path schemes for restricted classes and variants of 2-VASS, namely 1-VASS, 2-VASS with restrictions on the domains of the counters, and \mathbb{Z} -VASS (i.e., VASS where we use $\xrightarrow{*}_{\mathbb{Z}}$ as reachability relation). Those preliminary results are then pieced together to obtain Theorem 3.1 which is then used to prove Theorem 3.2.

4 FLATTABILITY: THEOREM 3.1

In this section, we will prove Theorem 3.1 which we recall from Section 3:

THEOREM 3.1. *For every 2-VASS $\mathcal{V} = (Q, T)$, there exists a set S of LPSs such that*




Fig. 2. Overview of the proofs of Theorems 3.1 and 3.2. An arrow from node u to node v indicates that results on u are used to prove results on v . The colored nodes hide more detailed lemmas on the respective topics and are unfolded in subsequent Figures 4 and 10.

- (1) $\xrightarrow{*} = \xrightarrow{\cup S}$, and
(2) every linear path scheme in S has $O(|Q|^2)$ many cycles and length $|\mathcal{V}|_1^{O(1)}$.

Instead of directly constructing linear path schemes for arbitrary runs, we will consider three restricted types of runs: (1) *runs staying close to the axes*, (2) *cyclic runs starting and ending far from the axes*, and (3) *runs staying far from the axes*. Here, *close* and *far* refer to whether counter values exceed a threshold D . These three types of runs are depicted in Figure 3 where $D = 5$. We will show that runs of each of these types can be captured by small linear path schemes. The proof of Theorem 3.1 will follow by observing that any run of a 2-VASS can be decomposed into polynomially many runs of the three types for a suitable threshold D .

This section is divided into two subsections. In the first subsection, we will prove the following proposition concerning flattening of runs of type 1:

PROPOSITION 4.1. *Let $D \in \mathbb{N}$ and $\mathbb{L} = ([0, D] \times \mathbb{N}) \cup (\mathbb{N} \times [0, D])$. For every 2-VASS $\mathcal{V} = (Q, T)$, there exists a finite set S of LPSs such that $\xrightarrow{*}_{\mathbb{L}} \subseteq \xrightarrow{\cup S}_{\mathbb{N}^2}$, and $|\Lambda| \leq (|Q| + \|T\| + D)^{O(1)}$ and $|\Lambda|_* \leq 2$ for every $\Lambda \in S$.*

In the second subsection, we will prove the following proposition concerning flattening of runs of types 2 and 3:

PROPOSITION 4.2. *For every 2-VASS $\mathcal{V} = (Q, T)$, there exist $\mathbb{O} = [D, \infty)^2$ and finite sets S, S' of LPSs such that $D \leq (|Q| + \|T\|)^{O(1)}$, and*

- (a) $q(\mathbf{u}) \xrightarrow{*}_{\mathbb{N}^2} q(\mathbf{v})$ if and only if $q(\mathbf{u}) \xrightarrow{\cup S}_{\mathbb{N}^2} q(\mathbf{v})$ for every $q \in Q$, $\mathbf{u}, \mathbf{v} \in \mathbb{O}$, and $|\Lambda| \leq (|Q| + \|T\|)^{O(1)}$ and $|\Lambda|_* \leq 2$ for every $\Lambda \in S$;
(b) $\xrightarrow{*}_{\mathbb{O}} \subseteq \xrightarrow{\cup S'}_{\mathbb{N}^2}$, and $|\Lambda| \leq (|Q| + \|T\|)^{O(1)}$ and $|\Lambda|_* \leq 2 \cdot |Q|$ for every $\Lambda \in S'$.

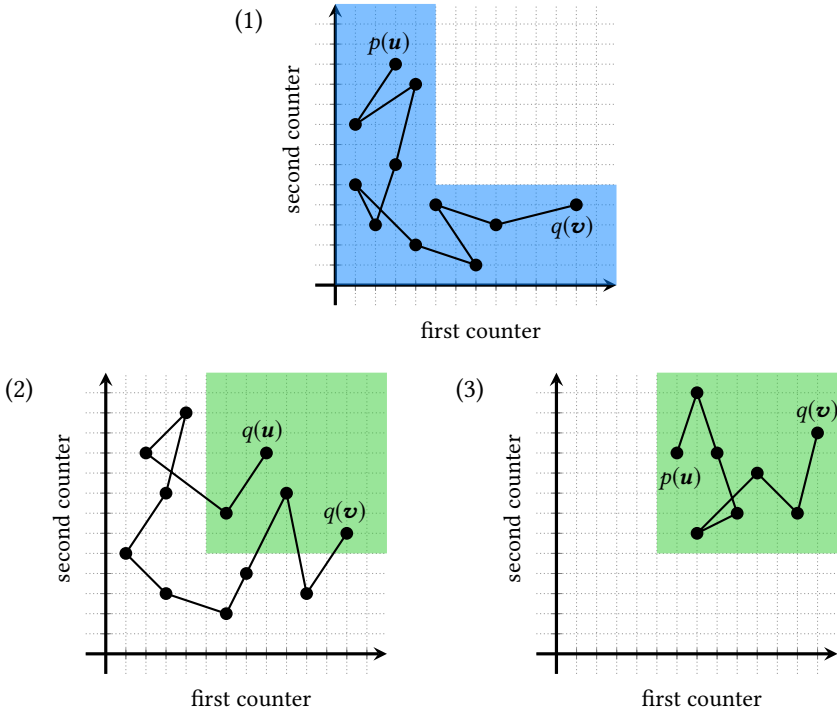


Fig. 3. Example of the three types of runs. (1) *top*: run staying close to at least one axis, i.e., running high on at most one component at a time; (2) *bottom-left*: run from q to q starting and ending sufficiently high; (3) *bottom-right*: run staying sufficiently high.

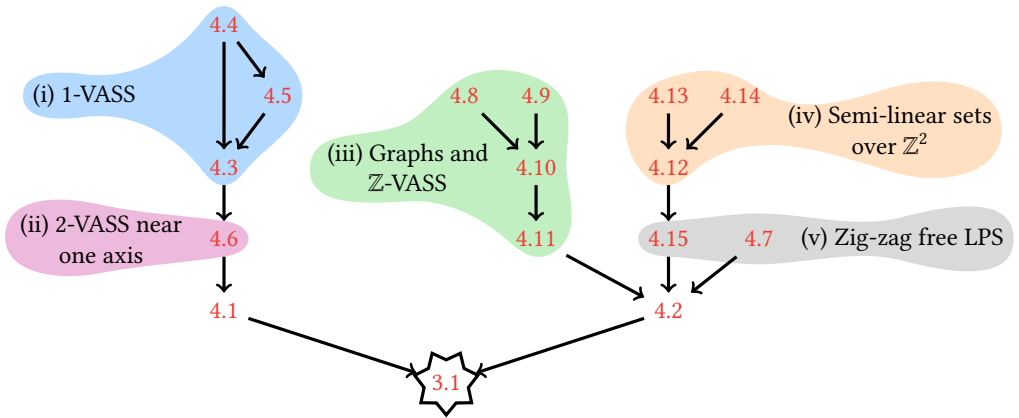


Fig. 4. Overview of the proof of Theorem 3.1. Each node labeled by x corresponds to Proposition, Theorem, Lemma or Corollary x . An arrow from node u to node v indicates that u is used in the proof of v . Each colored region corresponds to a theme which is depicted under the same color in the general overview of Figure 2. The order in which the five themes are presented are numbered from (i) to (v).

The structure of the proof of Theorem 3.1 is depicted in Figure 4. Before diving into the involved proofs of Propositions 4.1 and 4.2, let us immediately see how Theorem 3.1 follows from them:

PROOF OF THEOREM 3.1. Let $\mathcal{V} = (Q, T)$ be a 2-VASS and let D be the constant from Proposition 4.2 for \mathcal{V} . Let

$$\mathbb{L} \stackrel{\text{def}}{=} ([0, D + \|T\|] \times \mathbb{N}) \cup (\mathbb{N} \times [0, D + \|T\|]), \text{ and} \quad (2)$$

$$\mathbb{O} \stackrel{\text{def}}{=} [D, \infty)^2. \quad (3)$$

The regions \mathbb{L} and \mathbb{O} are depicted respectively in blue and green in Figure 5.

Let $R_{\mathbb{L}}$, $R_{\mathbb{O}}$ and $R'_{\mathbb{O}}$ be the sets of linear path schemes obtained for \mathcal{V} respectively from Proposition 4.1, Proposition 4.2 (a) and Proposition 4.2 (b). We claim that the following set S satisfies the claim made in the theorem:

$$S \stackrel{\text{def}}{=} \bigcup_{0 \leq h \leq |Q|} \{ \sigma_0 \Lambda_1 \sigma_1 \cdots \Lambda_h \sigma_h : \sigma_0, \sigma_1, \dots, \sigma_h \in R_{\mathbb{L}} \cup R'_{\mathbb{O}}, \Lambda_1, \Lambda_2, \dots, \Lambda_h \in R_{\mathbb{O}} \}.$$

In other words, S is made of linear path schemes obtained by concatenating alternately at most $2 \cdot |Q| + 1$ linear path schemes from $R_{\mathbb{L}} \cup R'_{\mathbb{O}}$ and $R_{\mathbb{O}}$.

Let us first prove that $\xrightarrow{*} = \xrightarrow{\cup S}$. Let $p(\mathbf{u}), q(\mathbf{v}) \in \mathbb{N}^2$ be such that $p(\mathbf{u}) \xrightarrow{*_{\mathbb{N}^2}} q(\mathbf{v})$. There exist $k \in \mathbb{N}$, $t_1, t_2, \dots, t_k \in T$ and $q_0(\mathbf{u}_0), q_1(\mathbf{u}_1), \dots, q_k(\mathbf{u}_k) \in Q \times \mathbb{N}^2$ such that

$$p(\mathbf{u}) = q_0(\mathbf{u}_0) \xrightarrow{t_1}_{\mathbb{N}^2} q_1(\mathbf{u}_1) \cdots \xrightarrow{t_k}_{\mathbb{N}^2} q_k(\mathbf{u}_k) = q(\mathbf{v}). \quad (4)$$

Intuitively, we decompose run (4) in terms of configurations whose vectors lie in $\mathbb{L} \cap \mathbb{O}$. First, we consider the smallest index i such that $\mathbf{u}_i \in \mathbb{L} \cap \mathbb{O}$ and the largest $j > i$ such that $q_j = q_i$ and $\mathbf{u}_j \in \mathbb{L} \cap \mathbb{O}$. The path from $p(\mathbf{u})$ to $q_i(\mathbf{u}_i)$ can be replaced by a path of $R_{\mathbb{L}}$ or $R'_{\mathbb{O}}$ since it remains entirely in \mathbb{L} if $\mathbf{u} \in \mathbb{L}$, or entirely in \mathbb{O} if $\mathbf{u} \in \mathbb{O}$. The path from $q_i(\mathbf{u}_i)$ to $q_j(\mathbf{u}_j)$ can be replaced by a path of $R_{\mathbb{O}}$, since $q_i = q_j$ and $\mathbf{u}_i, \mathbf{u}_j \in \mathbb{O}$. This process is repeated iteratively with the next index $i' > j$ such that $\mathbf{u}_{i'} \in \mathbb{L} \cap \mathbb{O}$, until all states have been considered.

More formally, let $I \stackrel{\text{def}}{=} \{i \in [0, k] : \mathbf{u}_i \in \mathbb{L} \cap \mathbb{O}\}$, let $\text{next} : I \rightarrow I$ be such that

$$\text{next}(i) \stackrel{\text{def}}{=} \begin{cases} \min\{j \in I : j > i\} & \text{if } i < \max(I), \\ i & \text{otherwise,} \end{cases}$$

and let $\ell : I \rightarrow I$ be such that $\ell(i) \stackrel{\text{def}}{=} \max\{j \in I : q_j = q_i\}$.

By a pigeonhole argument, there exist $0 \leq h \leq |Q|$, $i_1, i_2, \dots, i_h \in I$ and $\mathbb{C}_0, \mathbb{C}_1, \dots, \mathbb{C}_h \in \{\mathbb{L}, \mathbb{O}\}$ such that

$$\begin{aligned} q_0(\mathbf{u}_0) &\xrightarrow{*_{\mathbb{C}_0}} q_{i_1}(\mathbf{u}_{i_1}) \xrightarrow{*_{\mathbb{N}^2}} q_{\ell(i_1)}(\mathbf{u}_{\ell(i_1)}) \xrightarrow{*_{\mathbb{C}_1}} q_{i_2}(\mathbf{u}_{i_2}) \xrightarrow{*_{\mathbb{N}^2}} q_{\ell(i_2)}(\mathbf{u}_{\ell(i_2)}) \\ &\cdots \xrightarrow{*_{\mathbb{C}_{h-1}}} q_{i_h}(\mathbf{u}_{i_h}) \xrightarrow{*_{\mathbb{N}^2}} q_{\ell(i_h)}(\mathbf{u}_{\ell(i_h)}) \xrightarrow{*_{\mathbb{C}_h}} q_k(\mathbf{u}_k), \end{aligned}$$

and $i_{x+1} = \text{next}(\ell(i_x))$ for every $1 \leq x < h$. We illustrate this decomposition in Figure 5. Note that $q_{i_x} = q_{\ell(i_x)}$ and $\mathbf{u}_{i_x}, \mathbf{u}_{\ell(i_x)} \in \mathbb{O}$ for every $1 \leq x \leq h$. Therefore, by Proposition 4.2 (a), we have

$$q_{i_x}(\mathbf{u}_{i_x}) \xrightarrow{\cup R_{\mathbb{O}}}_{\mathbb{N}^2} q_{\ell(i_x)}(\mathbf{u}_{\ell(i_x)}) \text{ for every } 1 \leq x \leq h.$$

Together with Proposition 4.1 and Proposition 4.2 (b), we obtain

$$p(\mathbf{u}) \xrightarrow{(R_{\mathbb{L}} \cup R'_{\mathbb{O}}) R_{\mathbb{O}} (R_{\mathbb{L}} \cup R'_{\mathbb{O}}) \cdots R_{\mathbb{O}} (R_{\mathbb{L}} \cup R'_{\mathbb{O}})}_{\mathbb{N}^2} q(\mathbf{v}),$$

which in turn implies $p(\mathbf{u}) \xrightarrow{\cup S}_{\mathbb{N}^2} q(\mathbf{v})$ since $h \leq |Q|$.




Fig. 5. Example of the decomposition described in the proof of Theorem 3.1. Configurations corresponding to $I = \{3, 5, 6, 8, 9, 11, 12\}$ are marked by squares; regions \mathbb{L} and \mathbb{O} , defined in (2) and (3), are respectively colored in blue and green (their intersection appears in a mix of both colors); each run $q_{i_x}(u_{i_x}) \xrightarrow{*} q_{i_{x+1}}(u_{i_{x+1}})$ is colored in orange; and the remaining runs are colored in magenta. Here, $h = 3$, $C_0 = C_1 = C_3 = \mathbb{O}$, $C_2 = \mathbb{L}$, $i_1 = 3$, $\ell(i_1) = 6$, $i_2 = \ell(i_2) = 8$, $i_3 = 9$ and $\ell(i_3) = 12$.

It remains to show that S satisfies the required bounds. Let $\Lambda \in S$. By definition of S and by Propositions 4.1 and 4.2, we have

$$\begin{aligned} |\Lambda| &\leq (|Q| + 1) \cdot \max\{|\Lambda| : \Lambda \in R_{\mathbb{L}} \cup R'_{\mathbb{O}}\} + |Q| \cdot \max\{|\Lambda| : \Lambda \in R_{\mathbb{O}}\} \\ &\leq (|Q| + 1) \cdot (|Q| + \|T\| + D)^{O(1)} + |Q| \cdot (|Q| + \|T\|)^{O(1)} \\ &\leq (|Q| + \|T\|)^{O(1)} \end{aligned}$$

and

$$\begin{aligned} |\Lambda|_* &\leq (|Q| + 1) \cdot \max\{|\Lambda|_* : \Lambda \in R_{\mathbb{L}} \cup R'_{\mathbb{O}}\} + |Q| \cdot \max\{|\Lambda|_* : \Lambda \in R_{\mathbb{O}}\} \\ &\leq (|Q| + 1) \cdot (2 \cdot |Q|) + |Q| \cdot 2 \\ &\leq 6 \cdot |Q|^2. \end{aligned}$$

□

4.1 2-VASS Reachability Near the Axes

The proof strategy of Proposition 4.1 is as follows. Any run of a 2-VASS that remains close to the axes can be decomposed into polynomially many runs, each staying close to *one* axis as illustrated in Figure 6. Each run staying close to one axis has a bounded counter, and hence is essentially a run of an underlying 1-VASS. Therefore, it suffices to show that every 1-VASS can be flattened with small linear path schemes which can be lifted back to 2-VASS.

4.1.1 Flattening 1-VASS. While 1-VASS are known to be flattable [29, 42], here we show a stronger result: 1-VASS can be flattened through *small* linear path schemes. More formally, we prove the following:

PROPOSITION 4.3. *For every 1-VASS $\mathcal{V} = (Q, T)$, there exists a finite set S of LPSs such that $\xrightarrow{*}_{\mathbb{N}} = \bigcup_{S} \xrightarrow{*}_{\mathbb{N}}$, and $|\Lambda| \leq (|Q| + \|T\|)^{O(1)}$ and $|\Lambda|_* \leq 1$ for every $\Lambda \in S$.*




Fig. 6. Examples of 2-VASS runs staying close to a single axis.

We prove Proposition 4.3 as follows. First, we recall a lemma of Valiant and Paterson concerning the flattening of 1-VASS with ± 1 updates. Then, we use this lemma to bound paths witnessing reachability in such 1-VASS. Finally, we lift these results to 1-VASS without any assumption on $\|T\|$.

LEMMA 4.4 ([42, LEMMA 2]). *Let $\mathcal{V} = (Q, T)$ be a 1-VASS such that $\|T\| \leq 1$, and let $p(u), q(v) \in Q \times \mathbb{N}$. If $p(u) \xrightarrow{*}_{\mathbb{N}} q(v)$ and $|v - u| \geq |Q| + |Q|^2$, then there exist $\alpha, \beta, \gamma \in T^*$ and $i \in \mathbb{N}_{>0}$ such that*

- $p(u) \xrightarrow{\alpha\beta^i\gamma}_{\mathbb{N}} q(v)$,
- $\alpha\beta^*\gamma$ is a linear path scheme,
- $|\alpha\gamma| < |Q|^2$ and $1 \leq |\beta|, |\text{effect}(\beta)| \leq |Q|$.

LEMMA 4.5. *Let $\mathcal{V} = (Q, T)$ be a 1-VASS such that $\|T\| \leq 1$, and let $p(u), q(v) \in Q \times \mathbb{N}$. If $p(u) \xrightarrow{*}_{\mathbb{N}} q(v)$, then $p(u) \xrightarrow{\pi}_{\mathbb{N}} q(v)$ for some path π such that $|\pi| \leq O(|Q|^3) + |v - u| \cdot |Q|$.*

PROOF. Let $D \stackrel{\text{def}}{=} |v - u|$. Let us first consider the case where $D \geq |Q| + |Q|^2$. By Lemma 4.4, there exist $\alpha, \beta, \gamma \in T^*$ and $i \in \mathbb{N}_{>0}$ such that

- $p(u) \xrightarrow{\alpha\beta^i\gamma}_{\mathbb{N}} q(v)$,
- $\alpha\beta^*\gamma$ is a linear path scheme,
- $|\alpha\gamma| < |Q|^2$ and $1 \leq |\beta|, |\text{effect}(\beta)| \leq |Q|$.

Since $\|T\| \leq 1$ and $1 \leq |\text{effect}(\beta)| \leq |Q|$, we have $i \leq D + |\alpha\gamma|$. We are done since

$$\begin{aligned}
 |\alpha\beta^i\gamma| &< |Q|^2 + i \cdot |Q| \\
 &\leq |Q|^2 + (D + |\alpha\gamma|) \cdot |Q| \\
 &\leq |Q|^2 + (D + |Q|^2) \cdot |Q| \\
 &\leq 2 \cdot |Q|^3 + D \cdot |Q|.
 \end{aligned}$$

Let us now consider the case where $D < |Q| + |Q|^2$. Assume $p(u) \xrightarrow{*}_{\mathbb{N}} q(v)$, and let π be a minimal path such that $p(u) \xrightarrow{\pi}_{\mathbb{N}} q(v)$. If every configuration $r(w)$ along the run induced by π is such that $|w - v| < 2 \cdot (|Q| + |Q|^2)$, then by minimality of $|\pi|$ we have $|\pi| \leq 4 \cdot (|Q| + |Q|^2) \cdot |Q|$, and hence we are done. Therefore, there exists an intermediate configuration $r(w)$ such that $|w - v| = 2 \cdot (|Q| + |Q|^2)$.




