Improved Algorithm for Reachability in d-VASS

Yuxi Fu ✉
BASICS, Shanghai Jiao Tong University

Qizhe Yang ✉ 🆔
Shanghai Normal University

Yangluo Zheng ✉ 🆔
BASICS, Shanghai Jiao Tong University

— **Abstract** ———

An F_d upper bound for the reachability problem in vector addition systems with states (VASS) in fixed dimension is given, where F_d is the d-th level of the Grzegorczyk hierarchy of complexity classes. The new algorithm combines the idea of the linear path scheme characterization of the reachability in the 2-dimension VASSes with the general decomposition algorithm by Mayr, Kosaraju and Lambert. The result improves the F_{d+4} upper bound due to Leroux and Schmitz (LICS 2019).

1 Introduction

Petri nets, or equivalently vector addition system with states (VASS), are a well studied model of concurrency. A VASS consists of a finite state control where each state transition has as its effect an integer-valued vector, and its configurations are pairs of a state and a vector with *natural number* components. A transition may lead from one configuration to another by adding its effect component-wise, conditioned on that the components of the resulting vector remain non-negative. The *reachability problem*, which asks whether from one configuration there is a path reaching another configuration, lies in the center of the algorithmic theory of Petri nets and has found a wide range of applications due to its generic nature. Since the problem was shown to be decidable by Mayr [17] in 1981, its computational complexity had been a long-standing open problem in the field. In 2015, Leroux and Schmitz [14] presented the first complexity upper bound, stating that the reachability problem is cubic-Ackermannian. This was later improved to an Ackermannian upper bound in 2019, again by Leroux and Schmitz [15]. Regarding the hardness, in 2021 seminal works by Czerwiński and Orlikowski [6], and independently by Leroux [13], provided matching Ackermannian lower bounds, settling the exact complexity of the problem.

Concerning the parameterization by dimension, i.e. the reachability problem in d-dimensional VASSes where d is fixed, there is still a gap in the known complexity bounds. Currently, we only have exact complexity bounds for dimension one and two [7, 1]. For

dimension $d \geq 3$, the result of Leroux and Schmitz [15] shows that the problem is in F_{d+4}, the $(d+4)$-th level of the Grzegorczyk hierarchy of complexity classes. Note that a recent work by Yang and Fu [22] points out that the problem for dimension 3 is in $\mathsf{F}_3 = \mathsf{TOWER}$. On the other hand, the best known lower bound by Czerwiński, Jecker, Lasota, Leroux and Orlikowski [5] states that reachability in $(2d+3)$-dimensional VASSes is F_d-hard. Motivated by this gap, our paper focuses on the computational complexity of reachability problem in the fixed-dimensional VASSes.

Our contribution.

In this paper we show that the reachability problem in the d-dimensional VASS is in F_d for $d \geq 3$, improving the previous F_{d+4} upper bound by Leroux and Schmitz [15], and generalizing the tower upper bound for the reachability problem in 3-VASS [22]. The new upper bound is obtained with the help of two novel technical lemmas.

1. Our main technical tool (Theorem 3.4) is a generalization of the linear path scheme characterization for the reachability relation in the 2-dimensional VASSes [1]. By borrowing the key idea from the work of Yang and Fu [22], we show that as long as the "geometric dimension" of a VASS (that is, the dimension of the vector space spanned by the effects of cyclic paths) is bounded by 2, its reachability relation can be characterized by short linear path schemes. We then apply the lemma to simplify the KLMST algorithm so that (i) a VASS is replaced by a short linear path scheme whenever its geometric dimension is no more than 2 and (ii) the linear path schemes will not be decomposed further. It is then routine [15], using the tools from [21], to show that the reachability problem in the d-dimensional VASS is in F_{d+1} for all $d \geq 3$.
2. Our second lemma (Lemma 6.3) allows us to improve further the bound from F_{d+1} to F_d. This is done by a careful analysis of the properties of the fast-growing functions [21].

Due to space limitation the proofs of the two lemmas are placed in the appendices.

Organization.

Section 2 fixes notation, defines the VASS model and its reachability problem. Section 3 generalizes the linear path scheme characterization [1] to VASSes whose geometric dimension are bounded by 2. Section 4 recalls the characterization system of linear inequalities for linear path schemes. Section 5 makes use of the results of Section 3 to give an improved version of the classic KLMST decomposition algorithm. Section 6 analyzes the complexity of our modified algorithm, proving the main result. Section 7 concludes. Proofs omitted from the main text can be found in the appendices.

2   Preliminaries

We use $\mathbb{N}, \mathbb{Z}, \mathbb{Q}$ to denote respectively the set of non-negative integers, integers, and rational numbers. Let $n \in \mathbb{N}$ be a number, we write $[n]$ for the range $\{1, 2, \ldots, n\}$. Let $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{X}^d$ be d-dimensional vectors where $\mathbb{X}$ can be any one of $\mathbb{N}, \mathbb{Z}, \mathbb{Q}$. We write $\boldsymbol{v}(i)$ for the i-th component of $\boldsymbol{v}$ where $i \in [d]$, so $\boldsymbol{v} = (\boldsymbol{v}(1), \ldots, \boldsymbol{v}(d))$. The *maximum norm* of $\boldsymbol{v}$ is defined to be $\|\boldsymbol{v}\| := \max_{i \in [d]} |\boldsymbol{v}(i)|$. We extend component-wise the order $\leq$ for vectors in $\mathbb{X}^d$, so $\boldsymbol{u} \leq \boldsymbol{v}$ if and only if $\boldsymbol{u}(i) \leq \boldsymbol{v}(i)$ for all $i \in [d]$. Addition and subtraction of vectors are also component-wise, so $(\boldsymbol{u} + \boldsymbol{v})(i) = \boldsymbol{u}(i) + \boldsymbol{v}(i)$ for all $i \in [d]$. Define $\mathrm{supp}(\boldsymbol{v}) := \{i \in [d] : \boldsymbol{v}(i) \neq 0\}$ to be the set of indices of non-zero components of $\boldsymbol{v}$. This notation is extended to sets of vectors naturally, so $\mathrm{supp}(S) = \bigcup_{\boldsymbol{v} \in S} \mathrm{supp}(\boldsymbol{v})$ for any set $S \subseteq \mathbb{X}^d$.

For technical reasons we introduce the symbol ω that stands for the infinite element. Let $\mathbb{N}_\omega := \mathbb{N} \cup \{\omega\}$. We stipulate that $n < \omega$ for all $n \in \mathbb{N}$, and $x + \omega = \omega + x = \omega$ for all $x \in \mathbb{Z}$. Define the partial order $\sqsubseteq$ over $\mathbb{N}_\omega$ so that $x \sqsubseteq y$ if and only if $x = y$ or $y = \omega$ for all $x, y \in \mathbb{N}_\omega$. The relation $\sqsubseteq$ is also extended component-wise to vectors in $\mathbb{N}_\omega^d$.

Let Σ be a finite alphabet and $s, t \in \Sigma^*$ be two strings over Σ. We write st for the concatenation of s and t, and s^n for the concatenation of n copies of s where $n \in \mathbb{N}$. If $s = a_1 a_2 \ldots a_\ell$ where $a_1, \ldots, a_\ell \in \Sigma$, we write $|s| := \ell$ for the length of s, and $s[i \ldots j] := a_i a_{i+1} \ldots a_j$ for the substring of s between indices i and j where $1 \le i \le j \le |s|$.

2.1 Vector Addition Systems with States

Let $d \ge 0$ be an integer. A *d-dimensional vector addition system with states (d-VASS)* is a pair $G = (Q, T)$ where Q is a finite set of *states* and $T \subseteq Q \times \mathbb{Z}^d \times Q$ is a finite set of *transitions*. Clearly a VASS can also be viewed as a directed graph with edges labelled by integer vectors. Given a word $\pi = (p_1, \boldsymbol{a}_1, q_1)(p_2, \boldsymbol{a}_2, q_2) \ldots (p_n, \boldsymbol{a}_n, q_n) \in T^*$ over transitions, we say that π is a *path from p_1 to q_n* if $q_i = p_{i+1}$ for all $i = 1, \ldots, n-1$. It is a *cycle* if we further have $p_1 = q_n$. The *effect* of π is defined to be $\Delta(\pi) := \sum_{i=1}^n \boldsymbol{a}_i$, and the *action word* of π is the word $[\![\pi]\!] := \boldsymbol{a}_1 \boldsymbol{a}_2 \ldots \boldsymbol{a}_n$ over $\mathbb{Z}^d$. The *Parikh image* of π is the function $\phi \in \mathbb{N}^T$ mapping each transition to its number of occurrences in π. Given a function $\phi \in \mathbb{N}^T$ we also define $\Delta(\phi) := \sum_{t=(p,\boldsymbol{a},q) \in T} \phi(t) \cdot \boldsymbol{a}$. Note that $\Delta(\phi) = \Delta(\pi)$ if ϕ is the Parikh image of π. Let $L \subseteq T^*$ be a language (i.e. subset of words), we define its effect as $\Delta(L) := \{\Delta(\pi) : \pi \in L\}$.

The norm of a transition $t = (p, \boldsymbol{a}, q)$ is defined by $\|t\| := \|\boldsymbol{a}\|$. The norm of a path $\pi = t_1 t_2 \ldots t_n$ is $\|\pi\| := \max_{i \in [n]} \|t_i\|$. For a d-VASS $G = (Q, T)$ we write $\|T\| := \max\{\|t\| : t \in T\}$. The *size* of G is defined by

$$|G| := |Q| + |T| + d \cdot |T| \cdot \|T\|. \tag{1}$$

Semantics of VASSes.

Let $G = (Q, T)$ be a d-VASS. A *configuration* of G is a pair of a state $p \in Q$ and a vector $\boldsymbol{v} \in \mathbb{Z}^d$, written as $p(\boldsymbol{v})$. Let $\mathbb{D} \subseteq \mathbb{Z}^d$, we define the $\mathbb{D}$-semantics for G as follows. For each transition $t = (p, \boldsymbol{a}, q) \in T$, the one-step transition relation $\xrightarrow{t}_\mathbb{D}$ relates all pairs of configurations of the form $(p(\boldsymbol{u}), q(\boldsymbol{v}))$ where $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{D}$ and $\boldsymbol{v} = \boldsymbol{u} + \boldsymbol{a}$. Then for a word $\pi = t_1 t_2 \ldots t_n \in T^*$, the relation $\xrightarrow{\pi}_\mathbb{D}$ is the composition $\xrightarrow{\pi}_\mathbb{D} := \xrightarrow{t_1}_\mathbb{D} \circ \cdots \circ \xrightarrow{t_n}_\mathbb{D}$. So $p(\boldsymbol{u}) \xrightarrow{\pi}_\mathbb{D} q(\boldsymbol{v})$ if and only if there are configurations $p_0(\boldsymbol{u}_0), \ldots, p_n(\boldsymbol{u}_n) \in Q \times \mathbb{D}$ such that

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

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

For the above definitions, we shall often omit the subscript $\mathbb{D}$ when $\mathbb{D} = \mathbb{N}^d$.

We mention that in Section 5 we need to generalize the VASS semantics to configurations in $Q \times \mathbb{N}_\omega^d$, allowing ω components in vectors. The definitions of $\xrightarrow{t}_{\mathbb{N}_\omega^d}$, $\xrightarrow{\pi}_{\mathbb{N}_\omega^d}$, and $\xrightarrow{*}_{\mathbb{N}_\omega^d}$ are similar to the above.

Reachability problem.

The reachability problem in vector addition systems with states is formulated as follows:

> REACHABILITY IN d-DIMENSIONAL VECTOR ADDITION SYSTEM WITH STATES
> **Input:** A d-dimensional VASS $G = (Q, T)$, two configurations $p(\boldsymbol{u}), q(\boldsymbol{v}) \in Q \times \mathbb{N}^d$.
> **Question:** Does $p(\boldsymbol{u}) \xrightarrow{*}_{\mathbb{N}^d} q(\boldsymbol{v})$ hold in G?

Note that we study the reachability problem for fixed-dimensional VASSes, where the dimension d is treated as a constant to allow more fine-grained analysis. So we shall use the big-O notation to hide constants that may depend on d. The general problem where the dimension can be part of the input was already shown to be Ackermann-complete [6, 13].

Cycle spaces and geometric dimensions.

One of the key insights of [15] is a new termination argument for the KLMST decomposition algorithm based on the dimensions of vector spaces spanned by cycles in VASSes, which yielded the primitive recursive upper bound of VASS reachability problem in fixed dimensions. The vector spaces spanned by cycles also play an important role in our work.

▶ **Definition 2.1.** *Let G be a d-VASS. The* cycle space *of G is the vector space* $\mathrm{Cyc}(G) \subseteq \mathbb{Q}^d$ *spanned by the effects of all cycles in G, that is:*

$$\mathrm{Cyc}(G) := \mathrm{span}\{\Delta(\beta) : \beta \text{ is a cycle in } G\}. \tag{3}$$

The dimension of the cycle space of G is called the geometric dimension *of G. We say G is geometrically k-dimensional if* $\dim(\mathrm{Cyc}(G)) \leq k$.

3 Flattability of Geometrically 2-dimensional VASSes

A VASS is *flat* if each of its states lies on at most one cycle. Flat VASSes form an important subclass of VASSes due to its connection to Presburger arithmetic, and we refer the readers to [12] for a survey. In dimension 2, it was proved in [1] that 2-VASSes enjoy a stronger form of flat representation, known as the *linear path schemes*. A linear path scheme is a regular expression of the form $\alpha_0 \beta_1^* \alpha_1 \ldots \beta_k^* \alpha_k$ where $\alpha_0, \ldots \alpha_k$ are paths and $\beta_1, \ldots, \beta_k$ are cycles of the VASS, such that they form a path when joined together. The results of [1] show that the reachability relation of every 2-VASS can be captured by short linear path schemes.

Linear path schemes are extremely useful as they can be fully characterized by linear inequality systems so that standard tools for linear or integer programming can be applied. In this section, we generalize the results in [1] and show that the reachability relation of any d-VASS whose geometric dimension is bounded by 2 can also be captured by short linear path schemes.

Our proof follows closely to the lines of [1]. Given a geometrically 2-dimensional VASS G, we first cover $\mathbb{N}^d$ by the following two regions: one for the region far away from every axis:

$$\mathbb{O} := \{\boldsymbol{u} \in \mathbb{N}^d : \boldsymbol{u}(i) \geq D \text{ for all } i \in [d]\} \tag{4}$$

where D is some properly chosen threshold; the other one for the region close to some axis:

$$\mathbb{L} := \{\boldsymbol{u} \in \mathbb{N}^d : \boldsymbol{u}(i) \leq D' \text{ for some } i \in [d]\} \tag{5}$$

where D' is chosen slightly above D to create an overlap with $\mathbb{O}$. Let π be a run that we are going to capture by linear path schemes. We can extract its maximal prefix that
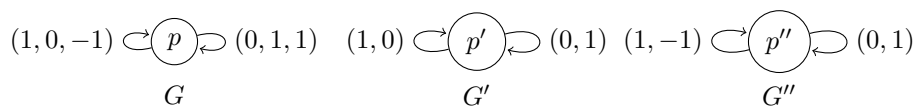
Figure 1 A geometrically 2-dimensional VASS G and two possible projections of it.

lies completely in either $\mathbb{O}$ or $\mathbb{L}$, depending on where π starts. This prefix must end at a configuration that lies in $\mathbb{L} \cap \mathbb{O}$, if we haven't touched the end of π. From this configuration we then extract a maximal cycle that also ends in $\mathbb{L} \cap \mathbb{O}$. Continuing this fashion, we can break π into segments of runs that lie completely in $\mathbb{O}$ or $\mathbb{L}$, interleaved by cycles that start and end in $\mathbb{O}$ (actually in $\mathbb{L} \cap \mathbb{O}$). Note that the number of such cycles cannot exceed the number of states of G since they are maximal. Now we only need to capture the following three types of runs by short linear path schemes:

1. Runs that are cycles starting and ending in $\mathbb{O}$.
   This will be handled in Section A.4 in the appendix. Briefly speaking, since the geometric dimension of G is 2, the effect of such a cycle must belong to a plane in $\mathbb{Z}^d$. We will find a clever way to project this plane to a coordinate plane, and then project the d-VASS G onto this plane to get a 2-VASS. This is made possible by a novel technique called the "sign-reflecting projection" developed in Section A.3. Intuitively speaking, for any vector in a plane we are able to determine whether it belongs to a certain orthant by observing only 2 entries of this vector. The d-VASS can then be projected onto these 2 coordinates. (See Example 3.1 for a more concrete demonstration.) Now we apply the results of [1] to obtain a linear path scheme that captures the projected cyclic run. The projection guarantees that we can safely project it back to get a linear path scheme for the run in G.
2. Runs that lie completely in $\mathbb{O}$.
   This is just an easy corollary of the first type, and will also be handled in Section A.4. Just note that any run can be broken into a series of simple paths interleaved by cycles.
3. Runs that lie completely in $\mathbb{L}$.
   This will be handled in Section A.5, by a long and tedious case analysis. In principle, we are going to argue that any such run essentially corresponds to a run in some $(d-1)$-VASS, so that we can use induction.

▶ **Example 3.1.** Consider the geometrically 2-dimensional 3-VASS G as shown in Figure 1. It consists of a single state p and two transitions with effects $(1, 0, -1)$ and $(0, 1, 1)$. In order to apply the results of [1], one would like to derive a 2-VASS from G that reflects runs in G. A simple idea is to discard one coordinate of G. Two possibilities of this idea are shown in Figure 1 as G', which discards the third coordinate, and G'', which discards the second coordinate. However, not all of them are satisfactory. For example, the legal run $p(0,0) \xrightarrow{(1,0)} p(1,0)$ in G' reflects an illegal run $p(0,0,0) \xrightarrow{(1,0,-1)} p(1,0,-1)$ in G where the third coordinate goes negative. On the other hand, all runs in G'' can be safely projected back to a run in G. To see this, just observe that for any vector $\boldsymbol{v}$ in the linear span of $(1, 0, -1)$ and $(0, 1, 1)$, $\boldsymbol{v} \geq \boldsymbol{0}$ if and only if $\boldsymbol{v}(1) \geq 0$ and $\boldsymbol{v}(3) \geq 0$. Thus we can safely discard the second coordinate.

In general, the "sign-reflecting projection" developed in Section A.3 shows that any geometrically 2-dimensional VASS can be projected onto two coordinates so that the signs of these two coordinates reflects the signs of other coordinates.

In the rest of this section we just state formally our main technical results. The detailed proofs are placed in the appendix.

3.1   Linear Path Schemes

A *linear path scheme (LPS)* is a pair (G, Λ) where G is a VASS and Λ is a regular expression of the form $\Lambda = \alpha_0 \beta_1^* \alpha_1 \ldots \beta_k^* \alpha_k$ such that $\alpha_0, \ldots, \alpha_k$ are paths in G and $\beta_1, \ldots, \beta_k$ are cycles in G, and $\alpha_0 \beta_1 \alpha_1 \ldots \beta_k \alpha_k$ is a path in G. We say an LPS (G, Λ) is *compatible to* a VASS G' if G' contains all states and transitions in Λ. Very often we shall omit the VASS G and say that Λ on its own is an LPS, with G understood as any VASS to which Λ is compatible. We write $|\Lambda| = |\alpha_0 \beta_1 \alpha_1 \ldots \beta_k \alpha_k|$ for the length of Λ, $\|\Lambda\| = \|\alpha_0 \beta_1 \alpha_1 \ldots \beta_k \alpha_k\|$ for its norm, and $|\Lambda|_* = k$ for the number of cycles in Λ.

We also treat Λ as the language defined by it, and thus for two configurations $p(\boldsymbol{u})$ and $q(\boldsymbol{v})$ we write $p(\boldsymbol{u}) \xrightarrow{\Lambda}_{\mathbb{D}} q(\boldsymbol{v})$ if and only if there exists $e_1, \ldots, e_k \in \mathbb{N}$ such that

$$p(\boldsymbol{u}) \xrightarrow{\alpha_0 \beta_1^{e_1} \alpha_1 \ldots \beta_k^{e_k} \alpha_k}_{\mathbb{D}} q(\boldsymbol{v}). \tag{6}$$

Positive linear path schemes.

A *positive LPS* is a regular expression of the form $\Lambda^+ = \alpha_0 \beta_1^+ \alpha_1 \ldots \beta_k^+ \alpha_k$ which is similar to a linear path scheme except that we require each cycle to be used at least once. We write $p(\boldsymbol{u}) \xrightarrow{\Lambda^+}_{\mathbb{D}} q(\boldsymbol{v})$ if and only if there are positive integers $e_1, \ldots, e_k \in \mathbb{N}_{>0}$ such that

$$p(\boldsymbol{u}) \xrightarrow{\alpha_0 \beta_1^{e_1} \alpha_1 \ldots \beta_k^{e_k} \alpha_k}_{\mathbb{D}} q(\boldsymbol{v}). \tag{7}$$

A path π is said to be *admitted by* Λ^+ if $\pi = \alpha_0 \beta_1^{e_1} \alpha_1 \ldots \beta_k^{e_k} \alpha_k$ for some $e_1, \ldots, e_k \in \mathbb{N}_{>0}$. We prefer positive LPSes as they can be easily characterized by linear inequality systems (see Section 4 for details). In fact, positive LPSes can be easily obtained from LPSes:

▶ **Lemma 3.2.** *For every linear path scheme Λ there exists a finite set S of positive linear path schemes compatible to the same VASSes with Λ, such that $\xrightarrow{\Lambda} = \bigcup_{\Lambda^+ \in S} \xrightarrow{\Lambda^+}$ and $|\Lambda^+| \le |\Lambda|$ for every $\Lambda^+ \in S$.*

Proof. Suppose $\Lambda = \alpha_0 \beta_1^* \alpha_1 \ldots \beta_k^* \alpha_k$. For each cycle component β_i^* in Λ we replace it by either β_i^+ or an empty word nondeterministically. Let S be the set of all resulting positive LPSes. It is obvious that S satisfies the desired properties. ◀

3.2   Main Results

The main results of this section is stated as follows.

▶ **Theorem 3.3.** *Let $G = (Q, T)$ be a geometrically 2-dimensional d-VASS. For every pair of configurations $p(\boldsymbol{u}), q(\boldsymbol{v}) \in Q \times \mathbb{N}^d$ with $p(\boldsymbol{u}) \xrightarrow{*} q(\boldsymbol{v})$ there exists a positive LPS Λ^+ compatible to G such that $p(\boldsymbol{u}) \xrightarrow{\Lambda^+} q(\boldsymbol{v})$ and $|\Lambda^+| \le |G|^{O(1)}$.*

We remark that the big-O term here and elsewhere in the paper hides constant that may depend on the dimension d, but does not depend on $G, \boldsymbol{u}, \boldsymbol{v}$ or anything else.

By Lemma 3.2 we know that positive LPSes can be obtained from LPSes. Thus theorem 3.3 follows easily from the following relaxed theorem, which will be proved in the appendix.

▶ **Theorem 3.4.** *Let $G = (Q, T)$ be a geometrically 2-dimensional d-VASS. For every pair of configurations $p(\boldsymbol{u}), q(\boldsymbol{v}) \in Q \times \mathbb{N}^d$ with $p(\boldsymbol{u}) \xrightarrow{*} q(\boldsymbol{v})$ there exists an LPS Λ compatible to G such that $p(\boldsymbol{u}) \xrightarrow{\Lambda} q(\boldsymbol{v})$ and $|\Lambda| \le |G|^{O(1)}$.*

4 Characteristic Systems for Linear Path Schemes

The property that linear path schemes can be fully characterized by linear inequality systems is exploited in [2] to derive the PSPACE upper bound of the reachability problem in 2-VASSes. Here we recall this linear inequality system and its properties.

We mainly focus on positive linear path schemes. Fix $\Lambda = \alpha_0 \beta_1^+ \alpha_1 \ldots \beta_k^+ \alpha_k$ to be a positive LPS from state p to q that is compatible to some d-VASS $G = (Q, T)$, where $k = |\Lambda|_*$ is the number of cycles in Λ.

▶ **Definition 4.1** (cf. [2, Lem. 14]). *The* characteristic system $\mathcal{E}_{\mathrm{LPS}}(\Lambda)$ *of the positive LPS* Λ *is the system of linear inequalities such that a triple* $\boldsymbol{h} = (\boldsymbol{u}, \boldsymbol{e}, \boldsymbol{v}) \in \mathbb{N}^d \times \mathbb{N}^k \times \mathbb{N}^d$ *satisfies* $\mathcal{E}_{\mathrm{LPS}}(\Lambda)$, *written* $\boldsymbol{h} \models \mathcal{E}_{\mathrm{LPS}}(\Lambda)$, *if and only if the following conditions hold:*

1. *for every* $i = 1, \ldots, k$, $\boldsymbol{e}(i) \geq 1$;
2. *for every* $i = 0, \ldots, k$ *and every* $j = 1, \ldots, |\alpha_i|$,

$$\boldsymbol{u} + \Delta(\alpha_0 \beta_1^{\boldsymbol{e}(1)} \alpha_1 \ldots \alpha_{i-1} \beta_i^{\boldsymbol{e}(i)}) + \Delta(\alpha_i[1 \ldots j]) \geq \boldsymbol{0}; \tag{8}$$

3. *for every* $i = 1, \ldots, k$ *and every* $j = 1, \ldots, |\beta_i|$,

$$\boldsymbol{u} + \Delta(\alpha_0 \beta_1^{\boldsymbol{e}(1)} \alpha_1 \ldots \beta_{i-1}^{\boldsymbol{e}(i-1)} \alpha_{i-1}) \qquad + \Delta(\beta_i[1 \ldots j]) \geq \boldsymbol{0}, \tag{9}$$

$$\boldsymbol{u} + \Delta(\alpha_0 \beta_1^{\boldsymbol{e}(1)} \alpha_1 \ldots \beta_{i-1}^{\boldsymbol{e}(i-1)} \alpha_{i-1} \beta_i^{\boldsymbol{e}(i)-1}) + \Delta(\beta_i[1 \ldots j]) \geq \boldsymbol{0}; \tag{10}$$

4. *and finally,* $\boldsymbol{u} + \Delta(\alpha_0 \beta_1^{\boldsymbol{e}(1)} \alpha_1 \ldots \beta_k^{\boldsymbol{e}(k)} \alpha_k) = \boldsymbol{v}$.

The readers can easily verify that these constraints are indeed linear in terms of $\boldsymbol{u}, \boldsymbol{e}, \boldsymbol{v}$. The next lemma shows that $\mathcal{E}_{\mathrm{LPS}}(\Lambda)$ indeed captures all runs admitted by Λ.

▶ **Lemma 4.2.** *Let G be a d-VASS and $\Lambda = \alpha_0 \beta_1^+ \alpha_1 \ldots \beta_k^+ \alpha_k$ be a positive LPS from state p to q compatible to G. Then for every $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{N}^d$, $p(\boldsymbol{u}) \xrightarrow{\Lambda} q(\boldsymbol{v})$ if and only if there exists $\boldsymbol{e} \in \mathbb{N}^k$ such that $(\boldsymbol{u}, \boldsymbol{e}, \boldsymbol{v}) \models \mathcal{E}_{\mathrm{LPS}}(\Lambda)$. Moreover, for every $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{N}^d$ and every $\boldsymbol{e} \in \mathbb{N}^k$ such that $(\boldsymbol{u}, \boldsymbol{e}, \boldsymbol{v}) \models \mathcal{E}_{\mathrm{LPS}}(\Lambda)$, we have $p(\boldsymbol{u}) \xrightarrow{\alpha_0 \beta_1^{\boldsymbol{e}(1)} \alpha_1 \ldots \beta_k^{\boldsymbol{e}(k)} \alpha_k} q(\boldsymbol{v})$.*

We also need to introduce the homogeneous version of $\mathcal{E}_{\mathrm{LPS}}(\Lambda)$ for technical reasons.