Fig. 7. Example of the transformation of transition $t = (p, -z, q)$ into an equivalent sequence of transitions of norm at most 1.

We have:

$$\begin{aligned}
 |Q| + |Q|^2 &< |w - v| - |v - u| && \text{(by } |w - v| = 2 \cdot (|Q| + |Q|^2) \text{ and } |v - u| < |Q| + |Q|^2) \\
 &\leq |w - u| && \text{(by the triangle inequality)} \\
 &\leq |w - v| + |v - u| \\
 &< 3 \cdot (|Q| + |Q|^2) && \text{(by } |w - v| = 2 \cdot (|Q| + |Q|^2) \text{ and } |v - u| < |Q| + |Q|^2).
 \end{aligned}$$

Let π_1 and π_2 be minimal paths such that $p(u) \xrightarrow{\pi_1}_{\mathbb{N}} r(w) \xrightarrow{\pi_2}_{\mathbb{N}} q(v)$. Since $|w - u| \geq |Q| + |Q|^2$ and $|v - w| \geq |Q| + |Q|^2$, the first case considered in this proof holds from $p(u)$ to $r(w)$, and from $r(w)$ to $q(v)$. Therefore, $|\pi_1| \leq 2 \cdot |Q|^3 + |w - u| \cdot |Q|$ and $|\pi_2| \leq 2 \cdot |Q|^3 + |v - w| \cdot |Q|$, and hence

$$\begin{aligned}
 |\pi| &\leq |\pi_1| + |\pi_2| \\
 &\leq (2 \cdot |Q|^3 + |w - u| \cdot |Q|) + (2 \cdot |Q|^3 + |v - w| \cdot |Q|) \\
 &\leq (2 \cdot |Q|^3 + 3 \cdot (|Q| + |Q|^2) \cdot |Q|) + (2 \cdot |Q|^3 + 2 \cdot (|Q| + |Q|^2) \cdot |Q|) \\
 &= 9 \cdot |Q|^3 + 5 \cdot |Q|^2 \\
 &\leq 14 \cdot |Q|^3.
 \end{aligned}$$

□

We may now prove Proposition 4.3:

PROOF OF PROPOSITION 4.3. We construct a 1-VASS $\mathcal{V}' = (Q', T')$ with ± 1 updates that mimics the behavior of \mathcal{V} . This will allow us to apply Lemmas 4.4 and 4.5 to \mathcal{V}' . Each transition $t = (q, z, q') \in T$ is associated to a sequence of $|z| + 2$ transitions of \mathcal{V}' , of which $|z|$ transitions increment or decrement the counter depending on whether z is positive or not. This transformation is illustrated in Figure 7.

For every $z \in \mathbb{Z}$, let $\text{sign}(z) \stackrel{\text{def}}{=} 1$ if $z \geq 0$ and $\text{sign}(z) \stackrel{\text{def}}{=} -1$ if $z < 0$. Formally, \mathcal{V}' is defined as follows:

$$\begin{aligned}
 Q' &\stackrel{\text{def}}{=} Q \cup \{t_i : t = (p, z, q) \in T, 0 \leq i \leq |z|\}, \\
 T' &\stackrel{\text{def}}{=} \{(p, 0, t_0) : t = (p, z, q) \in T\} \cup \\
 &\quad \{(t_{|z|}, 0, q) : t = (p, z, q) \in T\} \cup \\
 &\quad \{(t_i, \text{sign}(z), t_{i+1}) : t = (p, z, q) \in T, 0 \leq i < |z|\}.
 \end{aligned}$$

Let us define the morphism $h : T \rightarrow (T')^*$ such that for every $t = (p, z, q) \in T$,

$$h(t) \stackrel{\text{def}}{=} (p, 0, t_0) \cdot \left(\prod_{i=1}^{|z|} (t_{i-1}, \text{sign}(z), t_i) \right) \cdot (t_{|z|}, 0, q).$$

It is readily seen that the image of a run of \mathcal{V} under h is a run of \mathcal{V}' . In more details, it can be shown that:

- (1) $|h(t)| \leq \|T\| + 2$ for every $t \in T$,
- (2) if $p(u) \xrightarrow{\pi}_{\mathbb{N}} q(v)$ in \mathcal{V} , then $p(u) \xrightarrow{h(\pi)}_{\mathbb{N}} q(v)$ in \mathcal{V}' , and
- (3) if $p(u) \xrightarrow{\pi'}_{\mathbb{N}} q(v)$ in \mathcal{V}' and $p, q \in Q$, then there exists a unique path $\pi \in T^*$ such that $\pi' = h(\pi)$ and $p(u) \xrightarrow{\pi}_{\mathbb{N}} q(v)$ in \mathcal{V} .

For every $p(u), q(v) \in Q \times \mathbb{N}$ and $\pi' \in (T')^*$ such that $p(u) \xrightarrow{\pi'}_{\mathbb{N}} q(v)$ in \mathcal{V}' , we write $h^{-1}(\pi')$ to denote the unique $\pi \in T^*$ given by (3). Note that π' is a cycle in \mathcal{V}' if and only if $h^{-1}(\pi')$ is a cycle in \mathcal{V} .

We claim that whenever $p(u) \xrightarrow{*}_{\mathbb{N}} q(v)$ in \mathcal{V} , there exists a linear path scheme $\Lambda \subseteq T^*$ such that $p(u) \xrightarrow{\Lambda}_{\mathbb{N}} q(v)$ in \mathcal{V} , and $|\Lambda| \leq (|Q| + \|T\|)^{O(1)}$. Since there are only finitely many such linear path schemes, the validity of the claim completes the proof. Let $D \stackrel{\text{def}}{=} |Q'| + |Q'|^2$ and assume that $p(u) \xrightarrow{*}_{\mathbb{N}} q(v)$ in \mathcal{V} . By (2), we have $p(u) \xrightarrow{*}_{\mathbb{N}} q(v)$ in \mathcal{V}' . We prove the claim by making a case distinction on whether $|u - v| \leq D$ or not.

Case 1: $|u - v| \leq D$. By Lemma 4.5, we have $p(u) \xrightarrow{\pi'}_{\mathbb{N}} q(v)$ in \mathcal{V}' for some $\pi' \in (T')^*$ such that $|\pi'| \leq O(|Q'|^3) + D \cdot |Q'|$. We set $\Lambda \stackrel{\text{def}}{=} h^{-1}(\pi')$. By (3), $p(u) \xrightarrow{\Lambda}_{\mathbb{N}} q(v)$ in \mathcal{V} . Moreover,

$$|\Lambda| \leq |\pi'| \leq O(|Q'|^3) + D \cdot |Q'| \leq (|Q| + \|T\| + D)^{O(1)} \leq (|Q| + \|T\|)^{O(1)}.$$

Case 2: $|u - v| > D$. By Lemma 4.4, there exist $\alpha, \gamma, \beta \in (T')^*$ and some $i \in \mathbb{N}_{>0}$ such that

- $p(u) \xrightarrow{\alpha(\beta)^i \gamma}_{\mathbb{N}} q(v)$ in \mathcal{V}' ,
- $\alpha\beta^* \gamma$ is a linear path scheme,
- $|\alpha\gamma| < |Q'|^2$ and $|\beta| \leq |Q'|$.

Let $q' \in Q'$ be the first state of β . If $q' \in Q$, then α, β and γ are paths of \mathcal{V}' respectively from p to q' , q' to q' , and q' to q , which all belong to Q . Thus, by (3),

$$\Lambda \stackrel{\text{def}}{=} h^{-1}(\alpha) \cdot h^{-1}(\beta)^* \cdot h^{-1}(\gamma)$$

is a linear path scheme such that $p(u) \xrightarrow{\Lambda}_{\mathbb{N}} q(v)$ in \mathcal{V} . In this case, we are done since $|\Lambda| \leq |\alpha\beta\gamma| \leq |Q'|^2 + |Q'| \leq (|Q| + \|T\|)^{O(1)}$.

Otherwise, if $q' \in Q' \setminus Q$, we have $q' = t_i$ for some $t = (q_1, z, q_2) \in T$ and some $1 \leq i \leq |z|$. Since β' is a cycle from t_i to t_i , it follows from the definition of \mathcal{V}' that $\alpha = \alpha_1\alpha_2$ and $\beta = \beta_1\beta_2$ for some $\alpha_1, \alpha_2, \beta_1, \beta_2 \in (T')^*$ such that

$$\alpha_2 = \beta_2 = (q_1, 0, t_0) \cdot \prod_{j=1}^i (t_{j-1}, \text{sign}(z), t_j).$$

Therefore, we have

$$\begin{aligned} \alpha\beta^* \gamma &= (\alpha_1\alpha_2)(\beta_1\beta_2)^* \gamma \\ &= \alpha_1(\beta_2\beta_1)^* \beta_2 \gamma \end{aligned} \quad (\text{by } \alpha_2 = \beta_2).$$

Moreover, $\alpha_1, \beta_2\beta_1$ and $\beta_2\gamma$ are paths of \mathcal{V}' respectively from p to q_1 , q_1 to q_1 , and q_1 to q , which all belong to Q . Thus, by (3),

$$\Lambda \stackrel{\text{def}}{=} h^{-1}(\alpha_1) \cdot h^{-1}(\beta_2\beta_1)^* \cdot h^{-1}(\beta_2\gamma)$$

is a linear path scheme such that $p(u) \xrightarrow{\Lambda}_{\mathbb{N}} q(v)$ in \mathcal{V} . We are done since $|\Lambda| \leq |\alpha\beta\gamma| \leq |Q'|^2 + |Q'| \leq (|Q| + \|T\|)^{O(1)}$. \square

4.1.2 Flattening 2-VASS Reachability Along a Single Axis. Let us now lift the linear path schemes obtained for 1-VASS to runs of 2-VASS staying close to a single axis. Formally, we show the following:

LEMMA 4.6. *Let $D \in \mathbb{N}$ and $\mathbb{B} \in \{\mathbb{N} \times [0, D], [0, D] \times \mathbb{N}\}$. For every 2-VASS $\mathcal{V} = (Q, T)$, there exists a finite set S of LPSs such that $\xrightarrow{*}_{\mathbb{B}} = \xrightarrow{\cup S}_{\mathbb{B}}$, and $|\Lambda| \leq (|Q| + \|T\| + D)^{O(1)}$ and $|\Lambda|_* \leq 1$ for every $\Lambda \in S$.*

PROOF. We only consider the case where $\mathbb{B} = \mathbb{N} \times [0, D]$; the other case follows symmetrically. Let $\mathcal{V} = (Q, T)$ be a 2-VASS. We construct a 1-VASS $\overline{\mathcal{V}}$ that mimics \mathcal{V} over \mathbb{B} , by encoding the finitely many possible values of the second counter into the control states. Formally, $\overline{\mathcal{V}} \stackrel{\text{def}}{=} (\overline{Q}, \overline{T})$ where

$$\begin{aligned}\overline{Q} &\stackrel{\text{def}}{=} \{q_i : q \in Q, i \in [0, D]\}, \text{ and} \\ \overline{T} &\stackrel{\text{def}}{=} \{(p_n, i, q_{n+j}) : (p, (i, j), q) \in T \text{ and } n, n+j \in [0, D]\}.\end{aligned}$$

By applying Proposition 4.3 to $\overline{\mathcal{V}}$, we obtain a finite set \overline{S} of linear path schemes such that

$$p(u) \xrightarrow{*}_{\mathbb{N}} q(v) \text{ in } \overline{\mathcal{V}} \iff p(u) \xrightarrow{\cup \overline{S}}_{\mathbb{N}} q(v) \text{ in } \overline{\mathcal{V}}, \quad (5)$$

and $|\overline{\Lambda}| \leq (|\overline{Q}| + \|\overline{T}\|)^{O(1)}$ and $|\overline{\Lambda}|_* \leq 1$ for every $\overline{\Lambda} \in \overline{S}$.

We define the morphism $\phi : \overline{T}^* \rightarrow T^*$ such that $\phi(p_n, i, q_{n+j}) \stackrel{\text{def}}{=} (p, (i, j), q)$ for every $(p_n, i, q_{n+j}) \in \overline{T}$. A simple induction shows that

$$p(u_1, u_2) \xrightarrow{*}_{\mathbb{B}} q(v_1, v_2) \text{ in } \mathcal{V} \iff p_{u_2}(u_1) \xrightarrow{*}_{\mathbb{N}} q_{v_2}(v_1) \text{ in } \overline{\mathcal{V}}, \quad (6)$$

$$p_{u_2}(u_1) \xrightarrow{\pi}_{\mathbb{N}} q_{v_2}(v_1) \text{ in } \overline{\mathcal{V}} \implies p(u_1, u_2) \xrightarrow{\phi(\pi)}_{\mathbb{B}} q(v_1, v_2) \text{ in } \mathcal{V}. \quad (7)$$

Moreover, every linear path scheme $\overline{\Lambda} = \alpha_0 \beta_1^* \alpha_1 \cdots \beta_k^* \alpha_k$ of $\overline{\mathcal{V}}$ induces a linear path scheme $\phi(\overline{\Lambda}) = \phi(\alpha_0) \phi(\beta_1)^* \phi(\alpha_1) \cdots \phi(\beta_k)^* \phi(\alpha_k)$ of \mathcal{V} . We define the set of linear path schemes S as

$$S \stackrel{\text{def}}{=} \{\phi(\overline{\Lambda}) : \overline{\Lambda} \in \overline{S}\}.$$

Since $|\overline{Q}| = (D+1) \cdot |Q|$ and $|\overline{\Lambda}| \leq (|\overline{Q}| + \|\overline{T}\|)^{O(1)}$ for every $\overline{\Lambda} \in \overline{S}$, S satisfies the appropriate bounds.

It remains to prove that $\xrightarrow{*}_{\mathbb{B}} = \xrightarrow{\cup S}_{\mathbb{B}}$. For every $p(u_1, u_2), q(v_1, v_2) \in Q \times \mathbb{B}$, we have:

$$p(u_1, u_2) \xrightarrow{*}_{\mathbb{B}} q(v_1, v_2) \text{ in } \mathcal{V} \iff p_{u_2}(u_1) \xrightarrow{*}_{\mathbb{N}} q_{v_2}(v_1) \text{ in } \overline{\mathcal{V}} \quad (\text{by (6)})$$

$$\iff p_{u_2}(u_1) \xrightarrow{\cup \overline{S}}_{\mathbb{N}} q_{v_2}(v_1) \text{ in } \overline{\mathcal{V}} \quad (\text{by (5)})$$

$$\implies p(u_1, u_2) \xrightarrow{\cup S}_{\mathbb{B}} q(v_1, v_2) \text{ in } \mathcal{V} \quad (\text{by (7)})$$

$$\implies p(u_1, u_2) \xrightarrow{*}_{\mathbb{B}} q(v_1, v_2) \text{ in } \mathcal{V}. \quad \square$$

4.1.3 Flattening 2-VASS Reachability Near the Axes. We may now prove the main proposition of this section which we recall:

PROPOSITION 4.1. *Let $D \in \mathbb{N}$ and $\mathbb{L} = ([0, D] \times \mathbb{N}) \cup (\mathbb{N} \times [0, D])$. For every 2-VASS $\mathcal{V} = (Q, T)$, there exists a finite set S of LPSs such that $\xrightarrow{*}_{\mathbb{L}} \subseteq \xrightarrow{\cup S}_{\mathbb{N}^2}$, and $|\Lambda| \leq (|Q| + \|T\| + D)^{O(1)}$ and $|\Lambda|_* \leq 2$ for every $\Lambda \in S$.*




Fig. 8. Example of the decomposition of a run staying along the axes, where $D = 3$ and $D' = D + 2 = 5$. Regions \mathbb{B}_{\uparrow} and \mathbb{B}'_{\uparrow} are colored respectively in dark hue, and dark or light hue of ■; regions \mathbb{B}_{\leftarrow} and \mathbb{B}'_{\leftarrow} are colored respectively in dark hue, and dark or light hue of ■; \odot is the dotted square region, and \odot' is the dashed square region. The seven segments associated to $\pi = \pi_1\pi_2 \cdots \pi_7$ appear alternately in ■ and ■. Vectors $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_6$ are marked as squares.

PROOF. Let us define the three following regions depicted in Figure 8:

$$\begin{aligned}\mathbb{B}_{\uparrow} &\stackrel{\text{def}}{=} [0, D] \times \mathbb{N}, \\ \mathbb{B}_{\leftarrow} &\stackrel{\text{def}}{=} \mathbb{N} \times [0, D], \text{ and} \\ \odot &\stackrel{\text{def}}{=} \mathbb{B}_{\uparrow} \cap \mathbb{B}_{\leftarrow} = [0, D] \times [0, D].\end{aligned}$$

We will first bound the number of times a minimal run within \mathbb{L} can move back and forth from \mathbb{B}_{\uparrow} to \mathbb{B}_{\leftarrow} .

Let $\mathcal{V} = (Q, T)$ be a 2-VASS, let $p(\mathbf{u}), q(\mathbf{v}) \in Q \times \mathbb{L}$, and let $\pi \in T^*$ be a minimal path such that $p(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{L}} q(\mathbf{v})$. Note that π can alternate between \mathbb{B}_{\uparrow} and \mathbb{B}_{\leftarrow} without ever entering \odot . However, such alternations necessarily go through the region extending \odot by a width of $\|T\|$. More formally, let $D' \stackrel{\text{def}}{=} D + \|T\|$, and

$$\begin{aligned}\mathbb{B}'_{\uparrow} &\stackrel{\text{def}}{=} [0, D'] \times \mathbb{N}, \\ \mathbb{B}'_{\leftarrow} &\stackrel{\text{def}}{=} \mathbb{N} \times [0, D'], \\ \mathbb{L}' &\stackrel{\text{def}}{=} \mathbb{B}'_{\uparrow} \cup \mathbb{B}'_{\leftarrow}, \text{ and} \\ \odot' &\stackrel{\text{def}}{=} \mathbb{B}'_{\uparrow} \cap \mathbb{B}'_{\leftarrow} \stackrel{\text{def}}{=} [0, D'] \times [0, D'].\end{aligned}$$

Since π is minimal and $D' \geq D$, there exist $1 \leq k \leq |Q| \cdot |\odot'|$, $\pi_1, \pi_2, \dots, \pi_k \in T^*$, $p_0(\mathbf{u}_0), p_1(\mathbf{u}_1), \dots, p_k(\mathbf{u}_k) \in Q \times \mathbb{L}'$ and $\mathbb{C}_1, \mathbb{C}_2, \dots, \mathbb{C}_k \in \{\mathbb{B}'_{\uparrow}, \mathbb{B}'_{\leftarrow}\}$ such that:

- $\pi = \pi_1\pi_2 \cdots \pi_k$
- $p_0(\mathbf{u}_0) = p(\mathbf{u}), p_k(\mathbf{u}_k) = q(\mathbf{v})$,
- $\mathbf{u}_i \in \odot'$ for every $0 < i < k$, and
- $p_{i-1}(\mathbf{u}_{i-1}) \xrightarrow{\pi_i}_{\mathbb{C}_i} p_i(\mathbf{u}_i)$ for every $0 < i \leq k$.

This decomposition is illustrated in Figure 8.