▶ **Definition 4.3.** *The* homogeneous characteristic system $\mathcal{E}_{\mathrm{LPS}}^0(\Lambda)$ *of* Λ *is the system of linear inequalities such that a triple* $\boldsymbol{h}_0 = (\boldsymbol{u}_0, \boldsymbol{e}_0, \boldsymbol{v}_0) \in \mathbb{N}^d \times \mathbb{N}^k \times \mathbb{N}^d$ *satisfies* $\mathcal{E}_{\mathrm{LPS}}^0(\Lambda)$, *written* $\boldsymbol{h}_0 \models \mathcal{E}_{\mathrm{LPS}}^0(\Lambda)$, *if and only if the following conditions hold:*

1. *for every* $i = 0, \ldots, k$, $\boldsymbol{u}_0 + \Delta(\beta_1) \cdot \boldsymbol{e}_0(1) + \cdots + \Delta(\beta_i) \cdot \boldsymbol{e}_0(i) \geq \boldsymbol{0}$;
2. $\boldsymbol{u}_0 + \Delta(\beta_1) \cdot \boldsymbol{e}_0(1) + \cdots + \Delta(\beta_k) \cdot \boldsymbol{e}_0(k) = \boldsymbol{v}_0$.

5 The Modified KLMST Decomposition Algorithm

In this section we apply our results of Section 3 to improve the notoriously hard KLMST decomposition algorithm for VASS reachability. Our narration will base on the work of Leroux and Schmitz [15]. For readers familiar with [15], the major modifications are listed below:

- The decomposition structure is now a sequence of generalized VASS reachability instances linked by (positive) linear path schemes rather than by single transitions.

- We introduce a new "cleaning" step that replaces all VASS instances which are geometrically 2-dimensional by polynomial-length linear path schemes compatible to them.
- We do not guarantee the exact preservation of action languages at each decomposition step. Instead, we only preserve a subset of action languages. This is a compromise since linear path schemes capture only the reachability relation but not every possible run. Nonetheless, it is enough for the reachability problem.

In this section we focus on the effectiveness and correctness of the modified KLMST decomposition algorithm. Its complexity will be analyzed in Section 6.

5.1 Linear KLM Sequences

The underlying decomposition structure in the KLMST algorithm was known as KLM sequences, named after Mayr[17], Kosaraju[10], and Lambert[11].

▶ **Definition 5.1.** *A KLM tuple of dimension d is a tuple $\langle p(\boldsymbol{x})Gq(\boldsymbol{y})\rangle$ where $G = (Q, T)$ is a d-VASS and $p(\boldsymbol{x}), q(\boldsymbol{y}) \in Q \times \mathbb{N}_\omega^d$ are two (generalized) configurations of G. A KLM sequence of dimension d is a sequence of KLM tuples interleaved by transitions of the form*

$$\xi = \langle p_0(\boldsymbol{x}_0)G_0q_0(\boldsymbol{y}_0)\rangle\, t_1\, \langle p_1(\boldsymbol{x}_1)G_1q_1(\boldsymbol{y}_1)\rangle\, t_2 \ldots t_k\, \langle p_k(\boldsymbol{x}_k)G_kq_k(\boldsymbol{y}_k)\rangle, \tag{11}$$

where each tuple $\langle p_i(\boldsymbol{x}_i)G_iq_i(\boldsymbol{y}_i)\rangle$ is a KLM tuple of dimension d and each t_i is a transition of the form $(q_{i-1}, \boldsymbol{a}_i, p_i)$ from state q_{i-1} to p_i with effect $\boldsymbol{a}_i \in \mathbb{Z}^d$.

In this paper we generalize the definition of KLM sequences to allow (positive) linear path schemes to connect KLM tuples.

▶ **Definition 5.2.** *A linear KLM sequence of dimension d is a sequence*

$$\xi = \langle p_0(\boldsymbol{x}_0)G_0q_0(\boldsymbol{y}_0)\rangle\, \Lambda_1\, \langle p_1(\boldsymbol{x}_1)G_1q_1(\boldsymbol{y}_1)\rangle\, \Lambda_2 \ldots \Lambda_k\, \langle p_k(\boldsymbol{x}_k)G_kq_k(\boldsymbol{y}_k)\rangle, \tag{12}$$

where each tuple $\langle p_i(\boldsymbol{x}_i)G_iq_i(\boldsymbol{y}_i)\rangle$ is a KLM tuple of dimension d and each Λ_i is a positive linear path scheme from state q_{i-1} to p_i.

One immediately sees that KLM sequences are just special cases of linear KLM sequences. Let ξ be a linear KLM sequence given as (12), we write $\xi_i := \langle p_i(\boldsymbol{x}_i)G_iq_i(\boldsymbol{y}_i)\rangle$ for the ith KLM tuple occurring in ξ.

Action languages.

Let ξ be a linear KLM sequence given as (12). We say a path π from state p_0 to q_k is *admitted* by ξ, written $\xi \vdash \pi$, if π can be written as $\pi = \pi_0\rho_1\pi_1 \ldots \rho_k\pi_k$ where π_i is a path from p_i to q_i in G_i for each $i = 0, \ldots, k$, and ρ_i is a path admitted by Λ_i for each $i = 1, \ldots, k$, such that there exist vectors $\boldsymbol{m}_0, \boldsymbol{n}_0, \ldots, \boldsymbol{m}_k, \boldsymbol{n}_k \in \mathbb{N}^d$ such that

$$p_0(\boldsymbol{m}_0) \xrightarrow{\pi_0} q_0(\boldsymbol{n}_0) \xrightarrow{\rho_1} p_1(\boldsymbol{m}_1) \xrightarrow{\pi_1} q_1(\boldsymbol{n}_1) \xrightarrow{\rho_2} \cdots \xrightarrow{\rho_k} p_k(\boldsymbol{m}_k) \xrightarrow{\pi_k} q_k(\boldsymbol{n}_k) \tag{13}$$

and that $\boldsymbol{m}_i \sqsubseteq \boldsymbol{x}_i$, $\boldsymbol{n}_i \sqsubseteq \boldsymbol{y}_i$ for each $i = 0, \ldots, k$.

The *action language* L_ξ of ξ is the language over $\mathbb{Z}^d$ defined by $L_\xi := \big\{[\![\pi]\!] : \xi \vdash \pi\big\}$, where we recall that $[\![\cdot]\!]$ is the word morphism mapping each transition to its effect.

We are more interested in the action languages because in some decomposition steps we have to modify the set of transitions, and only the action word of admitted runs can be preserved. Notice that action languages preserve not only the effects of admitted runs, but also their lengths.

Ranks and Sizes.

Let t be a transition in a d-VASS G, we define $\mathrm{Cyc}(G/t)$ to be the vector space spanned by the effects of all cycles in G containing t. For the VASS G, let r_i be the number of transitions t in G such that $\dim(\mathrm{Cyc}(G/t)) = i$ for each $i = 0, \ldots, d$. Then the *rank* of G is defined as $\mathrm{rank}(G) = (r_d, \ldots, r_3) \in \mathbb{N}^{d-2}$. We also define the *full rank* of G as $\mathrm{rank}_{\mathrm{full}}(G) = (r_d, \ldots, r_0) \in \mathbb{N}^{d+1}$.

The following lemma was proved in [15], which shows that in a strongly connected VASS G, the space $\mathrm{Cyc}(G/t)$ corresponds to $\mathrm{Cyc}(G)$.

▶ **Lemma 5.3** ([15, Lem. 3.2]). *Let t be a transition of a strongly connected VASS G. Then $\mathrm{Cyc}(G/t) = \mathrm{Cyc}(G)$.*

The following corollary is immediate.

▶ **Corollary 5.4.** *Let G be a strongly connected d-VASS. Then $\mathrm{rank}(G) = \mathbf{0}$ if and only if G is geometrically 2-dimensional.*

Let ξ be a linear KLM sequence given as (12). We define the *rank* of ξ as $\mathrm{rank}(\xi) = \sum_{i=0}^{k} \mathrm{rank}(G_i)$, and the *full rank* of ξ as $\mathrm{rank}_{\mathrm{full}}(\xi) = \sum_{i=0}^{k} \mathrm{rank}_{\mathrm{full}}(G_i)$. We remark that the full rank corresponds to the rank defined in [15]. Ranks are ordered lexicographically: let $\boldsymbol{r} = (r_d, \ldots, r_0)$ and $\boldsymbol{r}' = (r_d', \ldots, r_0')$, we write $\boldsymbol{r} \leq_{\mathrm{lex}} \boldsymbol{r}'$ if $\boldsymbol{r} = \boldsymbol{r}'$ or the maximal i with $r_i \neq r_i'$ satisfies $r_i < r_i'$.

Recall that for a VASS G we write $|G|$ for its size as defined in (1). For a linear path scheme Λ, its length $|\Lambda|$ and norm $\|\Lambda\|$ are defined in Section 3.1. Let $\zeta = \langle p(\boldsymbol{x})Gq(\boldsymbol{y})\rangle$ be a KLM tuple of dimension d, its *size* is defined to be $|\zeta| = |G| + d \cdot (\|\boldsymbol{x}\| + \|\boldsymbol{y}\| + 1)$. Let ξ be a linear KLM sequence given as (12), we define its *size* as

$$|\xi| = \sum_{i=0}^{k} |\xi_i| + \sum_{i=1}^{k} d \cdot |\Lambda_i| \cdot (\|\Lambda_i\| + 1). \tag{14}$$

Note that the sizes defined in this paper reflect the sizes of unary encoding, thus have an exponential expansion in their binary encoding.

5.2 Characteristic Systems for Linear KLM Sequences

We define in this section the characteristic systems of linear KLM sequences, which are systems of linear inequalities that serve as an under-specification of admitted runs. Let $G = (Q, T)$ be a VASS, we first recall the *Kirchhoff system* $K_{G,p,q}$ of G with respect to states $p, q \in Q$, which is a system of linear equations such that a function $\phi \in \mathbb{N}^T$ is a model of $K_{G,p,q}$, written $\phi \models K_{G,p,q}$, if and only if

$$\mathbf{1}_q - \mathbf{1}_p = \sum_{t=(r,\boldsymbol{a},s)\in T} \phi(t) \cdot (\mathbf{1}_s - \mathbf{1}_r), \tag{15}$$

where $\mathbf{1}_p \in \{0,1\}^Q$ is the indicator function defined by $\mathbf{1}_p(q) = 1$ if $q = p$ and $\mathbf{1}_p(q) = 0$ otherwise. We also need the homogeneous version of $K_{G,p,q}$, denoted by $K_{G,p,q}^0$, where a function $\phi \in \mathbb{N}^T$ is a model of it, written $\phi \models K_{G,p,q}^0$, if and only if

$$\mathbf{0} = \sum_{t=(r,\boldsymbol{a},s)\in T} \phi(t) \cdot (\mathbf{1}_s - \mathbf{1}_r). \tag{16}$$

▶ **Definition 5.5.** *Let ξ be a linear KLM sequence given by*

$$\xi = \langle p_0(\boldsymbol{x}_0)G_0q_0(\boldsymbol{y}_0)\rangle \, \Lambda_1 \, \langle p_1(\boldsymbol{x}_1)G_1q_1(\boldsymbol{y}_1)\rangle \, \Lambda_2 \ldots \Lambda_k \, \langle p_k(\boldsymbol{x}_k)G_kq_k(\boldsymbol{y}_k)\rangle \tag{17}$$

The characteristic system $\mathcal{E}(\xi)$ *is a set of linear (in)equalities such that a sequence*

$$\boldsymbol{h} = (\boldsymbol{m}_0, \phi_0, \boldsymbol{n}_0), \boldsymbol{e}_1, (\boldsymbol{m}_1, \phi_1, \boldsymbol{n}_1), \boldsymbol{e}_2, \ldots, \boldsymbol{e}_k, (\boldsymbol{m}_k, \phi_k, \boldsymbol{n}_k), \tag{18}$$

where each $(\boldsymbol{m}_i, \phi_i, \boldsymbol{n}_i) \in \mathbb{N}^d \times \mathbb{N}^{T_i} \times \mathbb{N}^d$ and each $\boldsymbol{e}_i \in \mathbb{N}^{|\Lambda_i|_*}$, is a model of $\mathcal{E}(\xi)$, written $\boldsymbol{h} \models \mathcal{E}(\xi)$, if and only if
1. $\boldsymbol{m}_i \sqsubseteq \boldsymbol{x}_i$, $\phi_i \models K_{G,p,q}$, $\boldsymbol{n}_i \sqsubseteq \boldsymbol{y}_i$ *and* $\boldsymbol{n}_i = \boldsymbol{m}_i + \Delta(\phi_i)$ *for every* $i = 0, \ldots, k$;
2. $(\boldsymbol{n}_{i-1}, \boldsymbol{e}_i, \boldsymbol{m}_i) \models \mathcal{E}_{\mathrm{LPS}}(\Lambda_i)$ *for every* $i = 1, \ldots, k$.

Similarly, the homogeneous characteristic system $\mathcal{E}^0(\xi)$ *is a set of linear (in)equalities such that a sequence $\boldsymbol{h}$ of the form (18) is a model of $\mathcal{E}^0(\xi)$, written $\boldsymbol{h} \models \mathcal{E}^0(\xi)$, if and only if*
1. $\boldsymbol{m}_i(j) = 0$ *whenever* $\boldsymbol{x}_i(j) \neq \omega$, $\phi_i \models K^0_{G,p,q}$, $\boldsymbol{n}_i(j) = 0$ *whenever* $\boldsymbol{y}_i(j) \neq \omega$, *and* $\boldsymbol{n}_i = \boldsymbol{m}_i + \Delta(\phi_i)$ *for every* $i = 0, \ldots, k$;
2. $(\boldsymbol{n}_{i-1}, \boldsymbol{e}_i, \boldsymbol{m}_i) \models \mathcal{E}^0_{\mathrm{LPS}}(\Lambda_i)$ *for every* $i = 1, \ldots, k$.