By the last point of the above enumeration together with Lemma 4.6, for every $i \in [k]$, there exists a linear path scheme Λ_i such that $p_{i-1}(\mathbf{u}_{i-1}) \xrightarrow{\Lambda_i}_{\mathbb{C}_i} p_i(\mathbf{u}_i)$, $|\Lambda_i| \leq (|Q| + \|T\| + D')^{O(1)}$ and $|\Lambda_i|_* \leq 1$. Let $i \in [k]$, and let $\Lambda_i = \alpha_i \beta_i^* \gamma_i$ where β_i is the unique cycle of Λ_i if $|\Lambda_i|_* = 1$, and ε if $|\Lambda_i|_* = 0$. There exists $e_i \in \mathbb{N}$ such that

$$p_{i-1}(\mathbf{u}_{i-1}) \xrightarrow{\alpha_i \beta_i^{e_i} \gamma_i}_{\mathbb{C}_i} p_i(\mathbf{u}_i).$$

We claim that the following linear path scheme satisfies the proposition:

$$\Lambda \stackrel{\text{def}}{=} \alpha_1 \beta_1^* \gamma_1 \cdot \left(\prod_{i=2}^{k-1} \alpha_i \beta_i^{e_i} \gamma_i \right) \cdot \alpha_k \beta_k^* \gamma_k.$$

First, note that $p(\mathbf{u}) \xrightarrow{\Lambda}_{\mathbb{N}^2} q(\mathbf{v})$ and $|\Lambda|_* \leq 2$. Now, observe that for every $i \in [k]$, we have

$$\begin{aligned} e_i &\leq \|\mathbf{u}_i - \mathbf{u}_{i-1}\| + |\alpha_i \gamma_i| \cdot \|T\| \\ &\leq \|\mathbf{u}_i - \mathbf{u}_{i-1}\| + |\Lambda_i| \cdot \|T\| \\ &\leq \|\mathbf{u}_i - \mathbf{u}_{i-1}\| + (|Q| + \|T\| + D')^{O(1)} \cdot \|T\| \\ &\leq \|\mathbf{u}_i - \mathbf{u}_{i-1}\| + (|Q| + \|T\| + D')^{O(1)}. \end{aligned}$$

Thus, for every $2 \leq i \leq k-1$, we have $e_i \leq (|Q| + \|T\| + D')^{O(1)}$ since $\mathbf{u}_{i-1}, \mathbf{u}_i \in \mathbb{O}'$. This implies that

$$\begin{aligned} |\Lambda| &\leq k \cdot \max(1, e_2, e_3, \dots, e_{k-1}) \cdot \max(|\Lambda_1|, |\Lambda_2|, \dots, |\Lambda_k|) \\ &\leq |Q| \cdot |\mathbb{O}'| \cdot (|Q| + \|T\| + D')^{O(1)} \cdot (|Q| + \|T\| + D')^{O(1)} \\ &\leq (|Q| + \|T\| + D)^{O(1)}. \end{aligned}$$

To conclude, note that we have constructed a linear path scheme Λ for a specific pair of configurations $p(\mathbf{u})$ and $q(\mathbf{v})$. Nonetheless, the bounds on $|\Lambda|$ and $|\Lambda|_*$ are independent from \mathbf{u} and \mathbf{v} . Hence, taking S as the set of all linear path schemes satisfying these bounds proves the proposition. \square

4.2 2-VASS Reachability Far From the Axes

It remains to deal with the two other types of runs, illustrated at the bottom of Figure 3. As discussed at the beginning of the section, we will show the following:

PROPOSITION 4.2. *For every 2-VASS $\mathcal{V} = (Q, T)$, there exist $\mathbb{O} = [D, \infty)^2$ and finite sets S, S' of LPSs such that $D \leq (|Q| + \|T\|)^{O(1)}$, and*

- (a) $q(\mathbf{u}) \xrightarrow{*}_{\mathbb{N}^2} q(\mathbf{v})$ if and only if $q(\mathbf{u}) \xrightarrow{\bigcup S}_{\mathbb{N}^2} q(\mathbf{v})$ for every $q \in Q$, $\mathbf{u}, \mathbf{v} \in \mathbb{O}$, and $|\Lambda| \leq (|Q| + \|T\|)^{O(1)}$ and $|\Lambda|_* \leq 2$ for every $\Lambda \in S$;
- (b) $\xrightarrow{*}_{\mathbb{O}} \subseteq \xrightarrow{\bigcup S'}_{\mathbb{N}^2}$, and $|\Lambda| \leq (|Q| + \|T\|)^{O(1)}$ and $|\Lambda|_* \leq 2 \cdot |Q|$ for every $\Lambda \in S'$.

Proposition 4.2 (b) will follow easily from Proposition 4.2 (a) which will be proven as follows. First, we will show that the relation $\{(\mathbf{u}, \mathbf{v}) : p(\mathbf{u}) \xrightarrow{*}_{\mathbb{Z}^d} q(\mathbf{v})\}$ of any d -VASS can be flattened with small linear path schemes. Then, we will show that whenever $d = 2$ and $p = q$, these linear path schemes can be converted to equivalent so-called zigzag-free linear path schemes, by exploiting special properties of linear subsets of \mathbb{Z}^2 . Finally, we will make use of the fact that \mathbb{Z} -reachability and reachability coincide for runs induced by zigzag-free linear path schemes and taking place far enough from the axes.




Fig. 9. Example of two runs induced by the same path π of a zigzag-free linear path scheme $\Lambda = \alpha_0\beta_1^*\alpha_1\beta_2^*\alpha_2$ such that $\text{effect}(\beta_1), \text{effect}(\beta_2) \in \mathbb{N}^2$. Arrows colored in ■ correspond to the effects of α_0, α_1 and α_2 , and arrows colored in ■ correspond to the effects of β_1 and β_2 .

Let us explain this last observation in detail. We say that a linear path scheme $\Lambda = \alpha_0\beta_1^*\alpha_1 \cdots \beta_k^*\alpha_k$ of some d -VASS is *zigzag-free* [29] if for every $i \in [d]$, either

$$\bigwedge_{1 \leq j \leq k} \text{effect}(\beta_j)(i) \geq 0 \text{ or } \bigwedge_{1 \leq j \leq k} \text{effect}(\beta_j)(i) \leq 0.$$

In other words, Λ is zigzag-free if $\text{effect}(\text{cycles}(\Lambda)) \subseteq Z$ for some hyperoctant Z of \mathbb{Z}^d .

As an example, let us consider a zigzag-free linear path scheme $\Lambda = \alpha_0\beta_1^*\alpha_1\beta_2^*\alpha_2$ such that $\text{effect}(\beta_1), \text{effect}(\beta_2) \in \mathbb{N}^2$. A run over \mathbb{Z}^2 induced by a path $\pi \in \Lambda$ can only drift away from \mathbb{N}^2 by a constant distance since only α_0, α_1 and α_2 may contribute negatively to the effect of π . Therefore, if the initial and target configurations u and v are sufficiently high, then π induces a run over \mathbb{N}^2 , as illustrated in Figure 9. This intuition is formalized as follows:

LEMMA 4.7 ([29, LEMMA 4.6]). *For every d -VASS $\mathcal{V} = (Q, T)$, every $p(u), q(v) \in Q \times \mathbb{N}^d$, and every zigzag-free linear path scheme Λ , if $p(u) \xrightarrow{\Lambda}_{\mathbb{Z}^d} q(v)$ and $\|u\|, \|v\| \geq |\Lambda| \cdot \|T\|$, then $p(u) \xrightarrow{\Lambda}_{\mathbb{N}^d} q(v)$.*

4.2.1 Flattening \mathbb{Z} -Reachability. Here we show that the \mathbb{Z} -reachability relation of a d -VASS can be flattened with small linear path schemes. First we give some definitions and prove technical lemmas on finite directed graph Parikh images.

Let $G = (U, E)$ be a finite directed graph. For every $u \in U$, let

$$\text{in}(u) \stackrel{\text{def}}{=} \{(u', a, u'') \in E : u'' = u\}, \text{ and}$$

$$\text{out}(u) \stackrel{\text{def}}{=} \{(u', a, u'') \in E : u' = u\},$$

denote respectively the set of incoming and outgoing arcs of u . For every path π of G , we define the *Parikh image* of π as:

$$\text{Parikh}(\pi) \stackrel{\text{def}}{=} \sigma \in \mathbb{N}^E \text{ where } \sigma(e) \text{ is the number of occurrences of } e \text{ in } \pi.$$

Parikh images are naturally extended to path languages, i.e., for every $L \subseteq E^*$, we define $\text{Parikh}(L) \stackrel{\text{def}}{=} \{\text{Parikh}(\pi) : \pi \in L\}$. For every $\sigma \in \mathbb{N}^E$, we say that σ is a *flow* if for every $u \in U$ we have

$$\sum_{e \in \text{in}(u)} \sigma(e) = \sum_{e \in \text{out}(u)} \sigma(e).$$

We show that the Parikh image of a path can be decomposed into a short “base” path that visits each of its vertices, together with a flow:

PROPOSITION 4.8. *Let $G = (U, E)$ be a finite directed graph and let π be a path of G from $p \in U$ to $q \in U$. There exist a path π' from p to q and a flow σ such that*

- (a) $|\pi'| \leq |U|^2$ and π' visits each vertex of π at least once, and
- (b) $\text{Parikh}(\pi) = \text{Parikh}(\pi') + \sigma$.

PROOF. We construct π' and σ by repeatedly removing cycles β from π while keeping its set of vertices unchanged, and by repeatedly incrementing a vector by $\text{Parikh}(\beta)$. Similar constructions appear, e.g., in [39, proof of Lemma 4.5]. In more details, we construct a sequence of paths ρ_0, ρ_1, \dots and vectors $\mathbf{x}_0, \mathbf{x}_1, \dots$ that stabilizes at some index m , and we pick π' and σ respectively as ρ_m and \mathbf{x}_m . The sequences are defined as follows. Let $\rho_0 \stackrel{\text{def}}{=} \pi$ and $\mathbf{x}_0 \stackrel{\text{def}}{=} \mathbf{0}$. For every $i > 0$:

- ρ_{i-1} can be decomposed as $\rho_{i-1} = e_1 \pi_1 \cdots e_k \pi_k$ where $k \leq |U|$ and each $e_j = (u, a, u')$ is the first edge such that u or u' appears in ρ_{i-1} ;
- if each π_j is cycle-free, then $\rho_i \stackrel{\text{def}}{=} \rho_{i-1}$ and $\mathbf{x}_i \stackrel{\text{def}}{=} \mathbf{x}_{i-1}$ (we are done);
- otherwise if some π_j contains a cycle β , then ρ_i is defined as the path obtained by removing β from π_j , and $\mathbf{x}_i \stackrel{\text{def}}{=} \mathbf{x}_{i-1} + \text{Parikh}(\beta)$.

The resulting path π' is such that $|\pi'| \leq |U|^2$. Since σ is the sum of Parikh images of cycles, it is a flow. Moreover, $\text{Parikh}(\pi) = \text{Parikh}(\pi') + \sigma$ by construction. Therefore, (a) and (b) are satisfied. \square

We now show that any flow is a linear combination of few maps arising from simple cycles:

PROPOSITION 4.9. *Let $G = (U, E)$ be a finite directed graph and let $\sigma \in \mathbb{N}^E$ be a flow. There exist $h \leq |E|$, $c_1, \dots, c_h \in \mathbb{N}$ and $\sigma_1, \dots, \sigma_h \in \mathbb{N}^E$ such that $\sigma = c_1 \cdot \sigma_1 + \dots + c_h \cdot \sigma_h$ and each σ_i is induced by a simple cycle, i.e. $\sigma_i = \text{Parikh}(\beta_i)$ for some simple cycle β_i .*

PROOF. For every $\mathbf{x} \in \mathbb{N}^E$, let $E_{\mathbf{x}} \stackrel{\text{def}}{=} \{e \in E : \mathbf{x}(e) > 0\}$. We show a stronger claim, namely that the proposition holds for some $h \leq |E_{\sigma}|$. We proceed by induction on $|E_{\sigma}|$. If $|E_{\sigma}| = 0$, then $\sigma = \mathbf{0}$ and the claim holds trivially. Assume that $|E_{\sigma}| > 0$. Since σ is a flow, there exists a function $\chi : E_{\sigma} \rightarrow E_{\sigma}$ such that

$$\chi(x, a, y) = (x', a, y') \implies y = x'.$$

By the pigeonhole principle, there exist $e \in E_{\sigma}$ and $\ell \geq 0$ such that

$$\beta \stackrel{\text{def}}{=} e \cdot \chi(e) \cdot \chi^2(e) \cdots \chi^{\ell}(e)$$

is a simple cycle. Without loss of generality, we may assume that $\sigma(e)$ is minimal among the edges of β , i.e. $\sigma(e) = \min\{\sigma(\chi^j(e)) : 0 \leq j \leq \ell\}$.

Let $d \stackrel{\text{def}}{=} \sigma(e)$ and $\sigma' \stackrel{\text{def}}{=} \sigma - \text{Parikh}(\beta^d)$. Note that σ' is a flow since β is a cycle, and $\sigma' \in \mathbb{N}^E$ by minimality of d . Moreover, $|E_{\sigma'}| < |E_{\sigma}|$ by $e \in E_{\sigma} \setminus E_{\sigma'}$. Thus, by induction hypothesis, there exist $h' \leq |E_{\sigma'}|$, $c_1, \dots, c_{h'} \in \mathbb{N}$ and $\sigma_1, \dots, \sigma_{h'} \in \mathbb{N}^E$ such that $\sigma' = c_1 \cdot \sigma_1 + \dots + c_{h'} \cdot \sigma_{h'}$. Let $h \stackrel{\text{def}}{=} h' + 1$, $c_h \stackrel{\text{def}}{=} d$ and $\sigma_h \stackrel{\text{def}}{=} \text{Parikh}(\beta)$. We are done since $h \leq |E_{\sigma}|$, β is simple and

$$\begin{aligned} \sigma &= \sigma' + \text{Parikh}(\beta^d) \\ &= \sigma' + d \cdot \text{Parikh}(\beta) \\ &= c_1 \cdot \sigma_1 + \dots + c_h \cdot \sigma_h. \end{aligned}$$

\square

Let $\text{paths}(G, p, q)$ denote the set of all paths from vertex p to vertex q in a directed graph G . We derive the following lemma from Propositions 4.8 and 4.9:

LEMMA 4.10. For every finite directed graph $G = (U, E)$ and $p, q \in U$, there exists a finite set S of LPSs from p to q such that $\text{Parikh}(\text{paths}(G, p, q)) = \text{Parikh}(S)$, and $|\Lambda| \leq |U| \cdot (|U| + |E|)$ and $|\Lambda|_* \leq |E|$ for every $\Lambda \in S$.

PROOF. Let $G = (U, E)$ be a finite directed graph. We claim that the following set of linear path schemes satisfies the lemma:

$$S \stackrel{\text{def}}{=} \{\Lambda \subseteq \text{paths}(G, p, q) : \Lambda \text{ is an LPS, } |\Lambda| \leq |U| \cdot (|U| + |E|) \text{ and } |\Lambda|_* \leq |E|\}.$$

Obviously, S satisfies the right bounds, and $\text{Parikh}(S) \subseteq \text{Parikh}(\text{paths}(G, p, q))$. Thus, it suffices to show that $\text{Parikh}(\text{paths}(G, p, q)) \subseteq \text{Parikh}(S)$.

Let π be a path of G from p to q . Let $\pi' \in \text{paths}(G, p, q)$ and $\sigma \in \mathbb{N}^E$ be the path and the flow given by Proposition 4.8 for π . Let $c_1, \dots, c_h \in \mathbb{N}$ and $\sigma_1, \dots, \sigma_h \in \mathbb{N}^E$ be given by Proposition 4.9 for σ . Recall that each σ_i is induced by some simple cycle β_i . Moreover, π' visits all states of each β_i . Thus, we can insert β_i^* along π' , for each i , to obtain a linear path scheme Λ .

By Propositions 4.8 and 4.9, we have $|\pi'| \leq |U|^2$ and $h \leq |E|$. Consequently, $|\Lambda| \leq |\pi'| + h \cdot |U| \leq |U|^2 + |E| \cdot |U| = |U| \cdot (|U| + |E|)$ and $|\Lambda|_* \leq h \leq |E|$. Thus, we are done since the following holds:

$$\begin{aligned} \text{Parikh}(\pi) &= \text{Parikh}(\pi') + \sigma && \text{(by Proposition 4.8)} \\ &= \text{Parikh}(\pi') + c_1 \cdot \sigma_1 + \dots + c_h \cdot \sigma_h && \text{(by Proposition 4.9)} \\ &\in \text{Parikh}(\pi') + \text{Parikh}(\beta_1^*) + \dots + \text{Parikh}(\beta_h^*) && \text{(by } c_i \cdot \sigma_i = \text{Parikh}(\beta_i^{c_i}) \text{)} \\ &= \text{Parikh}(\Lambda) && \text{(by def. of } \Lambda \text{).} \quad \square \end{aligned}$$

Lemma 4.10 allows us to show that the \mathbb{Z} -reachability relation of a d -VASS can be flattened with small linear path schemes:

PROPOSITION 4.11. For every d -VASS $\mathcal{V} = (Q, T)$, there exists a finite set S of LPSs such that $\xrightarrow{*}_{\mathbb{Z}^d} = \xrightarrow{\cup S}_{\mathbb{Z}^d}$, and $|\Lambda| \leq |Q| \cdot (|Q| + |T|)$ and $|\Lambda|_* \leq |T|$ for every $\Lambda \in S$.

PROOF. Let $\mathcal{V} = (Q, T)$ be a d -VASS. For every $p, q \in Q$, let $S_{p,q}$ be the finite set of linear path schemes obtained from Lemma 4.10. We claim that the proposition is satisfied by

$$S \stackrel{\text{def}}{=} \bigcup_{p, q \in Q} S_{p,q}.$$

Clearly, S satisfies the required bounds, and $\xrightarrow{\cup S}_{\mathbb{Z}^d} \subseteq \xrightarrow{*}_{\mathbb{Z}^d}$. It remains to show that $\xrightarrow{*}_{\mathbb{Z}^d} \subseteq \xrightarrow{\cup S}_{\mathbb{Z}^d}$. Let $p(\mathbf{u}), q(\mathbf{v}) \in Q \times \mathbb{Z}^d$ and $\pi \in T^*$ be such that $p(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{Z}^d} q(\mathbf{v})$. We have

$$\mathbf{v} = \mathbf{u} + \sum_{t \in T} \text{Parikh}(\pi)(t) \cdot \text{effect}(t). \quad (8)$$

By Lemma 4.10, there exist $\Lambda \in S_{p,q}$ and $\pi' \in \Lambda$ such that $\text{Parikh}(\pi') = \text{Parikh}(\pi)$. By (8), this implies that

$$\mathbf{v} = \mathbf{u} + \sum_{t \in T} \text{Parikh}(\pi')(t) \cdot \text{effect}(t),$$

which in turn implies $p(\mathbf{u}) \xrightarrow{\pi'}_{\mathbb{Z}^d} q(\mathbf{v})$. \square

4.2.2 A Decomposition of Certain Linear Subsets of \mathbb{Z}^2 . Our main result here, namely Lemma 4.12 below, appears quite technical at first. However, it will play an essential role in transforming the flattenings of the \mathbb{Z} -reachability relation obtained previously into zigzag-free ones, and thus in obtaining flattenings of the reachability relation far from the axes.

LEMMA 4.12. Let $P \subset^{\text{fin}} \mathbb{Z}^2$ and $\mathbf{b} \in P$. For every quadrant Z , there exist

$$D \subseteq Z \cap L_{[0, O(\|P\|^2 \|P\|^8)]}(\mathbf{b}, P) \text{ and } Q \subseteq Z \cap (P \cup L_{[0, O(\|P\|^3)]}(\mathbf{b}, P))$$

such that $|Q| \leq 2$ and $L(\mathbf{b}, P) \cap Z = L(D, Q)$.

To assist the reader in making sense of the statement of Lemma 4.12, we sketch the context of its forthcoming application in the proof of Lemma 4.15: P should be thought of as consisting of the effects of some cycles $\alpha_0 \alpha_1 \cdots \alpha_k$, β_1, \dots, β_k , and \mathbf{b} as the effect of the cycle $\alpha_0 \alpha_1 \cdots \alpha_k$, where $\sigma = \alpha_0 \beta_1^* \alpha_1 \cdots \beta_k^* \alpha_k$ is a linear path scheme from some state to itself in some 2-VASS. Hence the linear set $L(\mathbf{b}, P)$ consists of the effects of all traversals of σ one or more times. Lemma 4.12 can then be seen as stating that, for any of the four quadrants Z , all such effects in Z are also in some hybrid-linear set $L(D, Q)$ such that:

- D and Q are in the quadrant Z ;
- their elements are effects of short (polynomial in $\|P\|$) traversals of σ , except for Q which may also contain elements of P ;
- the cardinality of Q is at most 2.

As the stepping stone towards Lemma 4.12, we prove the following relatively simple result about intersections of the rational cones spanned by P with quadrants. It says that they can be spanned by at most two vectors which are either from P or small linear combinations of elements of P where the coefficient of \mathbf{b} is at least 1.

The latter requirement may seem minor, but it will be crucial for proving Lemma 4.12. Moreover, it makes Lemma 4.13 difficult to generalise to dimensions beyond 2. Specifically, the infamous Hopcroft-Pansiot example [19, proof of Lemma 2.8] gives us

$$P = \{(1, 0, 0), (0, 1, -1), (0, -1, 2)\} \text{ and } \mathbf{b} = (1, 0, 0)$$

such that the intersection of the rational cone of P with the non-negative octant, which equals the non-negative octant, cannot be spanned only by vectors that are either from P or linear combinations of elements of P in which \mathbf{b} features positively (see [29, Remark 5.2]).

LEMMA 4.13. Let $P \subset^{\text{fin}} \mathbb{Z}^2$ and $\mathbf{b} \in P$. For every quadrant Z of the rational plane, there exists $Q \subseteq P \cup L_{[0, 6\|P\|^3]}(\mathbf{b}, P)$ such that $|Q| \leq 2$ and $L_{\mathbb{Q}_{\geq 0}}(\mathbf{0}, P) \cap Z = L_{\mathbb{Q}_{\geq 0}}(\mathbf{0}, Q)$.

PROOF. Without loss of generality, let us focus on the case where Z is the upper-right quadrant $L_{\mathbb{Q}_{\geq 0}}(\mathbf{0}, \{(1, 0), (0, 1)\})$.

Since we are in the plane, the intersection of the two rational cones $L_{\mathbb{Q}_{\geq 0}}(\mathbf{0}, P)$ and Z is a rational cone $L_{\mathbb{Q}_{\geq 0}}(\mathbf{0}, \{\mathbf{q}_1, \mathbf{q}_2\})$ such that, for each $i \in \{1, 2\}$:

- either $\mathbf{q}_i \in P$,
- or $\mathbf{q}_i \in \{(1, 0), (0, 1)\}$ and \mathbf{q}_i is a linear combination with positive rational coefficients of two linearly independent vectors in P .

To finish the proof, we show how, in each latter case, \mathbf{q}_i can be scaled by a positive integer to be a member of $L_{[0, 6\|P\|^3]}(\mathbf{b}, P)$. Again without loss of generality, we consider the case where $(1, 0)$ is a linear combination with positive rational coefficients of two linearly independent vectors $\mathbf{p}_1, \mathbf{p}_2 \in P$.

Let us recall Cramer's Rule: for all $\mathbf{c} \in \mathbb{Q}^2$, the unique solution $\lambda_1, \lambda_2 \in \mathbb{Q}$ of $\lambda_1 \mathbf{p}_1 + \lambda_2 \mathbf{p}_2 = \mathbf{c}$ is:

$$\lambda_1 = \frac{\mathbf{c}(1)\mathbf{p}_2(2) - \mathbf{p}_2(1)\mathbf{c}(2)}{\mathbf{p}_1(1)\mathbf{p}_2(2) - \mathbf{p}_2(1)\mathbf{p}_1(2)}, \quad \lambda_2 = \frac{\mathbf{p}_1(1)\mathbf{c}(2) - \mathbf{c}(1)\mathbf{p}_1(2)}{\mathbf{p}_1(1)\mathbf{p}_2(2) - \mathbf{p}_2(1)\mathbf{p}_1(2)}.$$

Note that $|\mathbf{p}_i(j)|, |\mathbf{b}(j)| \leq \|P\|$ for all $i, j \in \{1, 2\}$. Thus, we can invoke Cramer's Rule with $\mathbf{c} = (1, 0)$ and $\mathbf{c} = -\mathbf{b}$, and then multiply both sides by the (integral) denominator and by 1 or -1 to obtain:

- $x(1, 0) = x'_1 p_1 + x'_2 p_2$ for some $x \in [1, 2\|P\|^2]$ and $x'_1, x'_2 \in [1, \|P\|]$;
- $0 = yb + y'_1 p_1 + y'_2 p_2$ for some $y \in [1, 2\|P\|^2]$ and $y'_1, y'_2 \in [-2\|P\|^2, 2\|P\|^2]$.

Summing the multiple of the former equation by $2\|P\|^2$ and the latter equation, we get:

$$2x\|P\|^2(1, 0) = yb + (2x'_1\|P\|^2 + y'_1) p_1 + (2x'_2\|P\|^2 + y'_2) p_2.$$

Since $y \in [1, 2\|P\|^2]$ and, for each $i \in \{1, 2\}$, we have $2x'_i\|P\|^2 + y'_i \in [0, 4\|P\|^3]$, we conclude that $2x\|P\|^2(1, 0) \in L_{[0, 6\|P\|^3]}(b, P)$ as required. \square

We also recall a result on the sets of natural solutions of linear equalities that follows from results of Pottier [38]. We state it as generally as in the paper of Chistikov and Haase, although we shall apply it only with $d = 2$.

PROPOSITION 4.14 ([8, PROPOSITION 4]). *Let $\mathcal{E}_0: A \cdot \mathbf{x} = \mathbf{0}$ and $\mathcal{E}: A \cdot \mathbf{x} = \mathbf{b}$ be systems of linear Diophantine equations, where $A \in \mathbb{Z}^{d \times k}$. Then their sets of solutions in the naturals are of the forms $L(\mathbf{0}, R)$ and $L(C, R)$ (respectively), such that*

$$\|C\| \leq ((k+1)\|A\| + \|\mathbf{b}\| + 1)^d \text{ and } \|R\| \leq (k\|A\| + 1)^d.$$

LEMMA 4.12. *Let $P \subset^{fin} \mathbb{Z}^2$ and $\mathbf{b} \in P$. For every quadrant Z , there exist*

$$D \subseteq Z \cap L_{[0, O(\|P\|^2\|P\|^8)]}(\mathbf{b}, P) \text{ and } Q \subseteq Z \cap (P \cup L_{[0, O(\|P\|^3)]}(\mathbf{b}, P))$$

such that $|Q| \leq 2$ and $L(\mathbf{b}, P) \cap Z = L(D, Q)$.

PROOF. It is easy to see that the statement of the lemma is implied by the version obtained by replacing the last equality with the inclusion $L(\mathbf{b}, P) \cap Z \subseteq L(D, Q)$.

Let Q be obtained from Lemma 4.13. Since $L_{\mathbb{Q}_{\geq 0}}(\mathbf{0}, P) \cap Z = L_{\mathbb{Q}_{\geq 0}}(\mathbf{0}, Q)$, we have that $Q \subseteq Z$.

Suppose that $\mathbf{a} \in L(\mathbf{b}, P) \cap Z$. It will suffice to show that there exists $\mathbf{d} \in Z \cap L_{[0, O(\|P\|^2\|P\|^8)]}(\mathbf{b}, P)$ such that $\mathbf{a} \in L(\mathbf{d}, Q)$.

Recalling that $\mathbf{b} \in P$, we have that $\mathbf{a} \in L_{\mathbb{Q}_{\geq 0}}(\mathbf{0}, P) \cap Z$, and so $\mathbf{a} \in L_{\mathbb{Q}_{\geq 0}}(\mathbf{0}, Q)$. By considering the integral and fractional parts of the rational coefficients, it follows that $\mathbf{a} \in L(\mathbf{a}', Q)$ for some \mathbf{a}' arising from the fractional parts. Since $Q \subseteq Z$ and $\mathbf{a}' \in L_{\mathbb{Q}_{\geq 0}}(\mathbf{0}, Q)$, we have $\mathbf{a}' \in Z$. Moreover, since $\mathbf{a} \in \mathbb{Z}^2$ and $Q \subseteq \mathbb{Z}^2$, we have $\mathbf{a}' \in \mathbb{Z}^2$. Altogether, we obtain $\mathbf{a}' \in Z \cap \mathbb{Z}^2$ and $\|\mathbf{a}'\| \leq 2\|Q\| \leq 14\|P\|^4$.

Since also $\mathbf{a} \in L(\mathbf{b}, P)$, we have that the system of linear Diophantine equations

$$\begin{bmatrix} P & -Q \end{bmatrix} \begin{pmatrix} \mathbf{x} \\ \mathbf{y} \end{pmatrix} = \mathbf{a}' - \mathbf{b},$$

where P and Q are written as matrices whose columns are the vectors in P and Q (respectively), has a solution in the naturals such that $\mathbf{b} + P \cdot \mathbf{x} = \mathbf{a} = \mathbf{a}' + Q \cdot \mathbf{y}$.

By Proposition 4.14, there exist vectors of naturals $\mathbf{x}_C, \mathbf{x}_R, \mathbf{y}_C$ and \mathbf{y}_R such that $\mathbf{x} = \mathbf{x}_C + \mathbf{x}_R$, $\mathbf{y} = \mathbf{y}_C + \mathbf{y}_R$, $P \cdot \mathbf{x}_R = Q \cdot \mathbf{y}_R$ and

$$\begin{aligned} \|\mathbf{x}_C\| &\leq ((|P| + |Q| + 1) \cdot (\max\{\|P\|, \|Q\|\}) + \|\mathbf{a}'\| + \|\mathbf{b}\| + 1)^2 \leq \\ &((|P| + 3) \cdot 7\|P\|^4 + 14\|P\|^4 + \|P\| + 1)^2 = O(|P|^2\|P\|^8). \end{aligned}$$

Letting $\mathbf{d} = \mathbf{b} + P \cdot \mathbf{x}_C$, it remains to observe that $\mathbf{a} = \mathbf{b} + P \cdot (\mathbf{x}_C + \mathbf{x}_R) = \mathbf{d} + Q \cdot \mathbf{y}_R$ and to confirm that $\mathbf{d} = \mathbf{a}' + Q \cdot \mathbf{y}_C \in Z$. \square

4.2.3 Zigzag-Free Linear Path Schemes: Proof of Proposition 4.2. So far, we have seen that \mathbb{Z} -reachability can be flattened with small linear path schemes, and that some linear subsets of \mathbb{Z}^2 decompose nicely. To prove Proposition 4.2, it remains to make these linear path schemes zigzag-free by using this decomposition. This is done in the following lemma:

LEMMA 4.15. *For every 2-VASS $\mathcal{V} = (Q, T)$, every $q \in Q$, and every linear path scheme σ from q to q , there exists a finite set S of zigzag-free LPSs from q to q such that $\text{effect}(\sigma) \subseteq \text{effect}(S)$, and $|\Lambda| \leq (|\sigma| + \|T\|)^{O(1)}$ and $|\Lambda|_* \leq 2$ for every $\Lambda \in S$.*

PROOF. Let $\mathcal{V} = (Q, T)$ be a 2-VASS, let $q \in Q$, and let $\sigma = \alpha_0 \beta_1^* \alpha_1 \cdots \beta_k^* \alpha_k$ be a linear path scheme from q to q . Let

$$\sigma' \stackrel{\text{def}}{=} (\alpha_0 \cdots \alpha_k)^* \sigma.$$

Note that $\text{effect}(\sigma) \subseteq \text{effect}(\sigma') = \bigcup_{\text{quadrant } Z} \text{effect}(\sigma') \cap Z$. Thus, it suffices to exhibit a set S_Z of linear path schemes, for every quadrant Z , satisfying the desired bounds and such that $\text{effect}(\sigma') \cap Z = \text{effect}(S_Z)$. The proof is then completed by taking $S \stackrel{\text{def}}{=} \bigcup_{\text{quadrant } Z} S_Z$.

Let Z be a quadrant, let $\mathbf{b} \stackrel{\text{def}}{=} \text{effect}(\alpha_0 \cdots \alpha_k)$ and let $P \stackrel{\text{def}}{=} \text{effect}(\text{cycles}(\sigma'))$. Note that $\mathbf{b} \in P$ and $\text{effect}(\sigma') = L(\mathbf{b}, P)$. By Lemma 4.12, there exist $e \leq \|P\|^{O(1)}$, $D \subseteq Z \cap L_{[0, e]}(\mathbf{b}, P)$ and $Q \subseteq Z \cap (P \cup L_{[0, e]}(\mathbf{b}, P))$ such that $|Q| \leq 2$ and $L(\mathbf{b}, P) \cap Z = L(D, Q)$.

For every $\mathbf{u} \in D \cup (Q \setminus P)$, there exists a path $\pi_{\mathbf{u}} \in \sigma'$ such that $\text{effect}(\pi_{\mathbf{u}}) = \mathbf{u}$, and of the form $\alpha_0 \beta_1^{e_1} \alpha_1 \cdots \beta_k^{e_k} \alpha_k$ for some $0 \leq e_1, e_2, \dots, e_k \leq e$.

Let $\mathbf{d} \in D$. Then $\pi_{\mathbf{d}}$ is of the form $\alpha_0 \beta_1^{e_1} \alpha_1 \cdots \beta_k^{e_k} \alpha_k$. Let us define the linear path scheme $\Lambda_{\mathbf{d}}$ as

$$\Lambda_{\mathbf{d}} \stackrel{\text{def}}{=} \alpha_0 \beta_1^{e_1} \theta_1 \alpha_1 \cdots \beta_k^{e_k} \theta_k \alpha_k \cdot \prod_{\mathbf{u} \in Q \setminus P} \pi_{\mathbf{u}}^*$$

where

$$\theta_j \stackrel{\text{def}}{=} \begin{cases} \beta_j^* & \text{if } \text{effect}(\beta_j) \in Q \cap P \\ \varepsilon & \text{otherwise} \end{cases} \quad \text{for every } j \in [k].$$

By definition of $\Lambda_{\mathbf{d}}$, we have $\text{effect}(\Lambda_{\mathbf{d}}) = L(\mathbf{d}, Q)$. Moreover, $|\Lambda_{\mathbf{d}}|_* = |Q| \leq 2$, $\text{effect}(\text{cycles}(\Lambda_{\mathbf{d}})) = Q \subseteq Z$, and

$$\begin{aligned} |\Lambda_{\mathbf{d}}| &\leq 3 \cdot (1 + e) \cdot |\sigma'| + 2 \cdot |\sigma'| \\ &\leq 8 \cdot e \cdot |\sigma'| \\ &\leq 8 \cdot \|P\|^{O(1)} \cdot (2 \cdot |\sigma|) \\ &\leq 8 \cdot (|\sigma| + \|T\|)^{O(1)} \cdot (2 \cdot |\sigma|) \\ &\leq (|\sigma| + \|T\|)^{O(1)}. \end{aligned}$$

Therefore, we are done by taking $S_Z \stackrel{\text{def}}{=} \bigcup_{\mathbf{d} \in D} \Lambda_{\mathbf{d}}$, since

$$\text{effect}(\sigma') \cap Z = L(\mathbf{b}, P) \cap Z = L(D, Q) = \bigcup_{\mathbf{d} \in D} \text{effect}(\Lambda_{\mathbf{d}}) = \text{effect}(S_Z). \quad \square$$

We may finally prove Proposition 4.2:

PROPOSITION 4.2. *For every 2-VASS $\mathcal{V} = (Q, T)$, there exist $\mathbb{O} = [D, \infty)^2$ and finite sets S, S' of LPSs such that $D \leq (|Q| + \|T\|)^{O(1)}$, and*

- (a) $q(\mathbf{u}) \xrightarrow{*}_{\mathbb{N}^2} q(\mathbf{v})$ if and only if $q(\mathbf{u}) \xrightarrow{\cup S}_{\mathbb{N}^2} q(\mathbf{v})$ for every $q \in Q$, $\mathbf{u}, \mathbf{v} \in \mathbb{O}$, and $|\Lambda| \leq (|Q| + \|T\|)^{O(1)}$ and $|\Lambda|_* \leq 2$ for every $\Lambda \in S$;

$$(b) \xrightarrow{*}_{\mathbb{O}} \subseteq \xrightarrow{\cup S'}_{\mathbb{N}^2}, \text{ and } |\Lambda| \leq (|Q| + \|T\|)^{O(1)} \text{ and } |\Lambda|_* \leq 2 \cdot |Q| \text{ for every } \Lambda \in S'.$$

PROOF. Let $\mathcal{V} = (Q, T)$ be a 2-VASS, and let R be the set of linear path schemes obtained from Proposition 4.11 for \mathcal{V} . For every $q \in Q$ and every $\sigma \in R$ from q to q , let R_σ be the set of zigzag free linear path schemes obtained from Lemma 4.15 for q and σ . We claim that the following sets S, S' and value D satisfy the proposition:

$$\begin{aligned} S &\stackrel{\text{def}}{=} \bigcup_{q \in Q} \bigcup_{\substack{\sigma \in R \\ \sigma \text{ is from } q \text{ to } q}} R_\sigma, \\ S' &\stackrel{\text{def}}{=} \bigcup_{0 \leq k \leq |Q|} \bigcup_{\substack{\alpha_0, \alpha_1, \dots, \alpha_k \in T^* \\ \Lambda_1, \Lambda_2, \dots, \Lambda_k \in S}} \{ \alpha_0 \Lambda_1 \alpha_1 \cdots \Lambda_k \alpha_k \text{ is an LPS} : |\alpha_i| \leq |Q| \text{ for every } 0 \leq i \leq k \}, \\ D &\stackrel{\text{def}}{=} \max\{|\Lambda| : \Lambda \in S\} \cdot \|T\|. \end{aligned}$$

Proof of (a). Let us first show that $q(\mathbf{u}) \xrightarrow{*}_{\mathbb{N}^2} q(\mathbf{v}) \iff q(\mathbf{u}) \xrightarrow{\cup S}_{\mathbb{N}^2} q(\mathbf{v})$ for every $q(\mathbf{u}), q(\mathbf{v}) \in Q \times [D, \infty)^2$. Clearly, the left implication holds, hence we prove the right implication. Let $\mathbf{u}, \mathbf{v} \in [D, \infty)^2$ be such that $q(\mathbf{u}) \xrightarrow{*}_{\mathbb{N}^2} q(\mathbf{v})$. In particular, we have $q(\mathbf{u}) \xrightarrow{*}_{\mathbb{Z}^2} q(\mathbf{v})$. By Proposition 4.11, there exists $\sigma \in R$ such that $q(\mathbf{u}) \xrightarrow{\sigma}_{\mathbb{Z}^2} q(\mathbf{v})$. By Lemma 4.15, $\text{effect}(\sigma) \subseteq \text{effect}(R_\sigma)$. Therefore, $q(\mathbf{u}) \xrightarrow{\Lambda}_{\mathbb{Z}^2} q(\mathbf{v})$ for some $\Lambda \in R_\sigma$. Since Λ is zigzag-free, Lemma 4.7 and the choice of D imply that $q(\mathbf{u}) \xrightarrow{\Lambda}_{\mathbb{N}^2} q(\mathbf{v})$.

It remains to show that S satisfies the required bounds. Let $\Lambda \in S$. There exists $\sigma \in R$ such that $\Lambda \in R_\sigma$. By Lemma 4.15, $|\Lambda|_* \leq 2$. Moreover,

$$\begin{aligned} |\Lambda| &\leq (|\sigma| + \|T\|)^{O(1)} && \text{(by Lemma 4.15)} \\ &\leq (|Q| \cdot (|Q| + |T|) + \|T\|)^{O(1)} && \text{(by Proposition 4.11)} \\ &\leq (|Q| + \|T\|)^{O(1)}. \end{aligned}$$

Proof of (b). Let $p(\mathbf{u}), q(\mathbf{v}) \in Q \times \mathbb{O}$ be such that $p(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{O}} q(\mathbf{v})$ for some $\pi \in T^*$. By a pigeonhole argument, there exist $0 \leq k \leq |Q|$, $q_1, q_2, \dots, q_k \in Q$, $\mathbf{u}_1, \mathbf{u}'_1, \mathbf{u}_2, \mathbf{u}'_2, \dots, \mathbf{u}_k, \mathbf{u}'_k \in \mathbb{O}$, $\alpha_0, \alpha_1, \dots, \alpha_k \in T^*$ and $\beta_1, \beta_2, \dots, \beta_k \in T^*$ such that $\pi = \alpha_0 \beta_1 \alpha_1 \cdots \beta_k \alpha_k$, $|\alpha_i| \leq |Q|$ for every $i \in [k]$, and

$$p(\mathbf{u}) \xrightarrow{\alpha_0}_{\mathbb{O}} q_1(\mathbf{u}_1) \xrightarrow{\beta_1}_{\mathbb{O}} q_1(\mathbf{u}'_1) \xrightarrow{\alpha_1}_{\mathbb{O}} q_2(\mathbf{u}_2) \cdots q_k(\mathbf{u}_k) \xrightarrow{\beta_k}_{\mathbb{O}} q_k(\mathbf{u}'_k) \xrightarrow{\alpha_k}_{\mathbb{O}} q(\mathbf{v}).$$

Therefore, for every $i \in [k]$, there exists some linear path scheme $\Lambda_i \in S$ such that $q_i(\mathbf{u}_i) \xrightarrow{\Lambda_i}_{\mathbb{N}^2} q_i(\mathbf{u}'_i)$. Thus, $p(\mathbf{u}) \xrightarrow{\alpha_0 \Lambda_1 \alpha_1 \cdots \Lambda_k \alpha_k}_{\mathbb{O}} q(\mathbf{v})$ which implies that $p(\mathbf{u}) \xrightarrow{\cup S'}_{\mathbb{O}} q(\mathbf{v})$.