The sequence ξ is said to be satisfiable *if $\mathcal{E}(\xi)$ has a model, otherwise it's unsatisfiable.*

Let $\boldsymbol{h}$ be a model of $\mathcal{E}(\xi)$ (or $\mathcal{E}^0(\xi)$), we shall write $\boldsymbol{m}_i^{\boldsymbol{h}}, \phi_i^{\boldsymbol{h}}, \boldsymbol{n}_i^{\boldsymbol{h}}, \boldsymbol{e}_i^{\boldsymbol{h}}$ for the values of $\boldsymbol{m}_i, \phi_i, \boldsymbol{n}_i, \boldsymbol{e}_i$ assigned by $\boldsymbol{h}$, respectively. Recall that unsatisfiable linear KLM sequences have empty action languages.

▶ **Lemma 5.6** (cf. [15, Lem. 3.5])**.** *For any unsatisfiable linear KLM sequence ξ, $L_\xi = \emptyset$.*

5.2.1 Bounds on Bounded Values in $\mathcal{E}(\xi)$

We state here a lemma similar to [15, Lem. 3.7], which upper bounds the bounded values in the characteristic system $\mathcal{E}(\xi)$. Its proof can be found in the appendix.

▶ **Lemma 5.7.** *Assume that $\xi = \langle p_0(\boldsymbol{x}_0)G_0q_0(\boldsymbol{y}_0)\rangle \, \Lambda_1 \cdots \Lambda_k \, \langle p_k(\boldsymbol{x}_k)G_kq_k(\boldsymbol{y}_k)\rangle$ is satisfiable. Then for every $0 \leq i \leq k$ we have:*
- *For every $1 \leq j \leq d$, the set of values $\boldsymbol{m}_i^{\boldsymbol{h}}(j)$ where $\boldsymbol{h} \models \mathcal{E}(\xi)$ is unbounded if, and only if, there exists a model $\boldsymbol{h}_0$ of $\mathcal{E}^0(\xi)$ such that $\boldsymbol{m}_i^{\boldsymbol{h}_0}(j) > 0$.*
- *For every $t \in T_i$, the set of values $\phi_i^{\boldsymbol{h}}(t)$ where $\boldsymbol{h} \models \mathcal{E}(\xi)$ is unbounded if, and only if, there exists a model $\boldsymbol{h}_0$ of $\mathcal{E}^0(\xi)$ such that $\phi_i^{\boldsymbol{h}}(t) > 0$.*
- *For every $1 \leq j \leq d$, the set of values $\boldsymbol{n}_i^{\boldsymbol{h}}(j)$ where $\boldsymbol{h} \models \mathcal{E}(\xi)$ is unbounded if, and only if, there exists a model $\boldsymbol{h}_0$ of $\mathcal{E}^0(\xi)$ such that $\boldsymbol{n}_i^{\boldsymbol{h}_0}(j) > 0$.*

Moreover, every bounded value of $\mathcal{E}(\xi)$ is bounded by $(10|\xi|)^{12|\xi|}$.

5.3 Cleaning of Linear KLM Sequences

In this section we define three conditions that require a linear KLM sequence to be *strongly connected*, *pure*, and *saturated*. Together with the satisfiability condition, they make up the so-called "clean" condition of linear KLM sequences. Note that the purity condition is new compared to [15], which requires every geometrically 2-dimensional VASSes occur in a linear KLM sequence to be replaced by linear path schemes.

Strongly Connected Sequences

A linear KLM sequence $\xi = \langle p_0(\boldsymbol{x}_0)G_0q_0(\boldsymbol{y}_0)\rangle \Lambda_1 \cdots \Lambda_k \langle p_k(\boldsymbol{x}_k)G_kq_k(\boldsymbol{y}_k)\rangle$ is *strongly connected* if all the VASSes $G_0, \ldots, G_k$ are strongly connected (as they are understood as directed graphs). One can easily obtain strongly connected sequences by expanding the strongly connected components of each VASS:

▶ **Lemma 5.8** ([15, Lem. 4.2]). *For any linear KLM sequence ξ which is not strongly connected, one can compute in time $\exp(|\xi|)$ a finite set Ξ of strongly connected linear KLM sequences such that $L_\xi = \bigcup_{\xi' \in \Xi} L_{\xi'}$ and that $\mathrm{rank}(\xi') \leq_{lex} \mathrm{rank}(\xi)$ and $|\xi'| \leq (2d+1)|\xi|$ for every $\xi' \in \Xi$.*

Pure Sequences

A KLM tuple $\langle p(\boldsymbol{x})Gq(\boldsymbol{y})\rangle$ is called *trivial* if $p(\boldsymbol{x}) = q(\boldsymbol{y})$ and G contains no transition and only a single state p. In this case we simply write $\langle p(\boldsymbol{x})\rangle$ for this tuple. Note that the action language of a trivial tuple contains exactly the empty word.

A linear KLM sequence $\xi = \langle p_0(\boldsymbol{x}_0)G_0q_0(\boldsymbol{y}_0)\rangle \Lambda_1 \cdots \Lambda_k \langle p_k(\boldsymbol{x}_k)G_kq_k(\boldsymbol{y}_k)\rangle$ is said to be *pure* if ξ is strongly connected and for every $i = 0, \ldots, k$, $\mathrm{rank}(G_i) = \boldsymbol{0}$ implies that the tuple $\langle p_i(\boldsymbol{x}_i)G_iq_i(\boldsymbol{y}_i)\rangle$ is trivial. By Corollary 5.4, a rank-$\boldsymbol{0}$ strongly connected VASS is geometrically 2-dimensional, and thus can be replaced by linear path schemes in case it is not trivial.

▶ **Lemma 5.9.** *Let ξ be a strongly connected linear KLM sequence. Whether ξ is pure is in* PSPACE. *If ξ is not pure, one can compute in space $\mathrm{poly}(|\xi|)$ a finite set Ξ of pure linear KLM sequences such that $\bigcup_{\xi' \in \Xi} L_{\xi'} \subseteq L_\xi$ and $\bigcup_{\xi' \in \Xi} L_{\xi'} \neq \emptyset$ whenever $L_\xi \neq \emptyset$, and such that $\mathrm{rank}(\xi') = \mathrm{rank}(\xi)$ and $|\xi'| \leq |\xi|^{O(1)}$ for all $\xi' \in \Xi$.*

Saturated Sequences

Let $\xi = \langle p_0(\boldsymbol{x}_0)G_0q_0(\boldsymbol{y}_0)\rangle \Lambda_1 \cdots \Lambda_k \langle p_k(\boldsymbol{x}_k)G_kq_k(\boldsymbol{y}_k)\rangle$ be a linear KLM sequence. We say ξ is *saturated* if for every $0 \leq i \leq k$ and every $j \in [d]$, we have

- $\boldsymbol{x}_i(j) = \omega$ implies the set of values $\boldsymbol{m}_i^{\boldsymbol{h}}(j)$ where $\boldsymbol{h} \models \mathcal{E}(\xi)$ is unbounded; and
- $\boldsymbol{y}_i(j) = \omega$ implies the set of values $\boldsymbol{n}_i^{\boldsymbol{h}}(j)$ where $\boldsymbol{h} \models \mathcal{E}(\xi)$ is unbounded.

▶ **Lemma 5.10** ([15, Lem. 4.4]). *From any pure linear KLM sequence ξ, one can compute in time $\exp(|\xi|^{O(|\xi|)})$ a finite set Ξ of saturated pure linear KLM sequences such that $L_\xi = \bigcup_{\xi' \in \Xi} L_{\xi'}$, and such that $\mathrm{rank}(\xi') = \mathrm{rank}(\xi)$ and $|\xi'| \leq |\xi|^{O(|\xi|)}$ for every $\xi' \in \Xi$.*

Proof. By Lemma 5.7, if a variable $\boldsymbol{m}_i(j)$ or $\boldsymbol{n}_i(j)$ is bounded in $\mathcal{E}(\xi)$, we can replace the corresponding ω component in ξ by all possible values bounded by $(10|\xi|)^{12|\xi|} \leq |\xi|^{O(|\xi|)}$. ◀

The Cleaning Lemma

A linear KLM sequence ξ is called *clean* if it is satisfiable, strongly connected, pure and saturated. The lemmas 5.8 through 5.10 show how to make a linear KLM sequence clean.

▶ **Lemma 5.11.** *From any linear KLM sequence ξ, one can compute in time $\exp(g(|\xi|))$ a finite set $\mathrm{clean}(\xi)$ of clean linear KLM sequences such that $\bigcup_{\xi' \in \mathrm{clean}(\xi)} L_{\xi'} \subseteq L_\xi$ and $\bigcup_{\xi' \in \mathrm{clean}(\xi)} L_{\xi'} \neq \emptyset$ whenever $L_\xi \neq \emptyset$. Moreover, for every $\xi' \in \mathrm{clean}(\xi)$ we have $\mathrm{rank}(\xi') \leq_{lex} \mathrm{rank}(\xi)$ and $|\xi'| \leq g(|\xi|)$ where $g(x) = x^{x^{O(1)}}$.*

5.4 Decomposition of Linear KLM Sequences

In this section we recall three conditions that require a linear KLM sequence to be *unbounded*, *rigid*, and *pumpable*. If any one of them is violated, a decomposition into a set of linear KLM sequences with strictly smaller ranks can be performed. Essentially there is nothing new in this section compared to [15]. The decomposition operations in [15] can be directly applied here, since they operate on a single KLM tuple and produce KLM sequences that are just special cases of linear KLM sequences. The proofs in [15] can also be adapted easily, and we will omit the details here. Especially, the next lemma shows that the arguments of strict decrease in ranks are still valid even though we discard the lower three components of ranks.

▶ **Lemma 5.12.** *Let ξ' be a pure linear KLM sequence. For any linear KLM sequence ξ' with $\mathrm{rank}_{full}(\xi') <_{lex} \mathrm{rank}_{full}(\xi)$, we have $\mathrm{rank}(\xi') <_{lex} \mathrm{rank}(\xi)$.*

Unbounded Sequences

Let $\xi = \langle p_0(\boldsymbol{x}_0)G_0q_0(\boldsymbol{y}_0)\rangle \Lambda_1 \cdots \Lambda_k \langle p_k(\boldsymbol{x}_k)G_kq_k(\boldsymbol{y}_k)\rangle$ be a linear KLM sequence. We say ξ is *unbounded* if for all $i = 0, \ldots, k$ and every transition $t \in T_i$ where T_i is the set of transitions of G_i, the set of values $\phi_i^{\boldsymbol{h}}(t)$ where $\boldsymbol{h} \models \mathcal{E}(\xi)$ is unbounded. Bounded transitions can be expanded exhaustively according to the bounds given by Lemma 5.7.

▶ **Lemma 5.13** ([15, Lem. 4.6]). *Whether a linear KLM sequence ξ is unbounded is decidable in* NP. *Moreover, if ξ is pure and bounded, one can compute in time $\exp(|\xi|^{O(|\xi|)})$ a finite set Ξ of linear KLM sequences such that $L_\xi = \bigcup_{\xi' \in \Xi} L_{\xi'}$ and such that $\mathrm{rank}(\xi') <_{lex} \mathrm{rank}(\xi)$ and $|\xi'| < |\xi|^{O(|\xi|)}$ for every $\xi' \in \Xi$.*

Rigid Sequences

A coordinate $j \in [d]$ is said to be *fixed* by a VASS $G = (Q, T)$ if there exists a function $f_j : Q \to \mathbb{N}$ such that $f_j(q) = f_j(p) + \boldsymbol{a}(j)$ for every transition $(p, \boldsymbol{a}, q) \in T$. We also say that f_j *fixes* G at coordinate j in this case.

A KLM tuple $\langle p(\boldsymbol{x})Gq(\boldsymbol{y})\rangle$ is said to be *rigid* if for every coordinate j fixed by $G = (Q, T)$, there exists a function $g_j : Q \to \mathbb{N}$ that fixes G at coordinate j and such that $g_j(p) \sqsubseteq \boldsymbol{x}(j)$ and $g_j(q) \sqsubseteq \boldsymbol{y}(j)$.

A linear KLM sequence $\xi = \langle p_0(\boldsymbol{x}_0)G_0q_0(\boldsymbol{y}_0)\rangle \Lambda_1 \cdots \Lambda_k \langle p_k(\boldsymbol{x}_k)G_kq_k(\boldsymbol{y}_k)\rangle$ is said to be *rigid* if every tuple $\langle p_i(\boldsymbol{x}_i)G_iq_i(\boldsymbol{y}_i)\rangle$ in ξ is rigid.

▶ **Lemma 5.14** ([15, Lem. 4.9]). *From any pure linear KLM sequence ξ one can decide in polynomial time whether ξ is not rigid. Moreover, in that case one can compute in polynomial time a linear KLM sequence ξ' such that $L_\xi = L_{\xi'}$, $\mathrm{rank}(\xi') <_{lex} \mathrm{rank}(\xi)$, and $|\xi'| \leq |\xi|$.*

Pumpable Sequences

Given a KLM tuple $\langle p(\boldsymbol{x})Gq(\boldsymbol{y})\rangle$, recall the *forward and backward acceleration vectors* $\mathrm{FACC}_{G,p}(\boldsymbol{x}), \mathrm{BACC}_{G,q}(\boldsymbol{y}) \in \mathbb{N}_\omega^d$ defined by

$$\mathrm{FACC}_{G,p}(\boldsymbol{x})(j) = \begin{cases} \omega & \text{if } p(\boldsymbol{x}) \xrightarrow{*} p(\boldsymbol{x}') \text{ for some } \boldsymbol{x}' \text{ with } \boldsymbol{x}' \geq \boldsymbol{x}, \boldsymbol{x}'(j) > \boldsymbol{x}(j) \\ \boldsymbol{x}(j) & \text{otherwise} \end{cases} \tag{19}$$

$$\mathrm{BACC}_{G,q}(\boldsymbol{y})(j) = \begin{cases} \omega & \text{if } q(\boldsymbol{y}') \xrightarrow{*} q(\boldsymbol{y}) \text{ for some } \boldsymbol{y}' \text{ with } \boldsymbol{y}' \geq \boldsymbol{y}, \boldsymbol{y}'(j) > \boldsymbol{y}(j) \\ \boldsymbol{y}(j) & \text{otherwise} \end{cases} \tag{20}$$

A tuple $\langle p(\boldsymbol{x})Gq(\boldsymbol{y})\rangle$ is said to be *pumpable* if $\mathrm{FACC}_{G,p}(\boldsymbol{x})(j) = \mathrm{BACC}_{G,q}(\boldsymbol{y})(j) = \omega$ for every coordinate j *not* fixed by G.