It remains to bound the size of the linear path schemes of S' . Let $\Lambda \in S'$. We have

$$\begin{aligned} |\Lambda| &\leq (|Q| + 1) \cdot |Q| + |Q| \cdot \max\{|\Lambda| : \Lambda \in S\} \\ &\leq (|Q| + 1) \cdot |Q| + |Q| \cdot (|Q| + \|T\|)^{O(1)} && \text{(by (a))} \\ &\leq (|Q| + \|T\|)^{O(1)} \end{aligned}$$

and

$$\begin{aligned} |\Lambda|_* &\leq |Q| \cdot \max\{|\Lambda|_* : \Lambda \in S\} \\ &\leq |Q| \cdot 2. && \text{(by (a))} \end{aligned}$$

□




Fig. 10. Overview of the proof of Theorem 3.2. Each node labeled by x corresponds to Proposition, Theorem, Lemma or Corollary x . An arrow from node u to node v indicates that u is used in the proof of v . Each colored region corresponds to a theme which is depicted under the same color in the general overview of Figure 2. The order in which the three themes are presented are numbered from (i) to (iii).

5 SHORT REACHABILITY WITNESSES: THEOREM 3.2

This section is dedicated to the proof of our second main result, which we recall below.

THEOREM 3.2. *For every 2-VASS \mathcal{V} , if $p(\mathbf{u}) \xrightarrow{*} q(\mathbf{v})$, then $p(\mathbf{u}) \xrightarrow{\pi} q(\mathbf{v})$ for a path π of length $(|\mathcal{V}|_1 + \|\mathbf{u}\| + \|\mathbf{v}\|)^{O(1)}$.*

By Theorem 3.1, it suffices to consider only linear path schemes instead of arbitrary 2-VASS. To further simplify we consider here the subclass of *simple* linear path schemes: a LPS $\Lambda = \alpha_0 \beta_1^* \alpha_1 \beta_2^* \cdots \beta_k^* \alpha_k$ is *simple* (an SLPS) if all α_i and cycles β_i have length 1. The main insight we prove in this section is that shortest reachability witnesses in an SLPS will not leave a polynomially bounded area around the source and target points.

THEOREM 5.1. *For every SLPS Λ from state s to state t , there exists $B \leq (|\Lambda| \cdot \|\Lambda\|)^{O(1)}$ such that if $\pi \in \Lambda$ is minimal with $s(\mathbf{0}) \xrightarrow{\pi}_{\mathbb{N}^2} t(\mathbf{0})$ then $s(\mathbf{0}) \xrightarrow{\pi}_{(\mathbb{N}_{\leq B} \times \mathbb{N}_{\leq B})} t(\mathbf{0})$.*

The next subsection contains several preparatory lemmas and afterwards, in Section 5.2, we prove Theorem 5.1. In the statement of Theorem 5.1, like in that of Theorem 3.2, we write $O(1)$ instead of providing a constant exponent explicitly. However, most of the auxiliary results that we shall state and prove towards obtaining the theorem will feature specific multiplicative constants and exponents. We do not claim those to be optimal, but we display them explicitly in order to clarify the dependencies on the various parameters (such as the length versus the norm) and to make comparisons between the various bounds easier.

The structure of the proof of Theorem 3.2 is depicted in Figure 10. Let us first start by observing how to derive Theorem 3.2 from Theorems 3.1 and 5.1. We will use the following simple fact (recall the notion of admissibility from Section 2).

LEMMA 5.2. Suppose that π is a cycle from state s to s in a 2-VASS $\mathcal{V} = (Q, T)$ and let \mathcal{V}' be the 2-VASS resulting from \mathcal{V} by adding one new transition: $t \stackrel{\text{def}}{=} (s, \text{effect}(\pi), s)$. Then for all $m \in \mathbb{N}$ and $\mathbf{v} \in \mathbb{N}^2$, the path π^{m+2} is admissible from \mathbf{v} if, and only if, $\pi t^m \pi$ is admissible from \mathbf{v} .

PROOF. The ‘only if’ direction follows by observing that the set of all points visited by $\pi t^m \pi$ from \mathbf{v} is included in the set of all points visited by π^{m+2} from \mathbf{v} .

For the ‘if’ direction, suppose \mathbf{w}' is a point that is visited by π^{m+2} but not by $\pi t^m \pi$ from \mathbf{v} . We have that \mathbf{w}' is reached by a path of the form $\pi^i \rho$ from \mathbf{v} , where $1 \leq i \leq m$ and ρ is a prefix of π . Let $\mathbf{w}^{(1)}$ and $\mathbf{w}^{(m+2)}$ be the points reached by the paths ρ and $\pi t^m \rho$ (respectively) from \mathbf{v} , which are in \mathbb{N}^2 since we are assuming that $\pi t^m \pi$ is admissible from \mathbf{v} . Observing that $\mathbf{w}' = \mathbf{w}^{(1)} + i \cdot \text{effect}(\pi)$ and $\mathbf{w}^{(m+2)} = \mathbf{w}^{(1)} + (m+1) \cdot \text{effect}(\pi)$, we conclude that \mathbf{w}' is on the line segment from $\mathbf{w}^{(1)}$ to $\mathbf{w}^{(m+2)}$ and therefore also in \mathbb{N}^2 . \square

LEMMA 5.3. Let \mathcal{V} be a 2-VASS, $\Lambda = \alpha_0 \beta_1^* \alpha_1 \beta_2^* \cdots \beta_k^* \alpha_k$ be a LPS and \mathcal{V}' be the 2-VASS resulting from \mathcal{V} by adding k extra summary transitions t_i , one for each cycle β_i , as in the previous lemma, and by adding a one-step cycle with effect $\mathbf{0}$ at each state. There is a finite set S of SLPSs over \mathcal{V}' such that $\xrightarrow{\Lambda} = \xrightarrow{\bigcup S}$ and for all $\Lambda' \in S$ the following hold.

(1) $|\Lambda'| \leq 4|\Lambda|$ and $\|\Lambda'\| \leq \|\Lambda\| \cdot |\Lambda|$.

(2) For every path $\pi' \in \Lambda'$ there exists $\pi \in \Lambda$ with $|\pi| \leq |\pi'| \cdot |\Lambda|$ and $\xrightarrow{\pi'} = \xrightarrow{\pi}$.

PROOF. The idea is first to split Λ into a finite set S of LPSs such that each of them predetermines, for every cycle of Λ , if it can be used zero, one or more than one times. This is done by replacing expressions β_i^* by the empty word, β_i , or $\beta_i \beta_i^* \beta_i$, respectively. Clearly, $\bigcup S = \Lambda$. In each such LPS Λ' , we then replace occurrences of expressions $\beta_i \beta_i^* \beta_i$ by expressions $\beta_i t_i^* \beta_i$. This results in $|\Lambda'| \leq |\Lambda| + \left(\sum_{i=1}^k |\beta_i|\right) + k$ and $\|\Lambda'\| \leq \|\Lambda\| \cdot |\Lambda|$. By Lemma 5.2, this last replacement does not change the relation $\xrightarrow{\Lambda'}$ and guarantees the second claimed property (where π and π' correspond to the unfolded and summarized paths in \mathcal{V} and \mathcal{V}' , respectively). It remains to introduce one-step cycles with effect $\mathbf{0}$ into the unstarred segments of Λ' to make the LPS simple, which is required at most $|\Lambda| + \left(\sum_{i=1}^k |\beta_i|\right) - k$ times. This results in $|\Lambda'| \leq 2|\Lambda| + 2 \sum_{i=1}^k |\beta_i| \leq 4|\Lambda|$, no change to $\|\Lambda'\|$, and no change to the relation $\xrightarrow{\Lambda'}$. \square

PROOF OF THEOREM 3.2. To reduce to the case where the initial and final points are the origin, let $\mathcal{V}' = (Q', T')$ be the 2-VASS that extends \mathcal{V} by two new states, say p', q' , as well as transitions $p' \xrightarrow{u} p$ and $q \xrightarrow{v} q'$. Observe that:

$$\begin{aligned} |\mathcal{V}'|_1 &= |Q'| + 2|T'| \cdot \|T'\| + 1 = (|Q| + 2) + 2(|T| + 2) \cdot \max\{\|T\|, \|\mathbf{u}\|, \|\mathbf{v}\|\} + 1 \\ &\leq 2(|T| + 2)(|\mathcal{V}|_1 + \|\mathbf{u}\| + \|\mathbf{v}\|) \\ &\leq (|\mathcal{V}|_1 + \|\mathbf{u}\| + \|\mathbf{v}\|)^4; \text{ where} \end{aligned}$$

$$\|T'\| = \max\{\|T\|, \|\mathbf{u}\|, \|\mathbf{v}\|\} \leq |\mathcal{V}|_1 + \|\mathbf{u}\| + \|\mathbf{v}\|.$$

Assuming that $p(\mathbf{u}) \xrightarrow{*}_{\mathbb{N}^2} q(\mathbf{v})$, we have that \mathcal{V}' admits a run of the form $p'(0) \xrightarrow{*}_{\mathbb{N}^2} \xrightarrow{v} q'(0)$. By Theorem 3.1, there exists an LPS Λ over \mathcal{V}' such that $p'(0) \xrightarrow{\Lambda} q'(0)$ and $|\Lambda| \in |\mathcal{V}'|_1^{O(1)}$.

Let \mathcal{V}'' be the 2-VASS obtained from \mathcal{V}' and Λ as in Lemma 5.3, by which there exists an SLPS Λ' such that: $p'(0) \xrightarrow{\Lambda'} q'(0)$, $|\Lambda'| \leq 4|\Lambda|$, $\|\Lambda'\| \leq \|\Lambda\| \cdot |\Lambda|$, and for every path $\pi' \in \Lambda'$ there exists $\pi \in \Lambda$ with $|\pi| \leq |\pi'| \cdot |\Lambda|$ and $\xrightarrow{\pi'} = \xrightarrow{\pi}$.

Applying Theorem 5.1, we get $B \leq (|\Lambda'| \cdot \|\Lambda'\|)^{O(1)}$ and a path $\pi' \in \Lambda'$ such that $p'(\mathbf{0}) \xrightarrow{\pi'}_{(\mathbb{N}_{\leq B} \times \mathbb{N}_{\leq B})} q'(\mathbf{0})$. Hence every cycle in Λ' whose effect is non-zero can be repeated in π' at most B times, so we may assume that $|\pi'| \leq B \cdot |\Lambda'|$.

Let $\pi \in \Lambda$ be such that $|\pi| \leq |\pi'| \cdot |\Lambda|$ and $\xrightarrow{\pi'} = \xrightarrow{\pi}$. Then $p'(\mathbf{0}) \xrightarrow{\pi} q'(\mathbf{0})$ and

$$|\pi| \leq B \cdot |\Lambda'| \cdot |\Lambda| \leq (|\Lambda'| \cdot \|\Lambda'\|)^c \cdot |\Lambda'| \cdot |\Lambda| \leq (4|\Lambda|^2)^{c+1} \cdot \|\Lambda\|^c \leq |\Lambda|^{4(c+1)} \cdot \|\Lambda\|^c$$

for a constant c . Since $|\Lambda| \leq |\mathcal{V}'|_1^d$ for a constant d , and $\|\Lambda\| \leq \|T'\|$, we conclude that

$$|\pi| \leq |\mathcal{V}'|_1^{4(c+1)d} \cdot \|T'\|^c \leq (|\mathcal{V}|_1 + \|\mathbf{u}\| + \|\mathbf{v}\|)^{16(c+1)d+c}.$$

It remains to recall that the path π^\dagger obtained by removing the first and last transitions from π is a path of \mathcal{V} such that $p(\mathbf{u}) \xrightarrow{\pi^\dagger} q(\mathbf{v})$. \square

5.1 Lemmas On Cones and Linear Path Schemes

We present six lemmas here that are useful in the sequel.

We start with four lemmas about cones. They are simple consequences in the plane of Cramer's Rule and Farkas-Minkowski-Weyl's Theorem, but for completeness we provide their proofs.

From now on, unless otherwise stated, the cone of a subset C of \mathbb{Q}^2 , denoted by $\text{cone}(C)$, will mean the $\mathbb{Q}_{>0}$ -cone of C , i.e. the closure of C under addition and under multiplication by positive rationals. We remark that the cone of C contains the zero vector if and only if it contains a line or one of the vectors in C is zero.

We shall use subscripts 1 and 2 to denote the coordinates of two-dimensional vectors.

Our first lemma is obtained from Cramer's Rule: it states that, for cones that contain the zero vector, the latter is expressible using at most three vectors from the spanning set, moreover with small positive coefficients.

LEMMA 5.4. *Let C be a finite subset of \mathbb{Z}^2 . If the cone of C contains $\mathbf{0}$, then $\mathbf{0}$ is a non-empty linear combination of at most three vectors from C and with coefficients in $\{1, \dots, 2\|C\|^2\}$. Furthermore, if $\mathbf{0}$ cannot be expressed like this with fewer than three vectors, then the cone of C is equal to \mathbb{Q}^2 .*

PROOF. If C contains $\mathbf{0}$, the statement is trivial. If C contains a vector \mathbf{a} with a negative coordinate a_i as well as a vector $\mathbf{b} = -\lambda\mathbf{a}$ for some positive rational λ , then $\mathbf{0}$ can be expressed as $b_i\mathbf{a} - a_i\mathbf{b}$ and we are done. So now assume that C does not contain vectors \mathbf{a} and \mathbf{b} like this.

Consider a minimal subset $C' \subseteq C$ such that $\mathbf{0}$ can be expressed as a linear combination $\lambda_1\mathbf{a}^{(1)} + \dots + \lambda_{|C'|}\mathbf{a}^{(|C'|)}$ with positive rational coefficients λ_i of vectors $\mathbf{a}^{(i)} \in C'$. Assume for contradiction that $|C'| > 3$. Then, by Farkas-Minkowski-Weyl's Theorem, there must be a closed half-plane containing at least 3 vectors, say w.l.o.g. $\mathbf{a}^{(1)}$, $\mathbf{a}^{(2)}$, and $\mathbf{a}^{(3)}$, from C' . One of these three vectors can be expressed as a non-negative linear combination of the other two. Without loss of generality assume $\mathbf{a}^{(1)} = c_1\mathbf{a}^{(2)} + c_2\mathbf{a}^{(3)}$ with $c_1, c_2 \geq 0$. But then we can write

$$\mathbf{0} = \lambda_1(c_1\mathbf{a}^{(2)} + c_2\mathbf{a}^{(3)}) + \lambda_2\mathbf{a}^{(2)} + \lambda_3\mathbf{a}^{(3)} + \dots + \lambda_{|C'|}\mathbf{a}^{(|C'|)}$$

and express $\mathbf{0}$ as a linear combination with positive coefficients of only $|C'| - 1$ vectors contradicting the minimality of C' .

Therefore, recalling the first paragraph, we can choose three vectors $\mathbf{a}, \mathbf{b}, \mathbf{c} \in C$ such that there are strictly positive x_1, x_2, x_3 and $x_1\mathbf{a} + x_2\mathbf{b} + x_3\mathbf{c} = \mathbf{0}$.

The equation has infinitely many solutions since we can scale the coefficients. However, if we set x_3 to be, say, $|\mathbf{b}_1\mathbf{a}_2 - \mathbf{a}_1\mathbf{b}_2|$ the solution becomes unique (since \mathbf{a} and \mathbf{b} are linearly independent) and it can be easily checked that the solution obtained by Cramer's rule is $x_1 = |\mathbf{c}_1\mathbf{b}_2 - \mathbf{b}_1\mathbf{c}_2|$ and $x_2 = |\mathbf{a}_1\mathbf{c}_2 - \mathbf{c}_1\mathbf{a}_2|$.

For the second statement of the lemma observe that we can express $-a$ and $-b$ as linear combinations of a , b , and c with positive rationals. For example, $-a = (x_2b + x_3c)/x_1$. Since a and b are linearly independent, any vector in \mathbb{Q}^2 can be expressed as a linear combination of a and b using rational coefficients. Combined with the fact that we can express $-a$ and $-b$ the claim follows. \square

The next two lemmas apply to the other case, i.e. when the cone does not contain the zero vector: firstly, such cones are determined by pairs of outermost vectors in their spanning sets; and secondly, they are contained in open half-planes determined by small vectors.

For a vector $v = (v_1, v_2) \in \mathbb{Z}^2$ let us write $\text{right}(v) \stackrel{\text{def}}{=} (v_2, -v_1)$ and $\text{left}(v) \stackrel{\text{def}}{=} (-v_2, v_1)$ for the vectors derived by rotating v by 90° clockwise and anticlockwise, respectively.

LEMMA 5.5. *Let $C \subseteq \mathbb{Z}^2$ be non-empty and $0 \notin \text{cone}(C)$. There are two vectors $a, b \in C$ such that*

- $\text{cone}(\{a, b\}) = \text{cone}(C)$,
- for all $c \in \text{cone}(C)$, $\text{left}(a) \cdot c \geq 0$ (i.e. c is between six and twelve o'clock in relation to a), and
- for all $c \in \text{cone}(C)$, $\text{right}(b) \cdot c \geq 0$ (i.e. c is between zero and six o'clock in relation to b).

PROOF. Consider a subset $C' \subseteq C$ of minimum size that spans the same cone as C . Assume for contradiction that $|C'| > 2$. Then the set contains three vectors x, y , and z and because these vectors must be linearly dependent we have $z = \lambda_1 x + \lambda_2 y$ for some rationals λ_1 and λ_2 . We can assume, without loss of generality, that λ_1 and λ_2 do not have different signs (otherwise we can appropriately rename x, y , and z). If λ_1 and λ_2 are non-negative, $C' \setminus \{z\}$ still spans the same cone as C' since in any positive combination, z can be replaced by $\lambda_1 x + \lambda_2 y$. If however, λ_1 and λ_2 are non-positive, the cone spanned by C' contains 0 since $0 = z - \lambda_1 x - \lambda_2 y$. In both cases we get a contradiction to our assumptions.

So there must indeed be two vectors $a, b \in C$, not necessarily different, that span the same cone as C . Observe that $b \cdot \text{left}(a) < 0 \iff b \cdot \text{right}(a) > 0$ because $\text{right}(a) = -\text{left}(a)$. Further observe that $b \cdot \text{right}(a) = \text{left}(b) \cdot a$. Therefore, either $b \cdot \text{left}(a) \geq 0$ or $b \cdot \text{right}(a) = \text{left}(b) \cdot a \geq 0$ holds. We assume w.l.o.g. that $b \cdot \text{left}(a) \geq 0$, since otherwise we can swap the names of a and b .

Pick any $c \in C$. Since the cone of $\{a, b\}$ contains c , there exist $x, y \geq 0$ such that $c = xa + yb$ and therefore

$$c \cdot \text{left}(a) = (xa + yb) \cdot \text{left}(a) = xa \cdot \text{left}(a) + yb \cdot \text{left}(a) \geq 0,$$

because $xa \cdot \text{left}(a) = 0$. Analogously, using $yb \cdot \text{right}(b) = 0$, we get

$$c \cdot \text{right}(b) = (xa + yb) \cdot \text{right}(b) = xa \cdot \text{right}(b) + yb \cdot \text{right}(b) \geq 0. \quad \square$$

LEMMA 5.6. *Let $C \subset \mathbb{Z}^2$ be non-empty such that $0 \notin \text{cone}(C)$. There exists a vector $p \in \mathbb{Z}^2$ such that $\|p\| \leq 2\|C\|$ and $p \cdot c > 0$ for all $c \in C$.*

PROOF. According to Lemma 5.5 we have vectors $a, b \in C$ such that $\{a, b\}$ spans the same cone as C , and for all $c \in C$, $\text{left}(a) \cdot c \geq 0$ and $\text{right}(b) \cdot c \geq 0$.

If a alone already spans the same cone as C , we can choose $p = a$ and are done. Otherwise, a and b are linearly independent and we choose $p = \text{left}(a) + \text{right}(b)$. Clearly $\|p\| \leq 2\|C\|$. For any $c \in C$, $p \cdot c = (\text{left}(a) + \text{right}(b)) \cdot c$. Since a and b are linearly independent, $\text{left}(a) \cdot c \neq 0$ or $\text{right}(b) \cdot c \neq 0$ and therefore $(\text{left}(a) + \text{right}(b)) \cdot c > 0$. \square

Our last lemma dealing with cones gives some additional properties for the structure of the cones when it is known that the cone does not contain some vector. For simplicity, and because it is all we will need later, we focus on the case that $(0, 1)$ is not contained in the cone.

LEMMA 5.7. *Let $C \subset^{\text{fin}} \mathbb{Z}^2$ be a finite, non-empty set with $0 \notin C$ and such that $(0, 1) \notin \text{cone}(C)$. There is a vector $p \in \mathbb{Z}^2$ such that*

- $\|p\| \leq \|C\|$,
- $p \cdot (0, 1) < 0$,
- $p \cdot c \geq 0$ for all $c \in C$, and
- if $p_1 < 0$, then $\text{right}(p) \in C$.