A linear KLM sequence given by $\xi = \langle p_0(\boldsymbol{x}_0)G_0q_0(\boldsymbol{y}_0)\rangle \Lambda_1 \cdots \Lambda_k \langle p_k(\boldsymbol{x}_k)G_kq_k(\boldsymbol{y}_k)\rangle$ is said to be *pumpable* if every tuple $\langle p_i(\boldsymbol{x}_i)G_iq_i(\boldsymbol{y}_i)\rangle$ in ξ is pumpable.

▶ **Lemma 5.15** ([15, Lem. 4.15])**.** *Whether a linear KLM sequence ξ is pumpable is decidable in* EXPSPACE*. Moreover, if ξ is pure and unpumpable, one can compute in time $\exp(|\xi|^{O(1)})$ a finite set Ξ of linear KLM sequences such that $L_\xi = \bigcup_{\xi' \in \Xi} L_{\xi'}$ and such that $\mathrm{rank}(\xi') <_{lex} \mathrm{rank}(\xi)$ and $|\xi'| < |\xi|^{O(1)}$ for every $\xi' \in \Xi$.*

Note that the $O(1)$ term here hides a constant depending on d, which essentially arises from a result on the coverability problem by Rackoff [19]. The $O(1)$ term also captures the difference between the sizes of linear KLM sequences defined here and that in [15].

The Decomposition Lemma

A linear KLM sequence is *normal* if it is clean, unbounded, rigid, and pumpable. The lemmas 5.13 through 5.15 show that when a clean linear KLM sequence is not normal, we are able to decompose it into a finite set of linear KLM sequences of strictly smaller ranks.

▶ **Lemma 5.16.** *From any clean linear KLM sequences ξ that is not normal, one can compute in time $\exp(h(\xi))$ a finite set $\mathrm{dec}(\xi)$ of clean linear KLM sequence such that $\bigcup_{\xi' \in \mathrm{dec}(\xi)} L_{\xi'} \subseteq L_\xi$ and $\bigcup_{\xi' \in \mathrm{dec}(\xi)} L_{\xi'} \neq \emptyset$ whenever $L_\xi \neq \emptyset$. Moreover, for every $\xi' \in \mathrm{dec}(\xi)$ we have $\mathrm{rank}(\xi') <_{lex} \mathrm{rank}(\xi)$ and $|\xi'| \leq h(\xi)$ where $h(x) = x^{x^{x^{O(1)}}}$.*

5.5 Normal Sequences

The following lemma shows that a normal linear KLM sequence is guaranteed to have non-empty action language, thus one can terminate the decomposition process once a normal sequence is produced.

▶ **Lemma 5.17.** *Let ξ be a normal linear KLM sequence, then there is a word $\sigma \in L_\xi$ whose length is bounded by $|\sigma| \leq \ell(|\xi|)$ where $\ell(x) \leq x^{O(x)}$.*

5.6 Putting All Together: The Modified KLMST Algorithm

Here we describe the modified KLMST decomposition algorithm for VASS reachability problem. Suppose we are given a d-VASS $G = (Q, T)$ and two configurations $p(\boldsymbol{m}), q(\boldsymbol{n}) \in Q \times \mathbb{N}^d$. To decide whether $p(\boldsymbol{m}) \xrightarrow{*} q(\boldsymbol{n})$ holds in G, it is enough to decide whether L_ξ is non-empty where $\xi = \langle p(\boldsymbol{m})Gq(\boldsymbol{n})\rangle$. To start with, we use Lemma 5.11 to clean the sequence ξ, and then choose $\xi^0 \in \mathrm{clean}(\xi)$ non-deterministically. If ξ^0 is normal then we are done by Lemma 5.17. Otherwise, we decompose ξ^0 using Lemma 5.16 and choose $\xi^1 \in \mathrm{dec}(\xi^0)$ non-deterministically. The procedure continues to produce a series of linear KLM sequences $\xi^0, \xi^1, \xi^2, \ldots$ where $\xi^{i+1} \in \mathrm{dec}(\xi^i)$, until either we finally obtain a normal sequence ξ^n, or at some point we have to abort because the decomposition of a linear KLM sequence is the empty set. The procedure terminates because $\mathrm{rank}(\xi^0) >_{\mathrm{lex}} \mathrm{rank}(\xi^1) >_{\mathrm{lex}} \mathrm{rank}(\xi^2) >_{\mathrm{lex}} \cdots$ form a decreasing chain of the well-order $(\mathbb{N}^{d-2}, <_{\mathrm{lex}})$, which must be finite. If $L_\xi = \emptyset$ then we cannot get a normal sequence since the action languages $L_\xi \supseteq L_{\xi^0} \supseteq L_{\xi^1} \supseteq \cdots$ are all empty. On the other hand, if $L_\xi \neq \emptyset$ then there are non-deterministic choices that always choose the linear KLM sequences with non-empty action languages, which finally lead to a normal sequence. This shows the correctness of the algorithm.

6   Complexity Upper Bound

The termination of the modified KLMST decomposition algorithm is guaranteed by a ranking function that decreases along a well-ordering. In order to analyze the length of this decreasing chain, we recall the so-called "length function theorems" by Schmitz [20] in Section 6.1. After that, we can locate the complexity upper bound of the algorithm in the fast-growing complexity hierarchy [21] which we recall in Section 6.2. Readers familiar with [15] may realize that the complexity upper bound for d-VASS can be improved to F_{d+1}, i.e. the $(d+1)$-th level in the fast-growing hierarchy, with our ranking function. In fact, by a careful analysis on a property of fast-growing functions, we further improve this bound to F_d.

In this section we assume some familiarity with ordinal numbers (see, e.g. [9]). We write ω here for the first infinite ordinal, not to be confused with the infinite element in previous sections.

6.1   Length of Sequences of Decreasing Ranks

Let ξ be a linear KLM sequence of dimension d with $\mathrm{rank}(\xi) = (r_d, \ldots, r_3)$, we define the *ordinal rank* α_ξ of ξ as the ordinal number given by

$$\alpha_\xi := \omega^{d-3} \cdot r_d + \omega^{d-4} \cdot r_{d-1} + \cdots + \omega^0 \cdot r_3. \tag{21}$$

Note that $\mathrm{rank}(\xi) <_{\mathrm{lex}} \mathrm{rank}(\xi')$ if and only if $\alpha_\xi < \alpha_{\xi'}$. With this reformulation, we now focus on the decreasing chain of ordinal ranks.

Let $\alpha < \omega^\omega$ be an ordinal given in Cantor Normal Form as $\alpha = \omega^n \cdot c_n + \cdots + \omega^0 \cdot c_0$ where $n, c_0, \ldots, c_n \in \mathbb{N}$, we define the *size* of α as $N\alpha := \max\{n, \max_{0 \leq i \leq n} c_i\}$. For the linear KLM sequence ξ with $\mathrm{rank}(\xi) = (r_d, \ldots, r_3)$, we have $N\alpha_\xi = \max\{d-3, \max_{3 \leq i \leq d} r_i\} \leq |\xi|$.

Given a number $n_0 \in \mathbb{N}$ and a function $h : \mathbb{N} \to \mathbb{N}$ that is monotone inflationary (that is, $x \leq h(x)$ and $h(x) \leq h(y)$ whenever $x \leq y$), we say a sequence of ordinals $\alpha_0, \alpha_1, \ldots$ is (n_0, h)-*controlled* if $N\alpha_i \leq h^i(n_0)$ for all $i \in \mathbb{N}$, where $h^i(n_0)$ is the ith iteration of h on n_0.

Let $\xi^0, \xi^1, \ldots$ be the linear KLM sequences produced in the modified KLMST algorithm, by Lemma 5.16 we know that the sequence of ordinal ranks

$$\alpha_{\xi^0} > \alpha_{\xi^1} > \cdots \tag{22}$$

is $(|\xi^0|, h)$-controlled where h is defined in Lemma 5.16. Recall that $\xi^0 \in \mathrm{clean}(\xi)$ where $\xi = \langle p(\boldsymbol{m})Gq(\boldsymbol{n}) \rangle$ corresponds to the input reachability instance. Then $|\xi^0| \leq g(|\xi|)$ where g is defined in Lemma 5.11, and (22) is indeed $(g(n), h)$-controlled where $n := |\langle p(\boldsymbol{m})Gq(\boldsymbol{n}) \rangle|$.

Length function theorem.

The length of the controlled sequence of ordinals (22) can be bounded in terms of the hierarchies of fast-growing functions of Hardy and Cichoń [4]. First recall that given a limit ordinal $\lambda \leq \omega^\omega$, the standard *fundamental sequence* of λ is a sequence $(\lambda(x))_{x<\omega}$ defined inductively by

$$\omega^\omega(x) := \omega^{x+1}, \qquad (\beta + \omega^{k+1})(x) := \beta + \omega^k \cdot (x+1) \tag{23}$$

where $\beta + \omega^{k+1}$ is in Cantor Normal Form. Now given a function $h : \mathbb{N} \to \mathbb{N}$ that is monotone inflationary, we define the *Hardy hierarchy* $(h^\alpha)_{\alpha \leq \omega^\omega}$ and the *Cichoń hierarchy* $(h_\alpha)_{\alpha \leq \omega^\omega}$ as two families of functions $h^\alpha, h_\alpha : \mathbb{N} \to \mathbb{N}$ indexed by ordinals $\alpha \leq \omega^\omega$ given inductively by

$$h^0(x) := x, \qquad h^{\alpha+1}(x) := h^\alpha(h(x)), \qquad h^\lambda(x) := h^{\lambda(x)}(x), \tag{24}$$

$$h_0(x) := 0, \qquad h_{\alpha+1}(x) := 1 + h_\alpha(h(x)), \qquad h_\lambda(x) := h_{\lambda(x)}(x). \tag{25}$$

Observe that Cichoń hierarchy counts the number of iterations of h in Hardy hierarchy, that is, $h^{h_\alpha(x)}(x) = h^\alpha(x)$. Also note that as h is monotone inflationary, by induction on α we have $h_\alpha(x) \leq h^\alpha(x)$. Now we give the length function theorem as follows.

▶ **Theorem 6.1** (Length function theorem, [20, Thm. 3.3]). *Let $n_0 \geq d - 2$, then the maximal length of (n_0, h)-controlled decreasing sequences of ordinals in ω^{d-2} is $h_{\omega^{d-2}}(n_0)$.*

Small witness property.

By Theorem 6.1 we can bound the length of (22), which then yields a bound on the minimal length of runs witnessing reachability.

▶ **Lemma 6.2** (Small witnesses). *Let $G = (Q, T)$ be a d-VASS where $d \geq 3$, let $p(\boldsymbol{m}), q(\boldsymbol{n}) \in Q \times \mathbb{N}^d$ be two configurations, and let $n := |\langle p(\boldsymbol{m})Gq(\boldsymbol{n})\rangle|$. If $p(\boldsymbol{m}) \xrightarrow{*} q(\boldsymbol{n})$ holds in G, then there is a path σ such that $p(\boldsymbol{m}) \xrightarrow{\sigma} q(\boldsymbol{n})$ and $|\sigma| \leq \ell(h^{\omega^{d-2}}(g(n)))$, where g, h, ℓ are defined in lemmas 5.11, 5.16, and 5.17.*

Proof. Suppose $p(\boldsymbol{m}) \xrightarrow{*} q(\boldsymbol{n})$, then there is a sequence of linear KLM sequence $\xi^0, \xi^1, \ldots, \xi^L$ produced in the modified KLMST algorithm, such that ξ^L is normal. We have discussed that the sequence of their ordinal ranks $\alpha_{\xi^0} > \alpha_{\xi^1} > \cdots > \alpha_{\xi^L}$ is $(g(n), h)$-controlled, so by Theorem 6.1 we have $L \leq h_{\omega^{d-2}}(g(n))$. From Lemma 5.16 and the fact that $h^{h_\alpha(x)} = h^\alpha(x)$, the size of ξ^L is bounded by

$$|\xi^L| \leq h^L(|\xi^0|) \leq h^{h_{\omega^{d-2}}(g(n))}(g(n)) = h^{\omega^{d-2}}(g(n)). \tag{26}$$

Now Lemma 5.17 bounds the length of the minimal witnesses by $\ell(h^{\omega^{d-2}}(g(n)))$. ◀

6.2 Fast-Growing Complexity Hierarchy

We recall the fast-growing hierarchy formally introduced by Schmitz [21] that captures the complexity class high above elementary. Define $H(x) := x + 1$, we shall use the Hardy hierarchy $(H^\alpha)_{\alpha < \omega^\omega}$, where for example $H^{\omega^2}(x) = 2^{x+1}(x + 1)$ and $H^{\omega^3}(x)$ grows faster than the tower function. First we define the family $\mathcal{F}_\alpha := \bigcup_{\beta < \omega^\alpha} \mathsf{FDTIME}(H^\beta(n))$ which contains functions computable in deterministic time $O(H^\beta(n))$. Observe that, for example, $\mathcal{F}_3$ contains exactly the Kalmar elementary functions. Now we define

$$\mathsf{F}_\alpha := \bigcup_{p \in \mathcal{F}_\alpha} \mathsf{DTIME}(H^{\omega^\alpha}(p(n))) \tag{27}$$

which is the class of decision problems solvable in deterministic time $O(H^{\omega^\alpha}(p(n)))$. Note that non-deterministic time Turing machines can be made deterministic with an exponential overhead in $\mathcal{F}_3$, thus for $\alpha \geq 3$, we have equivalently that $\mathsf{F}_\alpha = \bigcup_{p \in \mathcal{F}_\alpha} \mathsf{NDTIME}(H^{\omega^\alpha}(p(n)))$. Observe that F_α is closed under reductions in $\mathcal{F}_\alpha$.