PROOF. We distinguish two basic cases based on whether the cone of C contains 0 or not. First suppose the cone of C does not contain 0 . Then, by Lemma 5.5, there are vectors $a, b \in C$ such that $\{a, b\}$ spans the same cone as C , and for all $c \in C$, $\text{left}(a) \cdot c \geq 0$ and $\text{right}(b) \cdot c \geq 0$. Informally, we have that the cone of C consists of all rational points in the closed acute angle from a anti-clockwise to b except the origin. We now consider the following three subcases.

$a_1 < 0$ and $a_2 > 0$, i.e. a is in the strict top-left quadrant: It is straightforward to check that $p = \text{left}(a)$ satisfies the conditions of the lemma.

$b_1 > 0$ and $b_2 > 0$, i.e. b is in the strict top-right quadrant: Similarly, $p = \text{right}(b)$ satisfies the conditions of the lemma.

otherwise: Here suppose $a_2 > 0$. Since we are not in the first subcase and $(0, 1) \notin \text{cone}(C)$, it must be that $a_1 > 0$. Further, since $b \in C$, we have that $0 \leq \text{left}(a) \cdot b = a_1 b_2 - a_2 b_1$. From that inequality, not being in the second subcase, and $(0, 1) \notin \text{cone}(C)$ again, it follows that $b_1 < 0$. But then $(0, a_1 b_2 - a_2 b_1) = a_1 b - b_1 a$ is in $\text{cone}(\{a, b\}) = \text{cone}(C)$ which does not contain 0 , so $a_1 b_2 - a_2 b_1 > 0$, and hence also $(0, 1) \in \text{cone}(C)$, which is a contradiction.

Thus $a_2 \leq 0$, and by a symmetric argument, $b_2 \leq 0$ as well. It is now easy to see that $p = (0, -1)$ satisfies the conditions of the lemma.

We now move to the second case in which we assume that the cone of C does contain 0 . Then there are two vectors $a, b \in C$ such that $a + \lambda b = 0$ for some positive rational λ . This is because if 0 could only be expressed with three or more vectors, according to Lemma 5.4, $(0, 1)$ would also be in the cone of C .

Observe that either $a_1 < 0$ or $b_1 < 0$, since otherwise we would need to have $a_1 = 0$ and $b_1 = 0$, which would imply that $(0, 1) \in \text{cone}(C)$. Without loss of generality let $a_1 < 0$. We choose $p = \text{left}(a)$.

The only condition of the lemma not trivially met is that $p \cdot c \geq 0$ for all $c \in C$. Assume that there is a $c \in C$ such that $\text{left}(a) \cdot c < 0$. Then $c_1 a_2 - a_1 c_2 > 0$. If $c_1 \geq 0$, $(0, 1)$ would be in the cone of C since it can be expressed as $(c_1 \cdot a - a_1 \cdot c) / (c_1 a_2 - a_1 c_2)$. Otherwise $(0, 1)$ would also be in the cone of C since it can be expressed as $(b_1 \cdot c - c_1 \cdot b) / (b_1 c_2 - c_1 b_2)$, where $b_1 = -\lambda a_1 > 0$ and $b_1 c_2 - c_1 b_2 = \lambda(c_1 a_2 - a_1 c_2) > 0$. Either way, we have a contradiction. \square

Moving from rational cones to paths of SLPSs, our remaining two lemmas pin down some relatively basic properties of SLPS paths in which some cycles are repeated ‘many’ times: firstly, if all those cycles are contained in a half-plane, then the effect of the path must point roughly in the same direction (we have a strict and a non-strict version here, cf. Figure 11); secondly, if the path when started at a point remains sufficiently far from both axes (i.e. respects a sufficiently wide margin), then it can be shortened admissibly by a range of multiples of any small vector that is in the cone spanned by the ‘often’ repeated cycles.

Observe that it is possible that the same path through a linear path scheme $\alpha_0 \beta_1^* \alpha_1 \beta_2^* \dots \beta_k^* \alpha_k$ can be written as $\alpha_0 \beta_1^{n_1} \alpha_1 \beta_2^{n_2} \dots \beta_k^{n_k} \alpha_k$ for several different tuples of naturals n_1, n_2, \dots, n_k . To avoid ambiguity, from now on, we assume that every such path is not simply a sequence of transitions, but is given by some such tuple of naturals.

Definition 5.8. Let $\Lambda = \alpha_0 \beta_1^* \alpha_1 \beta_2^* \dots \beta_k^* \alpha_k$ be an SLPS and $\pi \in \Lambda$ a path. We write

$$\text{cycles}_{\geq B}(\Lambda, \pi) \subseteq \{\beta_1, \beta_2, \dots, \beta_k\}$$

for the set of all cycles of Λ that are repeated in π at least B times.

LEMMA 5.9. Suppose that Λ is a simple linear path scheme with k cycles and $\pi \in \Lambda$ a path. For all $B \in \mathbb{N}$ and $\mathbf{p} \in \mathbb{Z}^2$:

- (1) if $\mathbf{p} \cdot \mathbf{a} > 0$ for all $\mathbf{a} \in \text{cycles}_{\geq B}(\Lambda, \pi)$, then $\mathbf{p} \cdot \text{effect}(\pi) \geq |\pi| - (k \cdot B + 1)(2\|\Lambda\| \|\mathbf{p}\| + 1)$;
- (2) if $\mathbf{p} \cdot \mathbf{a} \geq 0$ for all $\mathbf{a} \in \text{cycles}_{\geq B}(\Lambda, \pi)$, then $\mathbf{p} \cdot \text{effect}(\pi) \geq -(kB + 1)(2\|\Lambda\| \|\mathbf{p}\|)$.

PROOF. The effect of π can be decomposed as $\text{effect}(\pi) = \mathbf{v} + \mathbf{b}$, where \mathbf{v} is the combined effect of those cycles occurring at least B times and \mathbf{b} is the rest. Hence \mathbf{v} is a linear combination $\mathbf{v} = \sum_{i=1}^{\ell} \mathbf{a}^{(i)}$, where $\mathbf{a}^{(i)} \in \text{cycles}_{\geq B}(\Lambda, \pi)$ and \mathbf{b} is the effect of a path of length $|\pi| - \ell \leq k(B-1) + (k+1) = kB + 1$. We can therefore estimate

$$\mathbf{p} \cdot \mathbf{b} \geq -2(kB + 1)\|\Lambda\| \|\mathbf{p}\|. \quad (9)$$

If $\mathbf{p} \cdot \mathbf{a} > 0$ for all $\mathbf{a} \in \text{cycles}_{\geq B}(\Lambda, \pi)$ then

$$\mathbf{p} \cdot \mathbf{v} = \sum_{i=1}^{\ell} \mathbf{p} \cdot \mathbf{a}^{(i)} \geq \ell \geq |\pi| - (kB + 1). \quad (10)$$

Claim one of the lemma now follows by Equations (9) and (10), and because $\mathbf{p} \cdot \text{effect}(\pi) = \mathbf{p} \cdot \mathbf{v} + \mathbf{p} \cdot \mathbf{b}$.

For claim two observe that if $\mathbf{p} \cdot \mathbf{a} \geq 0$ for all $\mathbf{a} \in \text{cycles}_{\geq B}(\Lambda, \pi)$, then $\mathbf{p} \cdot \mathbf{v} = \sum_{i=1}^{\ell} \mathbf{p} \cdot \mathbf{a}^{(i)} \geq 0$. \square

Definition 5.10. Let $\pi = \alpha_0 \beta_1^{n_1} \alpha_1 \beta_2^{n_2} \cdots \beta_k^{n_k} \alpha_k$ and $\pi' = \alpha_0 \beta_1^{n'_1} \alpha_1 \beta_2^{n'_2} \cdots \beta_k^{n'_k} \alpha_k$ be paths of an SLPS Λ over the VASS (Q, T) , and $\mathbf{e} \in \mathbb{Z}^2$ some vector. Then π' is called a *shortening of π by \mathbf{e}* if $(n'_1, n'_2, \dots, n'_k) < (n_1, n_2, \dots, n_k)$ and $\text{effect}(\pi') = \text{effect}(\pi) - \mathbf{e}$.

LEMMA 5.11. Let Λ be an SLPS such that $\|\Lambda\| > 0$, $N \in \mathbb{N}$ and $\mathbb{D} \stackrel{\text{def}}{=} (\mathbb{N}_{\geq 6\|\Lambda\|^3 N})^2$. Suppose that $\pi \in \Lambda$ is a \mathbb{D} -run from some point \mathbf{s} , and $\mathbf{c} \in \mathbb{Z}^2$ is a vector such that

- $\|\mathbf{c}\| \leq \|\Lambda\|$ and
- the cone of cycles $_{\geq 2\|\Lambda\|^2 N}(\Lambda, \pi)$ contains \mathbf{c} .

Then there exists $\gamma \in \{1, \dots, 2\|\Lambda\|^2\}$ such that, for all $n \in \{1, \dots, N\}$, π has a shortening by $n\gamma\mathbf{c}$ which is admissible from \mathbf{s} .

PROOF. Let $C \stackrel{\text{def}}{=} \text{cycles}_{\geq 2\|\Lambda\|^2 N}(\Lambda, \pi)$. We claim that $\gamma\mathbf{c} = \lambda_1 \mathbf{a}^{(1)} + \dots + \lambda_j \mathbf{a}^{(j)}$ for some $j \in \{1, 2, 3\}$, $\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(j)} \in C$ and $\gamma, \lambda_1, \dots, \lambda_j \in \{1, \dots, 2\|\Lambda\|^2\}$. If $\mathbf{c} = \mathbf{0}$, this directly follows from Lemma 5.4. Otherwise, according to Lemma 5.5, there must be two vectors $\mathbf{a}^{(1)}, \mathbf{a}^{(2)} \in C$ such that the cone spanned by $\{\mathbf{a}^{(1)}, \mathbf{a}^{(2)}\}$ contains \mathbf{c} but not $\mathbf{0}$. Then the claim follows by Lemma 5.4 applied to the set $\{-\mathbf{c}, \mathbf{a}^{(1)}, \mathbf{a}^{(2)}\}$.

Now, we can subtract $n\gamma\mathbf{c}$ from the effect of π by deleting $n\lambda_i$ occurrences of the cycle $\mathbf{a}^{(i)}$ for all $i \in \{1, \dots, j\}$. Any such shortening π' is admissible from \mathbf{s} because, for any point visited by π , the differences between its coordinates and the coordinates of the corresponding point visited by π' are at most $6\|\Lambda\|^3 n$. \square

5.2 Reachability Witnesses in Simple Linear Path Schemes

Equipped with the lemmas from the previous subsection, we now focus on proving Theorem 5.1, which states that minimal runs of SLPSs from the origin to the origin only visit pseudo-polynomially bounded points.

We shall basically argue by contradiction, i.e. prove that every run from $\mathbf{0}$ to $\mathbf{0}$ that visits a point whose norm exceeds a pseudo-polynomial bound can be shortened while maintaining its initial and final points, and its admissibility.




Fig. 11. Illustration of Lemma 5.9: The vectors in $\text{cycles}_{\geq B}(\Lambda, \pi)$ are colored in magenta, the direction vector p is colored in blue. *Left: case (1).* If all cycles make strict progress in the direction of p , then so does the path π . *Right: case (2).* If in the direction of p no cycle effect is negative, then the effect of π can be lower-bounded (will end in the shaded area).

Most of the work is in obtaining the following four lemmas, which provide progressively more sophisticated tools for shortening runs whose final points have small first coordinates and large second coordinates, i.e. that have climbed up high to a point close to the vertical axis. The shortenings need to preserve admissibility, but their effect may be to move the final point vertically down by some amount.

In the first lemma below, Λ is to be viewed as a fragment of an SLPS that we are ultimately interested in, and M is a parameter which in applications of the lemma will be instantiated to relatively small pseudo-polynomial expressions. Like in most of the subsequent lemmas here, M is used as the width of a ‘margin’, i.e. points that have one or both coordinates less than M are thought of as ‘near’ one or both axes. In that terminology, the lemma states that every run that stays near the vertical axis but far from the horizontal axis, and climbs up a long way (by more than $(\|\Lambda\|_* M + 1)\|\Lambda\|$), can be shortened by a vertical vector $(0, n\gamma)$ (cf. Figure 12 on the left). Significantly, in the latter expression $(0, n\gamma)$, both γ and n have small bounds, but γ is quantified existentially followed by n quantified universally. The lemma thus provides a range of admissible shortenings of the run: all scalar multiples of some small vertical vector $(0, \gamma)$, up to some small bound for n .

The subsequent three lemmas will similarly provide ranges of shortenings that consist of all bounded scalar multiples of some bounded vertical vector. The reason is that, ultimately in the proof of Theorem 5.1, these lemmas will be applied both to a subrun that climbs up to a high point close to a vertical axis, and to the reversal of a continuing subrun that climbs down from the same high point. If the former subrun can be shortened by $(0, n\gamma)$ and the latter reversed subrun by $(0, n'\gamma')$, where the bounds on the universally quantified n and n' allow them to be set to the existentially quantified γ' and γ (respectively), then both can be shortened by $(0, \gamma\gamma')$, resulting in a composite shortening that moves the high point vertically down by $\gamma\gamma'$ but maintains the initial and final points of the overall run as well as its admissibility.

LEMMA 5.12. *Let Λ be an SLPS with k cycles, $M \geq \|\Lambda\|$ and $\mathbb{D} \stackrel{\text{def}}{=} \mathbb{N}_{<M} \times \mathbb{N}_{\geq M}$. Suppose that $\pi \in \Lambda$ is a \mathbb{D} -run from point s to point t such that $(t - s)_2 > (kM + 1)\|\Lambda\|$. Then there is a $\gamma \in \{1, \dots, \|\Lambda\|\}$ such that, for all $n \in \{1, \dots, \lfloor M/\gamma \rfloor\}$, π has a shortening by $(0, n\gamma)$ which is admissible from s .*

PROOF. There is a cycle c in π that is repeated at least M times. Otherwise, for the effect of π , $\|t - s\| \leq ((k + 1) + k \cdot (M - 1)) \cdot \|\Lambda\| = (kM + 1)\|\Lambda\|$, which contradicts the assumption of the lemma. Let u and v be the points visited right before the first, and right after the last repetitions of the cycle c , respectively. The first coordinate of c is 0 since otherwise $|(u - v)_1| \geq M$, which




Fig. 12. Illustrations of Lemma 5.12 (left) and Lemma 5.13 (right).

contradicts that we assumed a \mathbb{D} -run. Therefore $\mathbf{c} = (0, \gamma)$ for some $\gamma \in \{1, \dots, \|\Lambda\|\}$ and thus, π has a shortening by $(0, n\gamma)$ for all $n \in \{1, \dots, \lfloor M/\gamma \rfloor\}$. This shortening is admissible since it does not affect the first coordinate of any point visited, only decreases the second coordinates by at most $\lfloor M/\gamma \rfloor \cdot \gamma \leq M$, and all visited points have a second coordinate value of at least M prior to the shortening. \square

The second lemma is also about runs that climb up a long way. In contrast to the previous lemma, the run is assumed to remain far from both axes (by a margin that depends on the parameter N), and the shortenings identified in the proof will be provided by Lemma 5.11 and so in general involve more than one cycle. The assumption on the effect $\mathbf{t} - \mathbf{s}$ says that the climb is ‘long and almost vertical’, more precisely that it leads from \mathbf{s} to a point in the interior of the angle rooted at $\mathbf{s} + (0, (4\|\Lambda\|_*N + 2)\|\Lambda\|^4)$ and with rays determined relatively by $(1, \|\Lambda\|)$ and $(-1, \|\Lambda\|)$ (cf. Figure 12 on the right). We remark that restricting λ to range over $\{-\|\Lambda\|, \|\Lambda\|\}$ would not change the assumption.

LEMMA 5.13. *Let Λ be an SLPS with k cycles, $N \in \mathbb{N}$ and $\mathbb{D} \stackrel{\text{def}}{=} (\mathbb{N}_{\geq 6\|\Lambda\|^3N})^2$. Suppose that $\pi \in \Lambda$ is a \mathbb{D} -run from point \mathbf{s} to point \mathbf{t} such that for all $\lambda \in [-\|\Lambda\|, \|\Lambda\|]$,*

$$(\lambda, 1) \cdot (\mathbf{t} - \mathbf{s}) > (4kN + 2)\|\Lambda\|^4.$$

Then $(0, 1)$ is in the cone of cycles $_{\geq 2\|\Lambda\|^2N}(\Lambda, \pi)$ and there exists $\gamma \in \{1, \dots, 2\|\Lambda\|^2\}$ such that, for all $n \in \{1, \dots, N\}$, π has a shortening by $(0, n\gamma)$ which is admissible from \mathbf{s} .

PROOF. Note that, by taking $\lambda = 0$ in the last assumption of the lemma, we have in particular $(\mathbf{t} - \mathbf{s})_2 > 0$.

Let $C = \text{cycles}_{\geq 2\|\Lambda\|^2N}(\Lambda, \pi)$. Assume for contradiction that $(0, 1)$ is not in the cone of $C \setminus \{0\}$. Then, due to Lemma 5.7, there exists $\mathbf{p} \in \mathbb{Z}^2$ such that $\|\mathbf{p}\| \leq \|\Lambda\|$, $\mathbf{p} \cdot (0, 1) < 0$, and $\mathbf{p} \cdot \mathbf{a} \geq 0$ for all $\mathbf{a} \in C$. This implies $\mathbf{p}_2 < 0$ and therefore

$$-\mathbf{p} \cdot (\mathbf{t} - \mathbf{s}) \geq (-\mathbf{p}_1, 1) \cdot (\mathbf{t} - \mathbf{s}) > (4kN + 2)\|\Lambda\|^4.$$

But, by Lemma 5.9 (2),

$$\begin{aligned} \mathbf{p} \cdot (\mathbf{t} - \mathbf{s}) &\geq -(k2\|\Lambda\|^2N + 1)(2\|\Lambda\|^2) \\ &\geq -(4kN + 2)\|\Lambda\|^4. \end{aligned}$$

Therefore, $(0, 1)$ must be in the cone of C , and we conclude by Lemma 5.11. \square




Fig. 13. Illustrations of Lemma 5.14 (left) and Lemma 5.15 (right). The latter features cycle vectors \mathbf{v} and \mathbf{w} that occur in the proof and whose cone includes the unit vector $(0, 1)$.

The next lemma builds on the previous two, and considers runs that start near the horizontal axis, after the first transition remain far from the horizontal axis, and finish very high near the vertical axis. Their two parts are analysed, before and after visiting for the first time a point (t') near the vertical axis, and there are two cases (cf. Figure 13 on the left):

- If either the vertical distance between t' and the final point is large, or $(0, 1)$ is in the cone spanned by the cycles that are repeated many times on the way from the initial point to t' , then the lemma provides a bounded range of shortenings of the run by vertical vectors. The former subcase here is harder because the run may move out of and back into the vertical margin many times, but it is handled in the proof by arguing that either Lemma 5.12 or Lemma 5.13 is applicable to a segment.
- Otherwise, the lemma asserts that there must exist a cycle which is repeated many times before visiting t' , whose vector is in the interior of the upper left-hand quadrant, and whose angle (anti-clockwise from the positive horizontal axis) is not much greater than the angle of the effect of the run (i.e. the vector from the initial point to the final point).

LEMMA 5.14. *Suppose we have an SLPS Λ with $k > 0$ cycles and $N, M \in \mathbb{N}$ with $M \geq 6\|\Lambda\|^3 N$. Let $\pi \in \Lambda$ be a run from point s to point t such that*

- $s_1 \geq 0, s_2 < M$,
- *all points visited by π after s are in $(\mathbb{N} \times \mathbb{N}_{\geq M})$, and*
- $t_1 < M, t_2 \geq 12(k+1)(M+1)\|\Lambda\|^4$.

Let π' be the shortest nonempty prefix of π whose final point t' satisfies $t'_1 < M$. Suppose that $|\pi'| \geq 2$. Then let π^\dagger be π' without its first and last transitions, let Λ^\dagger be an SLPS one of whose paths is π^\dagger and whose length and norm are at most those of Λ , and let $C = \text{cycles}_{\geq 2\|\Lambda\|^2 N}(\Lambda^\dagger, \pi^\dagger)$.