6.2.1 Relativized Fast-Growing Functions

In order to express the complexity of the modified KLMST algorithm in terms of the hierarchy $(\mathsf{F}_\alpha)_{\alpha < \omega}$, one needs to locate the function $h^{\omega^{d-2}}$ in the Hardy hierarchy $(H^\alpha)_{\alpha < \omega^\omega}$ where $h \in \mathcal{F}_3$ is the elementary function from Lemma 5.16. Previously we can upper bound $h^{\omega^{d-2}}$ by $H^{\omega^{d+1}}$ with the help of [21, Lem. 4.2]. Here we show a slightly better result, from which we can bound $h^{\omega^{d-2}}(x)$ by $H^{\omega^d}(O(x))$.

▶ **Lemma 6.3** (cf. [21, Lem. A.5]). *Let $h : \mathbb{N} \to \mathbb{N}$ be a monotone inflationary function, let $a, b, c \geq 1$ and $x_0 \geq 0$ be natural numbers. If for all $x \geq x_0$, $h(x) \leq H^{\omega^b \cdot c}(x)$, then $h^{\omega^a}(x) \leq H^{\omega^{b+a}}((c+1)x)$ for all $x \geq \max\{2c, x_0\}$.*

6.3   Upper Bounds for VASS Reachability

Now we analyze the time complexity of the modified KLMST algorithm. Given as input the d-VASS $G = (Q, T)$ and two configurations $p(\boldsymbol{m}), q(\boldsymbol{n})$, let $\xi := \langle p(\boldsymbol{m})Gq(\boldsymbol{n})\rangle$ and $n := |\xi|$. The initial sequence $\xi^0 \in \mathrm{clean}(\xi)$ can be computed in (non-deterministic) time elementary in n by Lemma 5.11. Then the algorithm produces $\xi^0, \xi^1, \ldots, \xi^L$ with $L \le h_{\omega^{d-2}}(g(n))$, where g, h are defined in lemmas 5.11, 5.16. Note that in each step, the sequence $\xi^{i+1} \in \mathrm{dec}(\xi^i)$ can be computed in time elementary in $|\xi^i|$ by Lemma 5.16, and the sizes $|\xi^i|$ are all bounded by $h^{\omega^{d-2}}(g(n))$ as we have discussed above in the proof of Lemma 6.2. To sum up, the entire algorithm finishes in non-deterministic time elementary in $h^{\omega^{d-2}}(g(n))$.

▶ **Lemma 6.4.** *On input a d-VASS $G = (Q, T)$ where $d \ge 3$ and $p(\boldsymbol{m}), q(\boldsymbol{n}) \in Q \times \mathbb{N}^d$, the modified KLMST algorithm finishes in non-deterministic time $e(h^{\omega^{d-2}}(g(n)))$ where $n = |\langle p(\boldsymbol{m})Gq(\boldsymbol{n})\rangle|$, g, h are defined in lemmas 5.11, 5.16, and $e \in \mathcal{F}_3$ is some fixed function.*

Since h is an elementary function, there is a number $c \in \mathbb{N}$ such that h is eventually dominated by $H^{\omega^2 \cdot c}$. By Lemma 6.3 we can upper bound $h^{\omega^{d-2}}(x)$ by $H^{\omega^d}((c+1)x)$. Observe that the inner part $g(n)$ is elementary in the binary encoding size of the input $G, p(\boldsymbol{m}), q(\boldsymbol{n})$, thus can be captured by a function $p \in \mathcal{F}_3$. Finally, [21, Lem. 4.6] allows us to move the outermost function e to the innermost position. Hence we have the following upper bound.

▶ **Theorem 6.5.** *Reachability in d-dimensional VASS is in F_d for all $d \ge 3$.*

Also, by Lemma 6.2 there is a simple combinatorial algorithm for d-VASS reachability. We fist compute the bound $B := \ell(h^{\omega^{d-2}}(g(n)))$, which can be done in time elementary in B by [21, Thm. 5.1]. Then we can decide reachability by just enumerate all possible paths in G with length bounded by B.

7   Conclusion

We have shown that the reachability problem in d-dimensional vector addition system with states is in F_d, improving the previous F_{d+4} upper bound by Leroux and Schmitz [15]. By capturing reachability in geometrically 2-dimensional VASSes with linear path schemes, we are able to reduce significantly the number of decomposition steps in the KLMST decomposition algorithm. Combined with a careful analysis on fast-growing functions, we finally obtained the F_d upper bound. It should be noticed though, that our algorithm avoids computing the "full decomposition" [14] of KLM sequences, thus cannot improve the complexity of problems that essentially rely on the full decomposition, e.g., the VASS downward language inclusion problem [8, 23, 15].

It has been shown that the reachability problem in $(2d+3)$-VASS is F_d-hard [5]. In the case of 3-VASS, it is known that the reachability problem is PSPACE-hard. The gap between the lower bound and the upper bound $\mathsf{F}_3 = \mathsf{TOWER}$ [22] is huge. It is very unlikely that the problem is PSPACE-complete. Effort to uplift the lower bound is called for.

───── **References** ─────

1   Michael Blondin, Matthias Englert, Alain Finkel, Stefan Göller, Christoph Haase, Ranko Lazic, Pierre McKenzie, and Patrick Totzke. The reachability problem for two-dimensional vector addition systems with states. *J. ACM*, 68(5):34:1–34:43, 2021. `doi:10.1145/3464794`.

2   Michael Blondin, Alain Finkel, Stefan Göller, Christoph Haase, and Pierre McKenzie. Reachability in two-dimensional vector addition systems with states is pspace-complete. In *30th*

Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2015, Kyoto, Japan, July 6-10, 2015, pages 32–43. IEEE Computer Society, 2015. `doi:10.1109/LICS.2015.14`.

3 Dmitry Chistikov and Christoph Haase. The taming of the semi-linear set. In Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi, editors, *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*, volume 55 of *LIPIcs*, pages 128:1–128:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. URL: `https://doi.org/10.4230/LIPIcs.ICALP.2016.128`, `doi:10.4230/LIPICS.ICALP.2016.128`.

4 E. A. Cichon and Elias Tahhan-Bittar. Ordinal recursive bounds for higman's theorem. *Theor. Comput. Sci.*, 201(1-2):63–84, 1998. `doi:10.1016/S0304-3975(97)00009-1`.

5 Wojciech Czerwinski, Ismaël Jecker, Slawomir Lasota, Jérôme Leroux, and Lukasz Orlikowski. New lower bounds for reachability in vector addition systems. In Patricia Bouyer and Srikanth Srinivasan, editors, *43rd IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2023, December 18-20, 2023, IIIT Hyderabad, Telangana, India*, volume 284 of *LIPIcs*, pages 35:1–35:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023. URL: `https://doi.org/10.4230/LIPIcs.FSTTCS.2023.35`, `doi:10.4230/LIPICS.FSTTCS.2023.35`.

6 Wojciech Czerwinski and Lukasz Orlikowski. Reachability in vector addition systems is ackermann-complete. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 1229–1240. IEEE, 2021. `doi:10.1109/FOCS52979.2021.00120`.

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

8 Peter Habermehl, Roland Meyer, and Harro Wimmel. The downward-closure of petri net languages. In Samson Abramsky, Cyril Gavoille, Claude Kirchner, Friedhelm Meyer auf der Heide, and Paul G. Spirakis, editors, *Automata, Languages and Programming, 37th International Colloquium, ICALP 2010, Bordeaux, France, July 6-10, 2010, Proceedings, Part II*, volume 6199 of *Lecture Notes in Computer Science*, pages 466–477. Springer, 2010. `doi:10.1007/978-3-642-14162-1\_39`.

9 Thomas Jech. *Set Theory*. Springer Monographs in Mathematics. Springer, 2003. `doi:10.1007/3-540-44761-X`.

10 S. Rao Kosaraju. Decidability of reachability in vector addition systems (preliminary version). In *Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing*, STOC '82, page 267–281, New York, NY, USA, 1982. Association for Computing Machinery. `doi:10.1145/800070.802201`.

11 Jean-Luc Lambert. A structure to decide reachability in petri nets. *Theor. Comput. Sci.*, 99(1):79–104, 1992. `doi:10.1016/0304-3975(92)90173-D`.

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

13 Jérôme Leroux. The reachability problem for petri nets is not primitive recursive. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 1241–1252. IEEE, 2021. `doi:10.1109/FOCS52979.2021.00121`.

14 Jérôme Leroux and Sylvain Schmitz. Demystifying reachability in vector addition systems. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2015, Kyoto, Japan, July 6-10, 2015*, pages 56–67. IEEE Computer Society, 2015. `doi:10.1109/LICS.2015.16`.

15 Jérôme Leroux and Sylvain Schmitz. Reachability in vector addition systems is primitive-recursive in fixed dimension. In *34th Annual ACM/IEEE Symposium on Logic in Computer*

Science, LICS 2019, Vancouver, BC, Canada, June 24-27, 2019, pages 1–13. IEEE, 2019. `doi:10.1109/LICS.2019.8785796`.

16   Jérôme Leroux and Grégoire Sutre. On flatness for 2-dimensional vector addition systems with states. In Philippa Gardner and Nobuko Yoshida, editors, *CONCUR 2004 - Concurrency Theory, 15th International Conference, London, UK, August 31 - September 3, 2004, Proceedings*, volume 3170 of *Lecture Notes in Computer Science*, pages 402–416. Springer, 2004. `doi:10.1007/978-3-540-28644-8\_26`.

17   Ernst W. Mayr. An algorithm for the general petri net reachability problem. In *Proceedings of the Thirteenth Annual ACM Symposium on Theory of Computing*, STOC '81, page 238–246, New York, NY, USA, 1981. Association for Computing Machinery. `doi:10.1145/800076.802477`.

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

19   Charles Rackoff. The covering and boundedness problems for vector addition systems. *Theor. Comput. Sci.*, 6:223–231, 1978. `doi:10.1016/0304-3975(78)90036-1`.

20   Sylvain Schmitz. Complexity bounds for ordinal-based termination - (invited talk). In Joël Ouaknine, Igor Potapov, and James Worrell, editors, *Reachability Problems - 8th International Workshop, RP 2014, Oxford, UK, September 22-24, 2014. Proceedings*, volume 8762 of *Lecture Notes in Computer Science*, pages 1–19. Springer, 2014. `doi:10.1007/978-3-319-11439-2\_1`.

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

22   Qizhe Yang and Yuxi Fu. Reachability in 3-vass is in tower. *CoRR*, abs/2306.05710, 2023. URL: `https://doi.org/10.48550/arXiv.2306.05710`, `arXiv:2306.05710`, `doi:10.48550/ARXIV.2306.05710`.

23   Georg Zetzsche. The complexity of downward closure comparisons. In Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi, editors, *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*, volume 55 of *LIPIcs*, pages 123:1–123:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. URL: `https://doi.org/10.4230/LIPIcs.ICALP.2016.123`, `doi:10.4230/LIPICS.ICALP.2016.123`.

A   Proof of Theorem 3.4

In this appendix we prove the main result of Section 3.

▶ **Theorem 3.4.** *Let $G = (Q, T)$ be a geometrically 2-dimensional d-VASS. For every pair of configurations $p(\boldsymbol{u}), q(\boldsymbol{v}) \in Q \times \mathbb{N}^d$ with $p(\boldsymbol{u}) \xrightarrow{*} q(\boldsymbol{v})$ there exists an LPS Λ compatible to G such that $p(\boldsymbol{u}) \xrightarrow{\Lambda} q(\boldsymbol{v})$ and $|\Lambda| \leq |G|^{O(1)}$.*

First note that for a geometrically 2-dimensional d-VASS G, we can safely assume that $\dim(\mathrm{Cyc}(G)) = 2$ in the proof of Theorem 3.4. If it is the case that $\dim(\mathrm{Cyc}(G)) < 2$, we just add an isolated state $q_\perp$ with self-loops on it to enlarge $\mathrm{Cyc}(G)$. Note that $q_\perp$ will never be used in any LPS capturing reachability relation in the original G.

We will prove Theorem 3.4 by induction on d. The base case $d = 2$ is already studied in [1] and will be reviewed in the section A.2. First let's recall a special kind of linear path schemes, called *zigzag-free* LPSes.

A.1   Zigzag-free Linear Path Schemes

An *orthant* is one of the 2^d regions in $\mathbb{Q}^d$ split by the d coordinate axes. Formally speaking, given a tuple $t \in \{+1, -1\}^d$, the orthant Z_t defined by t is the following set:

$$Z_t := \{\boldsymbol{u} \in \mathbb{Q}^d : \boldsymbol{u}(i) \cdot t(i) \geq 0 \text{ for all } i \in [d]\}. \tag{28}$$

An LPS $\Lambda = \alpha_0 \beta_1^* \alpha_1 \ldots \beta_k^* \alpha_k$ is said to be *zigzag-free in the orthant Z_t* if $\Delta(\beta_j) \in Z_t$ for all $j = 1, \ldots, k$. We say Λ is *zigzag-free* if it is zigzag-free in some orthant. Zigzag free linear path schemes can cast $\mathbb{Z}^d$-reachability to $\mathbb{N}^d$-reachability when the start and end configurations are sufficiently high:

▶ **Lemma A.1** ([1, lemma 4.7] and [16, lemma 4.6])**.** *Let (G, Λ) be a zigzag-free LPS where $G = (Q, T)$ is a d-VASS. Let $p(\boldsymbol{u}), q(\boldsymbol{v}) \in Q \times \mathbb{N}^d$ be two configurations of G. If $p(\boldsymbol{u}) \xrightarrow{\Lambda}_{\mathbb{Z}^d} q(\boldsymbol{v})$ and $\boldsymbol{u}(i), \boldsymbol{v}(i) \geq |\Lambda| \cdot \|T\|$ for all $i \in [d]$, then $p(\boldsymbol{u}) \xrightarrow{\Lambda}_{\mathbb{N}^d} q(\boldsymbol{v})$.*

A.2   Review of Known Results in Dimension 2

The following is the main result of [1] and is the base case of Theorem 3.4:

▶ **Theorem A.2** ([1, Thm. 3.1])**.** *Let $G = (Q, T)$ be a 2-VASS. For every pair of configurations $p(\boldsymbol{u}), q(\boldsymbol{v}) \in Q \times \mathbb{N}^2$ with $p(\boldsymbol{u}) \xrightarrow{*} q(\boldsymbol{v})$ there exists an LPS Λ compatible to G such that $p(\boldsymbol{u}) \xrightarrow{\Lambda} q(\boldsymbol{v})$ and $|\Lambda| \leq |G|^{O(1)}$.*

We also need some other lemmas from [1] to handle cyclic runs far away from axes.