- (1) *If $(t - t')_2 > 6(k+1)(M+1)\|\Lambda\|^4$ or $(0, 1)$ is in the cone of C , then there exists $\gamma \in \{1, \dots, 2\|\Lambda\|^2\}$ such that, for all $n \in \{1, \dots, N\}$, π has a shortening by $(0, n\gamma)$ which is admissible from s .*
- (2) *Otherwise, there exists $\mathbf{v} \in C \cap (\mathbb{Z}_{<0} \times \mathbb{Z}_{>0})$ such that $\text{left}(\mathbf{v}) \cdot (s_1, -t_2) < 7(k+2)(M+1)\|\Lambda\|^5$.*

PROOF. We note that, from the first two bullet points in the lemma, we have $M > 0$ and $\|\Lambda\| > 0$.

If $(t - t')_2 > 6(k+1)(M+1)\|\Lambda\|^4$, let π'' be the rest of π after π' , i.e., the portion of π that starts at t' and ends at t . Then partition π'' into segments that visit only points in $\mathbb{N}_{<M} \times \mathbb{N}$ and segments for which all intermediate points are outside that set. Call these segments y -axis-close

and y -axis-far, respectively. In the following we argue that either Lemma 5.12 applies to one of the former segments, or Lemma 5.13 applies to one of the latter segments.

Let ℓ be the total number of segments and, for $i \in [1, \ell - 1]$, let $\mathbf{a}^{(i)}$ be the endpoint of the i -th segment and the start point of the $(i + 1)$ -th segment. Note that a path from an SLPS with at most k cycles will be split into at most $2(k + 1)$ segments and therefore $\ell \leq 2(k + 1)$. For convenience, define $\mathbf{a}^{(0)}$ to be \mathbf{t}' and $\mathbf{a}^{(\ell)}$ to be \mathbf{t} .

Each segment corresponds to an SLPS that is a fragment of the original SLPS. Let the SLPS fragment of the i -th segment contain k_i cycles. Note that each of the cycles in the original SLPS can only be part of two different segments. Therefore, $\sum k_i \leq 2k$. Since

$$\begin{aligned} \sum_{i=1}^{\ell} (\mathbf{a}^{(i)} - \mathbf{a}^{(i-1)})_2 &= (\mathbf{t} - \mathbf{t}')_2 \\ &> 6(k + 1)(M + 1)\|\Lambda\|^4 \\ &= 4(k + 1)M\|\Lambda\|^4 + 2(k + 1)(M + 1)\|\Lambda\|^4 + 4(k + 1)\|\Lambda\|^4 \\ &> 2kM\|\Lambda\| + 2(k + 1)(M + 1)\|\Lambda\| + 4(k + 1)\|\Lambda\|^4, \end{aligned}$$

there must be a segment i , going from $\mathbf{a}^{(i-1)}$ to $\mathbf{a}^{(i)}$, for which

$$(\mathbf{a}^{(i)} - \mathbf{a}^{(i-1)})_2 > (k_i M + M + 1)\|\Lambda\| + 2\|\Lambda\|^4.$$

If this segment i is y -axis-close, we observe that $(\mathbf{a}^{(i)} - \mathbf{a}^{(i-1)})_2 > (k_i M + 1)\|\Lambda\|$ and therefore Lemma 5.12 applies to it. If this segment i is y -axis-far then

$$\begin{aligned} (\mathbf{a}^{(i)} - \mathbf{a}^{(i-1)})_2 &> (k_i 6N\|\Lambda\|^3 + M + 1)\|\Lambda\| + 2\|\Lambda\|^4 \\ &> (k_i 4N + 2)\|\Lambda\|^4 + 2\|\Lambda\| + M\|\Lambda\|, \end{aligned}$$

since $M \geq 6\|\Lambda\|^3 N$.

Now consider the point $\mathbf{a}^{(i-1)'}$ visited right after $\mathbf{a}^{(i-1)}$ and the point $\mathbf{a}^{(i)'}$ visited right before $\mathbf{a}^{(i)}$ and consider the path between $\mathbf{a}^{(i-1)}$ and $\mathbf{a}^{(i)}$ without the first and last vector. Note that $\mathbf{a}_1^{(i-1)'}, \mathbf{a}_1^{(i)'} \in [M, M + \|\Lambda\|)$ and hence $|(\mathbf{a}^{(i)'} - \mathbf{a}^{(i-1)'})_1| \leq \|\Lambda\| < M$. Therefore, we have that $(\lambda, 1) \cdot (\mathbf{a}^{(i)'} - \mathbf{a}^{(i-1)'}) > (k_i 4N + 2)\|\Lambda\|^4$ for all $\lambda \in [-\|\Lambda\|, \|\Lambda\|]$ and hence Lemma 5.13 applies to this subpath, going from $\mathbf{a}^{(i-1)'}$ to $\mathbf{a}^{(i)'}$.

Note that the section of π going from \mathbf{s} to $\mathbf{a}^{(i-1)}$ (or $\mathbf{a}^{(i-1)'}$, respectively) is still admissible after the shortening carried out through Lemma 5.12 or Lemma 5.13. The shortened segment i is also admissible due to these lemmas. The section of π that started at $\mathbf{a}^{(i)}$ prior to the shortening is also admissible since the first coordinate of the corresponding points is not changed and the second coordinate is decreased by at most $N2\|\Lambda\|^2 < M$. Moreover, the second coordinate of all the points prior to the shortening was at least M .

In the remainder of the proof, assume

$$(\mathbf{t} - \mathbf{t}')_2 \leq 6(k + 1)(M + 1)\|\Lambda\|^4$$

and consequently, since $\mathbf{t}_2 \geq 12(k + 1)(M + 1)\|\Lambda\|^4$, that

$$\mathbf{t}'_2 \geq 6(k + 1)(M + 1)\|\Lambda\|^4.$$

Then $|\pi'| \geq 2$, so π^\dagger , Λ^\dagger and C are well defined. Let \mathbf{s}^\dagger be the first point visited by π' after \mathbf{s} , and let \mathbf{t}^\dagger be the final point of π^\dagger from \mathbf{s}^\dagger . Observe that $\mathbf{s}_1^\dagger \geq 0$, $\mathbf{s}_2^\dagger < M + \|\Lambda\|$, $\mathbf{t}_1^\dagger < M + \|\Lambda\|$, and $\mathbf{t}_2^\dagger \geq 6(k + 1)(M + 1)\|\Lambda\|^4 - \|\Lambda\|$. In particular, $\mathbf{t}_2^\dagger - \mathbf{s}_2^\dagger > 5(k + 1)(M + 1)\|\Lambda\|^4$, so $C \setminus \{\mathbf{0}\}$ must be non-empty, since otherwise we would have $\mathbf{t}_2^\dagger - \mathbf{s}_2^\dagger \leq (2kN\|\Lambda\|^2 + 1)\|\Lambda\| < 2(k + 1)(M + 1)\|\Lambda\|^3$.

If $(0, 1)$ is in the cone of C , we are done by Lemma 5.11 applied to π^\dagger from s^\dagger , which visits only points in $(\mathbb{N}_{\geq M})^2$. Note that all points of π after s have a second coordinate of at least M . Therefore, the shortening due to Lemma 5.11 can also be applied to π and results in an admissible path from s .

If $(0, 1)$ is not in the cone of $C \setminus \{0\}$ then Lemma 5.7 provides a vector $\mathbf{v} \in \mathbb{Z}^2$ such that $\|\mathbf{v}\| \leq \|\Lambda\|$, $\text{left}(\mathbf{v}) \cdot (0, 1) < 0$, $\text{left}(\mathbf{v}) \cdot \mathbf{a} \geq 0$ for all $\mathbf{a} \in C$, and such that $\mathbf{v}_2 > 0$ implies $\mathbf{v} \in C$. Hence, $\mathbf{v}_1 < 0$ and Lemma 5.9 (2) gives us

$$\text{left}(\mathbf{v}) \cdot (\mathbf{t}^\dagger - \mathbf{s}^\dagger) \geq -(2kN\|\Lambda\|^2 + 1)(2\|\Lambda\|^2) \geq -2k(2N + 1)\|\Lambda\|^4. \quad (11)$$

But then $\mathbf{v}_2 > 0$, since the contrary would contradict Equation 11 as follows:

$$\begin{aligned} \text{left}(\mathbf{v}) \cdot (\mathbf{t}^\dagger - \mathbf{s}^\dagger) &= -\mathbf{v}_2(\mathbf{t}_1^\dagger - \mathbf{s}_1^\dagger) + \mathbf{v}_1(\mathbf{t}_2^\dagger - \mathbf{s}_2^\dagger) \\ &< -\mathbf{v}_2(M + \|\Lambda\|) + \mathbf{v}_1(6(k + 1)(M + 1)\|\Lambda\|^4 - M - 2\|\Lambda\|) \\ &\leq \|\Lambda\|(M + \|\Lambda\|) - (6(k + 1)(M + 1)\|\Lambda\|^4 - M - 2\|\Lambda\|) \\ &\leq -6(k + 1)(M + 1)\|\Lambda\|^4 + \|\Lambda\|^2 + 2\|\Lambda\| + M\|\Lambda\| + M \\ &\leq (-6k(M + 1) - 6(M + 1) + 3 + 2M)\|\Lambda\|^4 \\ &\leq -6k(M + 1)\|\Lambda\|^4 \\ &\leq -2k(2N + 1)\|\Lambda\|^4, \end{aligned}$$

where the last step follows since $M \geq 6\|\Lambda\|^3 N$. Hence $\mathbf{v} \in C$.

Recalling $(\mathbf{t}^\dagger - \mathbf{s}^\dagger)_1 \geq -\mathbf{s}_1^\dagger \geq -s_1 - \|\Lambda\|$ and

$$\begin{aligned} (\mathbf{t}^\dagger - \mathbf{s}^\dagger)_2 &\geq (\mathbf{t}' - \mathbf{s})_2 - 2\|\Lambda\| \\ &\geq \mathbf{t}_2 - (\mathbf{t} - \mathbf{t}')_2 - M - 2\|\Lambda\| \\ &\geq \mathbf{t}_2 - 6(k + 1)(M + 1)\|\Lambda\|^4 - M - 2\|\Lambda\| \\ &\geq \mathbf{t}_2 - 6(k + 2)(M + 1)\|\Lambda\|^4, \end{aligned}$$

we conclude using $M \geq 5N + 1$ that

$$\begin{aligned} \text{left}(\mathbf{v}) \cdot (\mathbf{s}_1, -\mathbf{t}_2) &\leq (-\mathbf{v}_2, \mathbf{v}_1) \cdot (\mathbf{s}^\dagger - \mathbf{t}^\dagger - (\|\Lambda\|, 6(k + 2)(M + 1)\|\Lambda\|^4)) \\ &\leq 2k(2N + 1)\|\Lambda\|^4 + \mathbf{v}_2\|\Lambda\| - \mathbf{v}_1 6(k + 2)(M + 1)\|\Lambda\|^4 \\ &\leq 2k(2N + 1)\|\Lambda\|^4 + \|\Lambda\|^2 + 6(k + 2)(M + 1)\|\Lambda\|^5 \\ &\leq (k(4N + 2) + 1 + 6(k + 2)(M + 1))\|\Lambda\|^5 \\ &< 7(k + 2)(M + 1)\|\Lambda\|^5. \end{aligned} \quad \square$$

Our final lemma that focusses on runs whose final points have small first coordinates and large second coordinates is the most powerful. It considers runs that start and finish near the vertical axis, but which perform a ‘zig-zag’ by getting near the horizontal axis, and whose final point is very high and higher than the initial point. In the statement, the runs are split into two parts: before (ρ) and after (π) visiting a point s which is near the horizontal axis. The lemma assumes that ρ after its first transition remains far from the vertical axis, and that π after its first transition remains far from the horizontal axis. An illustration is in Figure 13 on the right.

LEMMA 5.15. *Suppose we have an SLPS Λ with $k > 0$ cycles and $N, M \in \mathbb{N}$ with $M \geq 8\|\Lambda\|^4 N$. Let $\rho\pi \in \Lambda$ be a run consisting of a prefix ρ from point \mathbf{r} to point s and a suffix π from point s to point \mathbf{t} such that*

- $\mathbf{r}_1 < M, \mathbf{r}_2 \geq 0$,

- $s_1 \geq 0, s_2 < M$,
- $t_1 < M, t_2 \geq 19(k+2)(M+1)\|\Lambda\|^6, t_2 \geq r_2$,
- all points visited by ρ after r are in $\mathbb{N}_{\geq M} \times \mathbb{N}$, and
- all points visited by π after s are in $\mathbb{N} \times \mathbb{N}_{\geq M}$.

There exists $\gamma \in \{0, \dots, 2\|\Lambda\|^3\}$ such that, for all $n \in \{1, \dots, N\}$, $\rho\pi$ has a shortening by $(0, n\gamma)$ which is admissible from r .

PROOF. If case (1) of Lemma 5.14 applies to π from s then we are done immediately, so assume case (2) applies to it. Hence, for some cycle $\mathbf{v} \in \mathbb{Z}_{<0} \times \mathbb{Z}_{>0}$ which occurs in π at least $2\|\Lambda\|^2 N$ times, we have

$$\text{left}(\mathbf{v}) \cdot (s_1, -t_2) < 7(k+2)(M+1)\|\Lambda\|^5.$$

This also implies $s_1 \geq 12(k+1)(M+1)\|\Lambda\|^4$, since otherwise

$$\begin{aligned} \text{left}(\mathbf{v}) \cdot (s_1, -t_2) &> -\mathbf{v}_2 12(k+1)(M+1)\|\Lambda\|^4 - \mathbf{v}_1 t_2 \\ &\geq -12(k+1)(M+1)\|\Lambda\|^5 + t_2 \\ &\geq 7(k+2)(M+1)\|\Lambda\|^5. \end{aligned}$$

Consequently, Lemma 5.14 with $N\|\Lambda\|$ for N and with the axes swapped applies to ρ from r .

Suppose that case (2) of Lemma 5.14 holds. That is, for some cycle $\mathbf{w} \in \mathbb{Z}_{>0} \times \mathbb{Z}_{<0}$ which occurs in ρ at least $2\|\Lambda\|^3 N$ times, we have

$$(-\mathbf{w}_1, \mathbf{w}_2) \cdot (r_2, -s_1) < 7(k+2)(M+1)\|\Lambda\|^5.$$

We will reduce the occurrence of cycle \mathbf{w} in ρ by $-\mathbf{v}_1 \cdot n$ resulting in a shortening by $-\mathbf{v}_1 \cdot n \cdot \mathbf{w}$.

If case (1) of Lemma 5.14 with $N\|\Lambda\|$ for N and with the axes swapped applies to ρ from r , there is a value $\gamma' \in \{1, \dots, 2\|\Lambda\|^2\}$ such that we can shorten ρ by $-\mathbf{v}_1 \cdot n \cdot (\gamma', 0)$. For convenience, we define $\mathbf{w} \stackrel{\text{def}}{=} (\gamma', 0)$ in this case.

Either way, the resulting shortened version of ρ is admissible from r . In both cases, the second coordinate of points cannot decrease due to the shortening (note that $\mathbf{w}_2 \leq 0$). The first coordinate may decrease but by at most $\|\Lambda\| \cdot N \cdot 2\|\Lambda\|^2 = 2\|\Lambda\|^3 N < M$. Therefore, the shortened version of ρ is still admissible since, prior to the shortening, all points visited by ρ after r have a first coordinate of at least M .

Note that, while ρ is still admissible after the shortening, $\rho\pi$ may not be admissible anymore. Therefore, we also need to shorten π appropriately to counter the effect that the shortening of ρ may have had on the first coordinate. We shorten π by reducing the number of occurrences of cycle \mathbf{v} by $\mathbf{w}_1 \cdot n$. We now argue that such a shortened version of π is admissible from $s + \mathbf{v}_1 \cdot n \cdot \mathbf{w}$.

Following Lemma 5.14, π consists of two parts: a prefix of π , π' for which all intermediate points lie in $(\mathbb{N}_{\geq M})^2$, and the remaining path after π' . Note that the cycle \mathbf{v} is part of the path π' . Therefore the final point of π' as well as all points on the second part of π experience an increase of their first coordinate by $-\mathbf{v}_1 \cdot n \cdot \mathbf{w}_1$. Hence, after the shortening, all points on π starting at s have a first coordinate of at least $\min\{M, -\mathbf{v}_1 \cdot n \cdot \mathbf{w}_1\} = -\mathbf{v}_1 \cdot n \cdot \mathbf{w}_1$. Reducing the repetitions of the cycle \mathbf{v} by $\mathbf{w}_1 \cdot n$ can decrease the second coordinates of points on the path by no more than $\mathbf{w}_1 \cdot n \cdot \mathbf{v}_2 \leq 2\|\Lambda\|^3 N < M$ but all points visited by π prior to the shortening lie in $\mathbb{N} \times \mathbb{N}_{\geq M}$. Altogether we conclude that the shortening of π is not only admissible from s , but even admissible from $s + \mathbf{v}_1 \cdot n \cdot \mathbf{w}$.

Overall, we have a shortened version of ρ going from r to $s + \mathbf{v}_1 \cdot n \cdot \mathbf{w}$ that is admissible. This is followed by a shortened version of π going from $s + \mathbf{v}_1 n \mathbf{w}$ to $t + \mathbf{v}_1 n \mathbf{w} - \mathbf{w}_1 n \mathbf{v} = t - n(0, \mathbf{w}_1 \mathbf{v}_2 - \mathbf{w}_2 \mathbf{v}_1)$ and which is admissible as well.

Since we successfully shortened $\rho\pi$ by $n(0, \mathbf{w}_1\mathbf{v}_2 - \mathbf{w}_2\mathbf{v}_1)$ it only remains to show that $\mathbf{w}_1\mathbf{v}_2 - \mathbf{w}_2\mathbf{v}_1 \in \{0, \dots, 2\|\Lambda\|^3\}$. Clearly, $\mathbf{w}_1\mathbf{v}_2 - \mathbf{w}_2\mathbf{v}_1 < \mathbf{w}_1\mathbf{v}_2 \leq 2\|\Lambda\|^3$. On the other hand, it cannot be that $\mathbf{v}_1\mathbf{w}_2 > \mathbf{v}_2\mathbf{w}_1$, because it implies

$$\begin{aligned} t_2 &\leq (\mathbf{v}_2\mathbf{w}_1, \mathbf{v}_1\mathbf{w}_2) \cdot (-r_2, t_2) \\ &= -r_2\mathbf{v}_2\mathbf{w}_1 + t_2\mathbf{v}_1\mathbf{w}_2 \\ &= \mathbf{v}_2 \cdot (-\mathbf{w}_1, \mathbf{w}_2) \cdot (r_2, -s_1) - \mathbf{w}_2 \cdot (-\mathbf{v}_2, \mathbf{v}_1) \cdot (s_1, -t_2) \\ &< (\mathbf{v}_2 - \mathbf{w}_2) \cdot 7(k+2)(M+1)\|\Lambda\|^5 \\ &\leq 14(k+2)(M+1)\|\Lambda\|^6. \end{aligned}$$

□

We are now almost equipped to establish that 2-dimensional simple linear path schemes have pseudo-polynomially bounded reachability witnesses. Our last lemma applies to runs that remain far from both axes, and provides a bound on the norm of the points they can visit without compromising their minimality.

LEMMA 5.16. *Suppose we have an SLPS Λ with k cycles and $\|\Lambda\| > 0$. Let $\pi \in \Lambda$ be an $(\mathbb{N}_{\geq 6\|\Lambda\|^3})^2$ -run from point s to point t which visits some point f such that $\|f\| > 3\|\Lambda\|^2 \cdot \|\{s, t\}\| + 7.5\|\Lambda\|^5k$. There is a shortening of ρ by $\mathbf{0}$ that is admissible from s .*

PROOF. Since f is distinct from s and t , we have that $k > 0$, so

$$\begin{aligned} |\pi| &\geq 2(\|f\| - \|\{s, t\}\|)/\|\Lambda\| \\ &> 4\|\Lambda\| \cdot \|\{s, t\}\| + 15\|\Lambda\|^4k \\ &\geq 4\|\Lambda\| \|\mathbf{t} - \mathbf{s}\| + (k2\|\Lambda\|^2 + 1)(4\|\Lambda\|^2 + 1). \end{aligned}$$

In particular, $C = \text{cycles}_{\geq 2\|\Lambda\|^2}(\Lambda, \pi)$ cannot be empty. Suppose the cone of C does not contain $\mathbf{0}$. Then Lemma 5.6 provides a vector \mathbf{p} with $\|\mathbf{p}\| \leq 2\|\Lambda\|$ and $\mathbf{p} \cdot \mathbf{c} > 0$ for all $\mathbf{c} \in C$. By Cauchy–Bunyakovsky–Schwarz’s inequality and Lemma 5.9 (1) we then get

$$\begin{aligned} 4\|\Lambda\| \|\mathbf{t} - \mathbf{s}\| &\geq \mathbf{p} \cdot (\mathbf{t} - \mathbf{s}) \\ &\geq |\pi| - (k2\|\Lambda\|^2 + 1)(4\|\Lambda\|^2 + 1), \end{aligned}$$

which contradicts the inequation above. So the cone of C contains $\mathbf{0}$ and we finish by Lemma 5.11 with $N = 1$ and $\mathbf{c} = \mathbf{0}$. □

THEOREM 5.1. *For every SLPS Λ from state s to state t , there exists $B \leq (\|\Lambda\| \cdot \|\Lambda\|)^{O(1)}$ such that if $\pi \in \Lambda$ is minimal with $s(\mathbf{0}) \xrightarrow{\pi} \mathbb{N}^2 t(\mathbf{0})$ then $s(\mathbf{0}) \xrightarrow{\pi} (\mathbb{N}_{\leq B} \times \mathbb{N}_{\leq B}) t(\mathbf{0})$.*

PROOF. We can assume $\|\Lambda\|, k > 0$ where k is the number of cycles in Λ . Consider any shortest admissible $\pi \in \Lambda$ from $\mathbf{0}$ to $\mathbf{0}$, and let $M \stackrel{\text{def}}{=} 16\|\Lambda\|^7$.