▶ **Lemma A.3** ([1, Prop. 4.2])**.** *For every 2-VASS $G = (Q, T)$ there exists a natural number $D \leq |G|^{O(1)}$ such that for any state $q \in Q$ and any $\boldsymbol{u}, \boldsymbol{v} \in [D, \infty)^2$ with $q(\boldsymbol{u}) \xrightarrow{*}_{\mathbb{N}^2} q(\boldsymbol{v})$ there exists an LPS Λ compatible to G such that $q(\boldsymbol{u}) \xrightarrow{\Lambda}_{\mathbb{N}^2} q(\boldsymbol{v})$ and $|\Lambda| \leq |G|^{O(1)}$.*

We can further require that the above LPS Λ is zigzag-free.

▶ **Lemma A.4** ([1, Lem. 4.15])**.** *For every 2-VASS $G = (Q, T)$, every $q \in Q$, and every LPS σ from q to q compatible to G, there exists a finite set S of zigzag-free LPSes from q to q compatible to G such that $\Delta(\sigma) \subseteq \Delta(S)$, and $|\Lambda| \leq (|\sigma| + \|T\|)^{O(1)}$ for every $\Lambda \in S$.*

If the effects of cycles in a zigzag-free LPS Λ capturing $q(\boldsymbol{u}) \xrightarrow{*} q(\boldsymbol{v})$ does not belong to any orthant containing $\boldsymbol{v} - \boldsymbol{u}$, they can only be fired for a bounded number of times. We can remove those cycles by expanding them.

▶ **Lemma A.5.** *Let $G = (Q, T)$ be a 2-VASS, let $q \in Q$ and $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{N}^2$. Let Z be any orthant containing $\boldsymbol{v} - \boldsymbol{u}$. For every zigzag-free linear path scheme σ compatible to G such that $q(\boldsymbol{u}) \xrightarrow{\sigma} q(\boldsymbol{v})$, there exists an LPS Λ from q to q compatible to G that is zigzag-free in Z, and such that $q(\boldsymbol{u}) \xrightarrow{\Lambda} q(\boldsymbol{v})$, and $|\Lambda| \leq (|\sigma| + \|T\|)^{O(1)}$.*

Proof. Let $\sigma = \alpha_0 \beta_1^* \alpha_1 \ldots \beta_k^* \alpha_k$ be a zigzag-free LPS such that $q(\boldsymbol{u}) \xrightarrow{\sigma} q(\boldsymbol{v})$. There exists numbers $e_1, \ldots, e_k \in \mathbb{N}$ such that $q(\boldsymbol{u}) \xrightarrow{\alpha_0 \beta_1^{e_1} \alpha_1 \ldots \beta_k^{e_k} \alpha_k} q(\boldsymbol{v})$. In particular,

$$\boldsymbol{v} - \boldsymbol{u} = \Delta(\alpha_0 \ldots \alpha_k) + \sum_{i=1}^{k} e_i \cdot \Delta(\beta_i). \tag{29}$$

Let Z be an orthant containing $\boldsymbol{v} - \boldsymbol{u}$. Assume w.l.o.g. that $Z = \mathbb{Q}^2$ is the first orthant. We claim that for all j such that $\Delta(\beta_j) \notin Z$, $e_j \leq |\sigma| \cdot \|T\|$. Indeed, since $\Delta(\beta_j) \notin Z$ and $\boldsymbol{v} - \boldsymbol{u} \in Z$, there must be an index $\ell \in \{1, 2\}$ such that $(\boldsymbol{v} - \boldsymbol{u})(\ell) \cdot \Delta(\beta_j)(\ell) \leq 0$ and $\Delta(\beta_j)(\ell) \neq 0$. We assume w.l.o.g. that $(\boldsymbol{v} - \boldsymbol{u})(\ell) \geq 0$ and $\Delta(\beta_j)(\ell) < 0$. Then $\Delta(\beta_i)(\ell) \leq 0$ for all $i \in [k]$ by the zigzag-free property of σ. In particular, $\Delta(\beta_j)(\ell) \leq -1$. From the ℓ-th component of equation (29),

$$0 \leq (\boldsymbol{v} - \boldsymbol{u})(\ell) = \Delta(\alpha_0 \ldots \alpha_k)(\ell) + \sum_{i=1}^{k} e_i \cdot \Delta(\beta_i)(\ell) \leq |\sigma| \cdot \|T\| - e_j \tag{30}$$

we have $e_j \leq |\sigma| \cdot \|T\|$ as we claimed. The proof is completed by taking Λ to be the LPS obtained from σ by replacing all cycles β_j^* where $\Delta(\beta_j) \notin Z$ with the path $\beta_j^{e_j}$. ◀

We combine Lemmas A.3, A.4, A.5 to get the following one:

▶ **Corollary A.6.** *For every 2-VASS $G = (Q, T)$ there exists a natural number $D \leq |G|^{O(1)}$ such that, for any state $q \in Q$, any $\boldsymbol{u}, \boldsymbol{v} \in [D, \infty)^2$ with $q(\boldsymbol{u}) \xrightarrow{*}_{\mathbb{N}^2} q(\boldsymbol{v})$ and for any orthant Z containing $\boldsymbol{v} - \boldsymbol{u}$, there exists an LPS Λ compatible to G that is zigzag-free in Z such that $q(\boldsymbol{u}) \xrightarrow{\Lambda}_{\mathbb{N}^2} q(\boldsymbol{v})$ and $|\Lambda| \leq |G|^{O(1)}$.*

Proof. Let D_0 be the parameter D given by Lemma A.3 and $B \leq |G|^{O(1)}$ be the bound on $|\Lambda|$ composed from Lemmas A.3 and A.4. We put $D = \max\{D_0, B \cdot \|T\|\} \leq |G|^{O(1)}$ here.

Fix $q \in Q$ and $\boldsymbol{u}, \boldsymbol{v} \in [D, \infty)^2$ with $q(\boldsymbol{u}) \xrightarrow{*}_{\mathbb{N}^2} q(\boldsymbol{v})$. By Lemma A.3 there is an LPS σ compatible to G such that $q(\boldsymbol{u}) \xrightarrow{\sigma}_{\mathbb{N}^2} q(\boldsymbol{v})$. Apply Lemma A.4 to σ, we get a zigzag-free LPS σ' compatible to G with $q(\boldsymbol{u}) \xrightarrow{\sigma'}_{\mathbb{Z}^2} q(\boldsymbol{v})$ and $|\sigma'| \leq B$. Lemma A.1 and our choice of D then show that indeed we have $q(\boldsymbol{u}) \xrightarrow{\sigma'}_{\mathbb{N}^2} q(\boldsymbol{v})$. Let Z be an orthant containing $\boldsymbol{v} - \boldsymbol{u}$. We finally apply Lemma A.5 to σ' to obtain an LPS Λ compatible to G that is zigzag-free in Z such that $q(\boldsymbol{u}) \xrightarrow{\Lambda}_{\mathbb{N}^2} q(\boldsymbol{v})$ and $|\Lambda| \leq (B + \|T\|)^{O(1)} \leq |G|^{O(1)}$. ◀

A.3   Sign Reflecting Projection

In this section we develop a technical tool called the "sign reflecting projection" which will be used to project a geometrically 2-dimensional d-VASS to a 2-VASS so that cyclic runs with effect belonging to a certain orthant can be safely projected back. We first introduce the notion of projection of vectors.

Projection.

Let $I \subseteq [d]$ be a subset of indices. Given a vector $\boldsymbol{u} \in \mathbb{X}^d$ where $\mathbb{X}$ can be one of $\mathbb{Q}, \mathbb{N}, \mathbb{Z}$, we define the *projection of $\boldsymbol{u}$ onto indices in I* as a function $\boldsymbol{u}|_I \in \mathbb{X}^I$ given by

$$(\boldsymbol{u}|_I)(i) = \boldsymbol{u}(i) \text{ for all } i \in I. \tag{31}$$

We tacitly identify the function $\boldsymbol{u}|_I \in \mathbb{X}^I$ as a vector in $\mathbb{X}^{|I|}$. Let $i \in [d]$, we abbreviate $\boldsymbol{u}^{-i}$ for the vector $\boldsymbol{u}|_{[d]\setminus\{i\}}$. The definition of projection is naturally extended to set of vectors V:

$$V|_I := \big\{ \boldsymbol{v}|_I : \boldsymbol{v} \in V \big\}. \tag{32}$$

It should be clear that the projection of a vector space onto indices in I is again a vector space in $\mathbb{Q}^{|I|}$, and the projection of an orthant Z_t onto I is an orthant in $\mathbb{Q}^{|I|}$ defined by $t|_I$.

Sign reflecting projection.

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

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

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

Proof. Suppose otherwise, there are vectors $\boldsymbol{v} \neq \boldsymbol{v}' \in P \cap Z$ with $\boldsymbol{v}|_I = \boldsymbol{v}'|_I$. Then there must be an index $j \in [d]$ with $\boldsymbol{v}(j) \neq \boldsymbol{v}'(j)$. Note that $\boldsymbol{v} - \boldsymbol{v}' \in P$ and $\boldsymbol{v}' - \boldsymbol{v} \in P$ cannot belong to the same orthant, as the signs of their jth component differ. On the other hand, $(\boldsymbol{v} - \boldsymbol{v}')|_I = (\boldsymbol{v}' - \boldsymbol{v})|_I = \boldsymbol{0} \in Z|_I$. Since I is a sign-reflecting projection, we must have $\boldsymbol{v} - \boldsymbol{v}' \in Z$ and $\boldsymbol{v}' - \boldsymbol{v} \in Z$, a contradiction. ◀

We mainly care about finding sign-reflecting projections for a plane P, that is, a 2-dimensional subspace of $\mathbb{Q}^d$. The main technical theorem is stated as follows.

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

The proof of A.9 makes use of the Farkas-Minkowski-Weyl's Theorem. We first introduce some additional definitions about *cones*.

A *polyhedral cone* is a set of the form $\{\boldsymbol{v} \in \mathbb{Q}^d : A\boldsymbol{v} \leq 0\}$ where A is a rational matrix. Let $V \subseteq \mathbb{Q}^d$ be a set of vectors, the *cone generated by V* is the following set:

$$\text{cone}(V) := \bigcup_{k \in \mathbb{N}} \{\lambda_1 x_1 + \cdots + \lambda_k x_k : \lambda_1, \ldots, \lambda_k \in \mathbb{Q}_{\geq 0}, x_1, \ldots, x_k \in V\}. \tag{33}$$

A *finitely generated cone* is a cone generated by a finite set. The Farkas-Minkowski-Weyl's Theorem states that:

▶ **Theorem A.10** (Farkas-Minkowski-Weyl). *A set $C \subseteq \mathbb{Q}^d$ is a polyhedral cone if and only if it is a finitely generated cone.*

Let P be a plane and Z be an orthant in $\mathbb{Q}^d$, one easily verifies that $P \cap Z$ is a polyhedral cone, where the inequality $A\boldsymbol{v} \leq 0$ asserts that (i) $\boldsymbol{v}$ is orthogonal to all normal vectors of P, and (ii) $\boldsymbol{v}$ belongs to the orthant Z. Theorem A.10 then shows that the intersection of a plane and an orthant is a finitely generated cone. We need a stronger fact that it is generated by at most 2 vectors.

▶ **Lemma A.11.** *Let $P \subseteq \mathbb{Q}^d$ be a plane and Z be an orthant in $\mathbb{Q}^d$. Then there exists $V \subseteq \mathbb{Q}^d$ such that $P \cap Z = \mathrm{cone}(V)$ and $|V| \leq 2$.*

Proof. Let P be a plane and Z be an orthant in $\mathbb{Q}^d$. The existence of a finite set $V \subseteq \mathbb{Q}^d$ such that $P \cap Z = \mathrm{cone}(V)$ is guaranteed by Theorem A.10. Let V have minimal size. Towards a contradiction we assume that V contains 3 distinct non-zero vectors $\boldsymbol{u}, \boldsymbol{v}, \boldsymbol{w}$. These vectors must be linearly dependent since $V \subseteq P$ and P is 2-dimensional. Suppose $\boldsymbol{w} = \lambda_1 \boldsymbol{u} + \lambda_2 \boldsymbol{v}$. We can make sure that λ_1 and λ_2 are of the same sign since otherwise we just rename $\boldsymbol{u}, \boldsymbol{v}, \boldsymbol{w}$ properly. If $\lambda_1, \lambda_2 \geq 0$ then $\mathrm{cone}(V) = \mathrm{cone}(V \setminus \{\boldsymbol{w}\})$, contradicting to the minimality of $|V|$. Thus $\lambda_1, \lambda_2 \leq 0$. But then $\boldsymbol{u}, \boldsymbol{v}$ and $\boldsymbol{w}$ cannot belong to the same orthant, contradicting to $V \subseteq Z$. Therefore, we must have $|V| \leq 2$. ◀

We can now prove Theorem A.9.

Proof of Theorem A.9. Let $P \subseteq \mathbb{Q}^d$ be a plane, and Z be an orthant in $\mathbb{Q}^d$ such that $P \cap Z$ contains two linearly independent vectors. Without loss of generality we assume that $Z = \mathbb{Q}_{\geq 0}$ is the first orthant. By Lemma A.11 there are non-zero vectors $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{Q}^d$ such that $P \cap Z = \mathrm{cone}(\{\boldsymbol{u}, \boldsymbol{v}\})$. We emphasize that $\boldsymbol{u}, \boldsymbol{v} \geq \boldsymbol{0}$.