First, we show that at all points visited by π where one coordinate is less than M , the other coordinate must be less than

$$M' \stackrel{\text{def}}{=} 969k\|\Lambda\|^{13} = 19(3k)(17\|\Lambda\|^7)\|\Lambda\|^6 \geq 19(k+2)(M+1)\|\Lambda\|^6. \quad (12)$$

To see this, assume the contrary and let $\mathbf{t} \in \mathbb{N}^2$ be a point visited by π from $\mathbf{0}$ such that, w.l.o.g., $t_1 < M$ and $t_2 \geq M'$. Further assume that \mathbf{t} is a point with maximum t_2 among all points with this property. Then we can extract a subpath ρ by following π backwards, starting in \mathbf{t} until for the first time a point s is visited that satisfies $s_2 < M$ and then further, until for the first time a point \mathbf{r} is visited with $r_1 < M$. (Here it may be the case that s and \mathbf{r} are the same point, i.e. the latter path segment is empty.) On this path Lemma 5.15 is applicable with $N = 2\|\Lambda\|^3$. So there exist

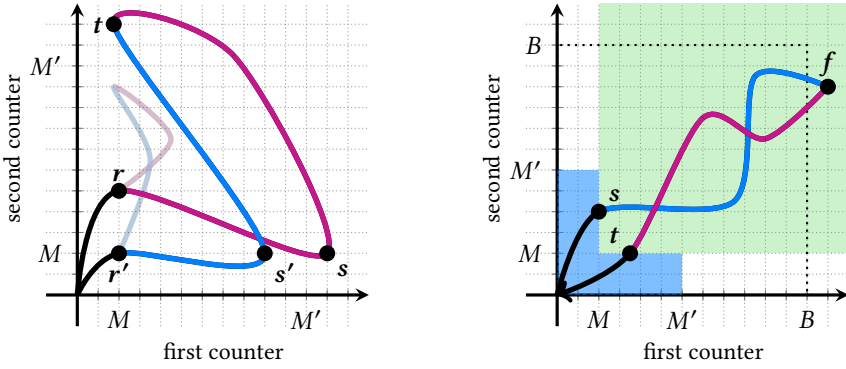


Fig. 14. Illustrations of the proof of Theorem 5.1. *Left*: we identify two path segments, one from r to t via s (in pink), and one from t to r' via s' (in blue). Both can be shortened via Lemma 5.15 (applied to the reverse of the second segment). Composing the shortenings produces a new path from 0 to 0 , in which the point corresponding to t has moved vertically down. *Right*: shortest paths cannot leave the shaded area. The path from s to t via a sufficiently large point f can be shortened using Lemma 5.16.

$\gamma \in \{0, \dots, 2\|\Lambda\|^3\}$ and shortenings by $(0, n\gamma)$ for all $n \in \{1, \dots, N\}$, admissible from the point r . If $\gamma = 0$ then this directly contradicts the minimality of π . Otherwise we can, analogously, extract a subpath ρ' by following π forwards from t to some r' and then reversing, so that Lemma 5.15 provides $\gamma' \in \{0, \dots, 2\|\Lambda\|^3\}$ and shortenings by $(0, n\gamma')$ for all $n \in \{1, \dots, N\}$, admissible backwards from r' . Together, this means there is a shortening of π by 0 ; a contradiction with the minimality assumption. (An illustration is in Figure 14 on the left.)

To show the claim of the theorem, assume that π visits some point f whose norm exceeds

$$B \stackrel{\text{def}}{=} 2914.5k\|\Lambda\|^{15} \geq 3\|\Lambda\|^2 M' + 7.5\|\Lambda\|^5 k.$$

Then we can partition $0 \xrightarrow{\pi} 0$ as $0 \xrightarrow{\rho} s \xrightarrow{\sigma} f \xrightarrow{\sigma'} t \xrightarrow{\tau} 0$ where $\|s\|, \|t\| < M'$ and all other points visited by $\sigma\sigma'$ from s are in $(\mathbb{N}_{\geq M})^2$. (This is illustrated in Figure 14 on the right, where the dotted line indicates the bound B above.) Now Lemma 5.16 provides a shortening of $\sigma\sigma'$ that is admissible from s , and thus a shortening of π admissible from 0 , again contradicting the minimality assumption. \square

6 CONCLUSION

We finish by highlighting four directions for further work:

- (1) Finkel et al. have recently completed a classification of extensions of 2-VASS by zero test, reset and transfer operations on their counters, in terms of computability and semi-linearity of their reachability sets and reachability relations [16]. However, in spite of the subsequent progress by Leroux and Sutre [30], the complexity of the reachability problem remains open in several cases.
- (2) The complexity of the reachability problem for branching VASS in dimension one has been settled both in the unary encoding and the binary encoding cases, namely P-complete [17] and PSPACE-complete [15] (respectively). However, in dimension two, it is not even known whether the problem is decidable.
- (3) Another interesting extension of 2-VASS is by replacing one of the counters by a pushdown stack. The resulting systems, which can be seen as one-dimensional *grammar-controlled vector*

addition systems (GVAS), are challenging: their coverability problem is known to be PSPACE-hard [12] and in EXPSPACE [31], and whether their reachability problem is decidable is open.

- (4) In spite of major recent progress on the complexity of the general VASS reachability problem [10, 28], relatively little is known specifically for dimension three; from the latter reference, 3-VASS reachability is in the non-elementary fast-growing complexity class F_7 [41]. As we have remarked, a substantial obstacle to applying the techniques from this article to 3-VASS is that their reachability relations (in fact, already reachability sets) are not semi-linear [19, proof of Lemma 2.8].

ACKNOWLEDGMENTS

We thank Jérôme Leroux for a helpful discussion about Section 4.2.2, and the referees for their careful reading and useful comments.

Michael Blondin acknowledges support by the Fonds de recherche du Québec – Nature et technologies (FRQNT), by the Technical University of Munich, by the Centre national de la recherche scientifique (CNRS), by the “Chaire Digiteo, ENS Cachan – École Polytechnique”, and by the Natural Sciences and Engineering Research Council of Canada (NSERC). Matthias Englert, Ranko Lazić and Patrick Totzke acknowledge support by the EPSRC, grants EP/M011801/1 and EP/M027651/1. Alain Finkel and Christoph Haase acknowledge support by the Agence nationale de la recherche, REACHARD (grant ANR-11-BS02-001). Stefan Göller and Christoph Haase acknowledge support by Labex Digicosme, Université Paris-Saclay, project VERICONISS. Pierre McKenzie acknowledges support by the Natural Sciences and Engineering Research Council of Canada (NSERC) and by the “Chaire Digiteo, ENS Cachan – École Polytechnique”.

REFERENCES

- [1] David Angeli, Patrick De Leenheer, and Eduardo D. Sontag. 2011. Persistence Results for Chemical Reaction Networks with Time-Dependent Kinetics and No Global Conservation Laws. *SIAM Journal of Applied Mathematics* 71, 1 (2011), 128–146. <https://doi.org/10.1137/090779401>
- [2] Paolo Baldan, Nicoletta Cocco, Andrea Marin, and Marta Simeoni. 2010. Petri nets for modelling metabolic pathways: a survey. *Natural Computing* 9, 4 (2010), 955–989. <https://doi.org/10.1007/s11047-010-9180-6>
- [3] Michael Blondin, Alain Finkel, Stefan Göller, Christoph Haase, and Pierre McKenzie. 2015. Reachability in Two-Dimensional Vector Addition Systems with States Is PSPACE-Complete. In *Proc. 30th Symposium on Logic in Computer Science (LICS)*. IEEE Computer Society, 32–43. <https://doi.org/10.1109/LICS.2015.14>
- [4] Stanislav Böhm, Stefan Göller, and Petr Jancar. 2013. Equivalence of deterministic one-counter automata is NL-complete. In *Proc. 45th Symposium on Theory of Computing (STOC)*. ACM, 131–140. <https://doi.org/10.1145/2488608.2488626>
- [5] Tomáš Brázdil, Stefan Kiefer, Antonín Kucera, and Petr Novotný. 2015. Long-Run Average Behaviour of Probabilistic Vector Addition Systems. In *Proc. 30th Symposium on Logic in Computer Science (LICS)*. 44–55. <https://doi.org/10.1109/LICS.2015.15>
- [6] E. Cardoza, Richard J. Lipton, and Albert R. Meyer. 1976. Exponential Space Complete Problems for Petri Nets and Commutative Semigroups: Preliminary Report. In *Proc. 8th Symposium on Theory of Computing (STOC)*. ACM, 50–54. <https://doi.org/10.1145/800113.803630>
- [7] Dmitry Chistikov, Wojciech Czerwinski, Piotr Hofman, Michal Pilipczuk, and Michael Wehar. 2019. Shortest paths in one-counter systems. *Logical Methods in Computer Science* 15, 1 (2019). <https://lmcs.episciences.org/5251>
- [8] Dmitry Chistikov and Christoph Haase. 2016. The Taming of the Semi-Linear Set. In *Proc. 43rd International Colloquium on Automata, Languages, and Programming (ICALP)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 128:1–128:13. <https://doi.org/10.4230/LIPIcs.ICALP.2016.128>
- [9] Wojciech Czerwinski and Slawomir Lasota. 2017. Regular separability of one counter automata. In *Proc. 32nd Symposium on Logic in Computer Science (LICS)*. 1–12. <https://doi.org/10.1109/LICS.2017.8005079>
- [10] Wojciech Czerwinski, Slawomir Lasota, Ranko Lazic, Jérôme Leroux, and Filip Mazowiecki. 2019. The Reachability Problem for Petri Nets is Not Elementary. In *Proc. 51st Symposium on Theory of Computing (STOC)*. ACM, 24–33. <https://doi.org/10.1145/3313276.3316369>

- [11] Wojciech Czerwinski, Sławomir Lasota, Christof Löding, and Radosław Piórkowski. 2019. New Pumping Technique for 2-Dimensional VASS. In *Proc. 44th International Symposium on Mathematical Foundations of Computer Science (MFCS) (LIPIcs)*, Vol. 138. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 62:1–62:14. <https://doi.org/10.4230/LIPIcs.MFCS.2019.62>
- [12] Matthias Englert, Piotr Hofman, Sławomir Lasota, Ranko Lazic, Jérôme Leroux, and Juliusz Straszynski. 2021. A lower bound for the coverability problem in acyclic pushdown VAS. *Inf. Process. Lett.* 167 (2021). <https://doi.org/10.1016/j.ipl.2020.106079>
- [13] Matthias Englert, Ranko Lazic, and Patrick Totzke. 2016. Reachability in Two-Dimensional Unary Vector Addition Systems with States is NL-Complete. In *Proc. 31st Symposium on Logic in Computer Science (LICS)*. ACM, 477–484. <https://doi.org/10.1145/2933575.2933577>
- [14] John Fearnley and Marcin Jurdziński. 2015. Reachability in two-clock timed automata is PSPACE-complete. *Information and Computation* 243 (2015), 26–36. <https://doi.org/10.1016/j.ic.2014.12.004>
- [15] Diego Figueira, Ranko Lazic, Jérôme Leroux, Filip Mazowiecki, and Grégoire Sutre. 2017. Polynomial-Space Completeness of Reachability for Succinct Branching VASS in Dimension One. In *Proc. 44th International Colloquium on Automata, Languages, and Programming (ICALP) (LIPIcs)*, Vol. 80. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 119:1–119:14. <https://doi.org/10.4230/LIPIcs.ICALP.2017.119>
- [16] Alain Finkel, Jérôme Leroux, and Grégoire Sutre. 2018. Reachability for Two-Counter Machines with One Test and One Reset. In *Proc. 38th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS) (LIPIcs)*, Vol. 122. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 31:1–31:14. <https://doi.org/10.4230/LIPIcs.FSTTCS.2018.31>
- [17] Stefan Göller, Christoph Haase, Ranko Lazic, and Patrick Totzke. 2016. A Polynomial-Time Algorithm for Reachability in Branching VASS in Dimension One. In *Proc. 43rd International Colloquium on Automata, Languages, and Programming (ICALP) (LIPIcs)*, Vol. 55. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 105:1–105:13. <https://doi.org/10.4230/LIPIcs.ICALP.2016.105>
- [18] Christoph Haase, Stephan Kreutzer, Joël Ouaknine, and James Worrell. 2009. Reachability in Succinct and Parametric One-Counter Automata. In *Proc. 20th International Conference on Concurrency Theory (CONCUR) (Lecture Notes in Computer Science)*, Vol. 5710. Springer, 369–383. https://doi.org/10.1007/978-3-642-04081-8_25
- [19] John E. Hopcroft and Jean-Jacques Pansiot. 1979. On the Reachability Problem for 5-Dimensional Vector Addition Systems. *Theoretical Computer Science* 8 (1979), 135–159. [https://doi.org/10.1016/0304-3975\(79\)90041-0](https://doi.org/10.1016/0304-3975(79)90041-0)
- [20] Rodney R. Howell, Louis E. Rosier, Dung T. Huynh, and Hsu-Chun Yen. 1986. Some Complexity Bounds for Problems Concerning Finite and 2-Dimensional Vector Addition Systems with States. *Theoretical Computer Science* 46, 3 (1986), 107–140. [https://doi.org/10.1016/0304-3975\(86\)90026-5](https://doi.org/10.1016/0304-3975(86)90026-5)
- [21] Alexander Kaiser, Daniel Kroening, and Thomas Wahl. 2014. A Widening Approach to Multithreaded Program Verification. *ACM Trans. Program. Lang. Syst.* 36, 4 (2014), 14:1–14:29. <https://doi.org/10.1145/2629608>
- [22] S. Rao Kosaraju. 1982. Decidability of Reachability in Vector Addition Systems (Preliminary Version). In *Proc. 14th Symposium on Theory of Computing (STOC)*. ACM, 267–281. <https://doi.org/10.1145/800070.802201>
- [23] Jean-Luc Lambert. 1992. A Structure to Decide Reachability in Petri Nets. *Theoretical Computer Science* 99, 1 (1992), 79–104. [https://doi.org/10.1016/0304-3975\(92\)90173-D](https://doi.org/10.1016/0304-3975(92)90173-D)
- [24] Hélène Leroux, David Andreu, and Karen Godary-Dejean. 2015. Handling Exceptions in Petri Net-Based Digital Architecture: From Formalism to Implementation on FPGAs. *IEEE Trans. Industrial Informatics* 11, 4 (2015), 897–906. <https://doi.org/10.1109/TII.2015.2435696>
- [25] Jérôme Leroux. 2010. The General Vector Addition System Reachability Problem by Presburger Inductive Invariants. *Logical Methods in Computer Science* 6, 3 (2010). [https://doi.org/10.2168/LMCS-6\(3:22\)2010](https://doi.org/10.2168/LMCS-6(3:22)2010)
- [26] Jérôme Leroux. 2011. Vector addition system reachability problem: a short self-contained proof. In *Proc. 38th Symposium on Principles of Programming Languages (POPL)*. ACM, 307–316. <https://doi.org/10.1145/1926385.1926421>
- [27] Jérôme Leroux. 2012. Vector Addition Systems Reachability Problem (A Simpler Solution). In *Turing-100 (EPIc Series in Computing)*, Vol. 10. EasyChair, 214–228. <http://www.easychair.org/publications/paper/106497>
- [28] Jérôme Leroux and Sylvain Schmitz. 2019. Reachability in Vector Addition Systems is Primitive-Recursive in Fixed Dimension. In *Proc. 34th Symposium on Logic in Computer Science (LICS)*. IEEE Computer Society. <https://doi.org/10.1109/LICS.2019.8785796>
- [29] Jérôme Leroux and Grégoire Sutre. 2004. On Flatness for 2-Dimensional Vector Addition Systems with States. In *Proc. 15th International Conference on Concurrency Theory (CONCUR) (Lecture Notes in Computer Science)*, Vol. 3170. Springer, 402–416. https://doi.org/10.1007/978-3-540-28644-8_26
- [30] Jérôme Leroux and Grégoire Sutre. 2020. Reachability in Two-Dimensional Vector Addition Systems with States: One Test Is for Free. In *Proc. 31st International Conference on Concurrency Theory (CONCUR) (LIPIcs)*, Vol. 171. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 37:1–37:17. <https://doi.org/10.4230/LIPIcs.CONCUR.2020.37>
- [31] Jérôme Leroux, Grégoire Sutre, and Patrick Totzke. 2015. On the Coverability Problem for Pushdown Vector Addition Systems in One Dimension. In *Proc. 42nd International Colloquium on Automata, Languages, and Programming (ICALP)*,

- Part II (Lecture Notes in Computer Science)*, Vol. 9135. Springer, 324–336. https://doi.org/10.1007/978-3-662-47666-6_26
- [32] Yuliang Li, Alin Deutsch, and Victor Vianu. 2017. VERIFAS: A Practical Verifier for Artifact Systems. *PVLDB* 11, 3 (2017), 283–296. <http://www.vldb.org/pvldb/vol11/p283-li.pdf>
 - [33] Richard J. Lipton. 1976. *The reachability problem requires exponential space*. Technical Report 62. Yale University. <http://cpsc.yale.edu/sites/default/files/files/tr63.pdf>
 - [34] Ernst W. Mayr. 1981. An Algorithm for the General Petri Net Reachability Problem. In *Proc. 13th Symposium on Theory of Computing (STOC)*. ACM, 238–246. <https://doi.org/10.1145/800076.802477>
 - [35] Jakub Michaliszyn, Jan Otop, and Piotr Wiecek. 2017. Querying Best Paths in Graph Databases. In *Proc. 37th LARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*. 43:1–43:15. <https://doi.org/10.4230/LIPIcs.FSTTCS.2017.43>
 - [36] Christos H. Papadimitriou. 1994. *Computational complexity*. Addison-Wesley.
 - [37] Carl Adam Petri. 1962. *Kommunikation mit Automaten*. Ph.D. Dissertation. Universität Hamburg. <http://edoc.sub.uni-hamburg.de/informatik/volltexte/2011/160/>
 - [38] Loïc Pottier. 1991. Minimal Solutions of Linear Diophantine Systems: bounds and Algorithms. In *Proc. 4th International Conference on Rewriting Techniques and Applications (RTA) (Lecture Notes in Computer Science)*, Vol. 488. Springer, 162–173. https://doi.org/10.1007/3-540-53904-2_94
 - [39] Charles Rackoff. 1978. The Covering and Boundedness Problems for Vector Addition Systems. *Theoretical Computer Science* 6 (1978), 223–231. [https://doi.org/10.1016/0304-3975\(78\)90036-1](https://doi.org/10.1016/0304-3975(78)90036-1)
 - [40] George S. Sacerdote and Richard L. Tenney. 1977. The Decidability of the Reachability Problem for Vector Addition Systems (Preliminary Version). In *Proc. 9th Symposium on Theory of Computing (STOC)*. ACM, 61–76. <https://doi.org/10.1145/800105.803396>
 - [41] Sylvain Schmitz. 2016. Complexity Hierarchies beyond Elementary. *ACM Transactions on Computation Theory* 8, 1 (2016), 3:1–3:36. <https://doi.org/10.1145/2858784>
 - [42] Leslie G. Valiant and Mike Paterson. 1975. Deterministic One-Counter Automata. *J. Comput. System Sci.* 10, 3 (1975), 340–350. [https://doi.org/10.1016/S0022-0000\(75\)80005-5](https://doi.org/10.1016/S0022-0000(75)80005-5)

Received ; revised ; accepted