Clearly $\boldsymbol{u}, \boldsymbol{v}$ are linearly independent since the cone generated by them contains two linearly independent vectors. We claim that $\mathrm{supp}(\boldsymbol{u})$ and $\mathrm{supp}(\boldsymbol{v})$ are incomparable under $\subseteq$. Otherwise, suppose $\mathrm{supp}(\boldsymbol{v}) \subseteq \mathrm{supp}(\boldsymbol{u})$, then $\lambda \boldsymbol{u} - \boldsymbol{v} \in P \cap Z$ for sufficiently large $\lambda > 0$. As $P \cap Z = \mathrm{cone}(\{\boldsymbol{u}, \boldsymbol{v}\})$, we have $\lambda \boldsymbol{u} - \boldsymbol{v} = \alpha \boldsymbol{u} + \beta \boldsymbol{v}$ for some $\alpha, \beta \geq 0$. Then $(\beta + 1)\boldsymbol{v} = (\lambda - \alpha)\boldsymbol{u}$, contradicting to the linear independence of $\boldsymbol{u}$ and $\boldsymbol{v}$. Symmetrically, we cannot have $\mathrm{supp}(\boldsymbol{u}) \subseteq \mathrm{supp}(\boldsymbol{v})$. Now take $i \in \mathrm{supp}(\boldsymbol{u}) \backslash \mathrm{supp}(\boldsymbol{v})$ and $j \in \mathrm{supp}(\boldsymbol{v}) \backslash \mathrm{supp}(\boldsymbol{u})$. We show that $I = \{i, j\}$ is a sign-reflecting projection for P w.r.t. Z. Indeed, let $\boldsymbol{w} \in P = \mathrm{span}(\{\boldsymbol{u}, \boldsymbol{v}\})$ be such that $\boldsymbol{w}(i), \boldsymbol{w}(j) \geq 0$. Suppose $\boldsymbol{w} = \alpha \boldsymbol{u} + \beta \boldsymbol{v}$ for some $\alpha, \beta \in \mathbb{Q}$. Then $\boldsymbol{w}(i) = \alpha \boldsymbol{u}(i) \geq 0$ implies $\alpha \geq 0$, and $\boldsymbol{w}(j) = \beta \boldsymbol{v}(j) \geq 0$ implies $\beta \geq 0$. Thus $\boldsymbol{w} \in \mathrm{cone}(\{\boldsymbol{u}, \boldsymbol{v}\}) \subseteq Z$. ◀

Finally, we remark that the requirement in Theorem A.9 that $P \cap Z$ contains two linearly independent vectors is not essential:

▶ **Lemma A.12.** *Let $P \subseteq \mathbb{Q}^d$ be a plane, and $\boldsymbol{v} \in P$ be a non-zero vector. There exists an orthant Z such that $\boldsymbol{v} \in Z$ and $P \cap Z$ contains two linearly independent vectors.*

Proof. Let $\boldsymbol{v} \in P$ be non-zero. Pick $\boldsymbol{u} \in P$ such that $P = \mathrm{span}(\{\boldsymbol{u}, \boldsymbol{v}\})$. Choose $t > 0$ large enough such that $t|\boldsymbol{v}(i)| > |\boldsymbol{u}(i)|$ for all $i \in \mathrm{supp}(\boldsymbol{v})$. Observe that $\boldsymbol{v}$ and $t\boldsymbol{v} + \boldsymbol{u}$ are linearly independent and lie in the same orthant Z as we desired. ◀

A.4   Reachability Far Away from Axes

In this section, we prove that runs far away the axes in $\mathbb{N}^d$ can be captured by small linear path schemes. We shall focus on cycles at first:

▶ **Lemma A.13.** *For every d-VASS $G = (Q, T)$ that is geometrically 2-dimensional, there exists a natural numbers $D \leq |G|^{O(1)}$ such that for every state $q \in Q$ and $\boldsymbol{u}, \boldsymbol{v} \in [D, \infty)^d$ with $q(\boldsymbol{u}) \xrightarrow{*} q(\boldsymbol{v})$, there exists an LPS Λ compatible to G such that $q(\boldsymbol{u}) \xrightarrow{\Lambda} q(\boldsymbol{v})$ and $|\Lambda| \leq |G|^{O(1)}$.*

Proof. Let D_0 be the constants D as in Corollary A.6 and $B \leq |G|^{O(1)}$ be the upper bound of lengths of LPSes given by Corollary A.6. We put $D = \max\{D_0, B \cdot \|T\|\} \leq |G|^{O(1)}$ here.

Let $q \in Q$ and $\boldsymbol{u}, \boldsymbol{v} \in [D, \infty)^d$ be such that $q(\boldsymbol{u}) \xrightarrow{*} q(\boldsymbol{v})$. Clearly $\boldsymbol{v} - \boldsymbol{u} \in \mathrm{Cyc}(G)$. Let Z be the orthant given by Lemma A.12 such that Z contains $\boldsymbol{v} - \boldsymbol{u}$ and that $\mathrm{Cyc}(G) \cap Z$ contains two linearly independent vectors. Now Theorem A.9 shows that there exists a sign-reflecting projection $I \subseteq [d]$ for $\mathrm{Cyc}(G)$ w.r.t. Z, where $|I| = 2$.

We define the 2-VASS $G|_I = (Q, T|_I)$ as the projection of G onto indices in I, where the transition set $T|_I$ is given by

$$T|_I = \{(p, \boldsymbol{a}|_I, q) : (p, \boldsymbol{a}, q) \in T\}. \tag{34}$$

For every transition $t' = (p, \boldsymbol{a}', q) \in T|_I$ we fix $g(t')$ to be an arbitrary transition $t = (p, \boldsymbol{a}, q) \in T$ such that $\boldsymbol{a}|_I = \boldsymbol{a}'$. The mapping g is extended to paths and cycles naturally.

Note that $\boldsymbol{u}|_I, \boldsymbol{v}|_I \in [D, \infty)^2$ and $q(\boldsymbol{u}|_I) \xrightarrow{*} q(\boldsymbol{v}|_I)$ in $G|_I$. By Corollary A.6 there is an LPS $\sigma = \alpha_0 \beta_1^* \alpha_1 \ldots \beta_k^* \alpha_k$ compatible to $G|_I$ that is zigzag-free in $Z|_I$ such that $q(\boldsymbol{u}|_I) \xrightarrow{\sigma} q(\boldsymbol{v}|_I)$. Moreover, $|\sigma| \leq B \leq D$. Consider the LPS Λ given by

$$\Lambda = g(\alpha_0) g(\beta_1)^* g(\alpha_1) \ldots g(\beta_k)^* g(\alpha_k). \tag{35}$$

Clearly Λ is an LPS compatible to G from state q to q. Note that for all $i \in [k]$ we have $\Delta(g(\beta_i)) \in \mathrm{Cyc}(G)$ and $\Delta(g(\beta_i))|_I = \Delta(\beta_i) \in Z|_I$, thus $\Delta(g(\beta_i)) \in Z$ by the fact that I is sign-reflecting. We have shown that Λ is zigzag-free in Z. Next we prove that $q(\boldsymbol{u}) \xrightarrow{\Lambda} q(\boldsymbol{v})$.

As $q(\boldsymbol{u}|_I) \xrightarrow{\sigma} q(\boldsymbol{v}|_I)$, there must be numbers $e_1, \ldots, e_k \in \mathbb{N}$ such that

$$q(\boldsymbol{u}|_I) \xrightarrow{\alpha_0 \beta_1^{e_1} \alpha_1 \ldots \beta_k^{e_k} \alpha_k} q(\boldsymbol{v}|_I). \tag{36}$$

Let ρ be the path in G given by

$$\rho := g(\alpha_0) g(\beta_1)^{e_1} g(\alpha_1) \ldots g(\beta_k)^{e_k} g(\alpha_k), \tag{37}$$

We shall show that $q(\boldsymbol{u}) \xrightarrow{\rho}_{\mathbb{Z}^d} q(\boldsymbol{v})$ holds in G. Then by Lemma A.1 and our choice of D we have indeed $q(\boldsymbol{u}) \xrightarrow{\rho}_{\mathbb{N}^d} q(\boldsymbol{v})$, witnessing the fact that $q(\boldsymbol{u}) \xrightarrow{\Lambda} q(\boldsymbol{v})$. Now we only need to check that $\Delta(\rho) = \boldsymbol{v} - \boldsymbol{u}$. Recall that ρ is a cycle, so $\Delta(\rho) \in \mathrm{Cyc}(G)$. By (36) and the definition of g we have $\Delta(\rho)|_I = \boldsymbol{v}|_I - \boldsymbol{u}|_I \in Z|_I$. We then have $\Delta(\rho) \in Z$ since I is sign-reflecting. By Lemma A.8, any vector in $\mathrm{Cyc}(G) \cap Z$ is uniquely determined by its components in I. As $\Delta(\rho)|_I = (\boldsymbol{v} - \boldsymbol{u})|_I$ we must have $\Delta(\rho) = \boldsymbol{v} - \boldsymbol{u}$ as claimed.

Finally, we bound the length of Λ. Just note that $|\Lambda| = |\sigma| \leq |G|^{O(1)}$. ◀

Runs that stay far away from the axes can be decomposed into a short path interleaved by at most $|Q|$ cycles, where each cycle can be handled by Lemma A.13. This leads to the second result we will prove in this section:

▶ **Lemma A.14.** *For every d-VASS $G = (Q, T)$ that is geometrically 2-dimensional, there exists a natural number $D \leq |G|^{O(1)}$ such that for every pair of configurations $p(\boldsymbol{u}), q(\boldsymbol{v}) \in Q \times [D, \infty)^d$ with $p(\boldsymbol{u}) \xrightarrow{*}_{[D,\infty)^d} q(\boldsymbol{v})$, there exists an LPS Λ compatible to G such that $p(\boldsymbol{u}) \xrightarrow{\Lambda} q(\boldsymbol{v})$ and $|\Lambda| \leq |G|^{O(1)}$.*

Proof. We put D to be the same constant as in Lemma A.13. Denote $\mathbb{O} := [D, \infty)^d$. Let $p(\boldsymbol{u}), q(\boldsymbol{v}) \in Q \times \mathbb{O}$ be such that $p(\boldsymbol{u}) \xrightarrow{*}_{\mathbb{O}} q(\boldsymbol{v})$.

Let π be a run witnessing $p(\boldsymbol{u}) \xrightarrow{*}_{\mathbb{O}} q(\boldsymbol{v})$. By extracting maximal cycles from π, we can write π in the following form:

$$\pi : p(\boldsymbol{u}) \xrightarrow{\alpha_0}_{\mathbb{O}} q_1(\boldsymbol{u}_1) \xrightarrow{\beta_1}_{\mathbb{O}} q_1(\boldsymbol{v}_1) \xrightarrow{\alpha_1}_{\mathbb{O}} q_2(\boldsymbol{u}_2) \cdots q_k(\boldsymbol{u}_k) \xrightarrow{\beta_k}_{\mathbb{O}} q_k(\boldsymbol{v}_k) \xrightarrow{\alpha_k}_{\mathbb{O}} q(\boldsymbol{v}) \tag{38}$$

where

- $q_1, \ldots, q_k \in Q$ are pairwise distinct, so $k \leq |Q|$;
- $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_k, \boldsymbol{v}_1, \ldots, \boldsymbol{v}_k \in \mathbb{O}$;
- $\alpha_0 \alpha_1 \ldots \alpha_k$ is a simple path, so $|\alpha_0| + \cdots + |\alpha_k| \leq |Q|$.

By Lemma A.13, for each cycle β_i there exists an LPS Λ_i compatible to G such that $q_i(\boldsymbol{u}_i) \xrightarrow{\Lambda_i} q_i(\boldsymbol{v}_i)$ and $|\Lambda_i| \leq |G|^{O(1)}$. Let

$$\Lambda = \alpha_0 \Lambda_1 \alpha_1 \ldots \alpha_{k-1} \Lambda_k \alpha_k, \tag{39}$$

it is easy to see that Λ is compatible to G, and $p(\boldsymbol{u}) \xrightarrow{\Lambda} q(\boldsymbol{v})$ holds. Finally, we bound the length of Λ:

$$|\Lambda| = \sum_{i=0}^{k} |\alpha_i| + \sum_{i=1}^{k} |\Lambda_i| \leq |Q| + |Q| \cdot |G|^{O(1)} \leq |G|^{O(1)}. \tag{40}$$

◄

We remark that Lemma A.14 captures only the runs that stay in the region $[D, \infty)^d$, while the run derived from the LPS Λ given by it may go out of $[D, \infty)^d$ but stays in $\mathbb{N}^d$.

A.5 Reachability Near the Axes

We have captured runs that stay far away from axes by small LPSes in Section A.4. In this section we study the complementary region. Let $D \in \mathbb{N}$ be any natural number, we define the region

$$\mathbb{L}[D] := \{\boldsymbol{u} \in \mathbb{N}^d : \boldsymbol{u}(i) \leq D \text{ for some } i \in [d]\}. \tag{41}$$

The following is to be proved:

▶ **Lemma A.15.** *Let $D \in \mathbb{N}$ and let $G = (Q, T)$ be a d-VASS that is geometrically 2-dimensional. For every pair of configurations $p(\boldsymbol{u}), q(\boldsymbol{v}) \in Q \times \mathbb{L}[D]$ with $p(\boldsymbol{u}) \xrightarrow{*}_{\mathbb{L}[D]} q(\boldsymbol{v})$ there exists an LPS Λ compatible to G such that $p(\boldsymbol{u}) \xrightarrow{\Lambda} q(\boldsymbol{v})$, and $|\Lambda| \leq (|G| + D)^{O(1)}$.*

Let's first subdivide the region $\mathbb{L}[D]$ into simpler regions. For any $D \in \mathbb{N}$ and $i \in [d]$, we define

$$\mathbb{L}_i[D] := \{\boldsymbol{u} \in \mathbb{N}^d : \boldsymbol{u}(i) \leq D\}. \tag{42}$$

Clearly $\mathbb{L}[D] = \bigcup_{i \in [d]} \mathbb{L}_i[D]$. We shall prove a stronger result that runs in a union of any number of $\mathbb{L}_i[D]$'s can be captured by LPSes. For $\mathcal{I} \subseteq [d]$ we write $\mathbb{L}_{\mathcal{I}}[D] := \bigcup_{i \in \mathcal{I}} \mathbb{L}_i[D]$.

▶ **Lemma A.16.** *Let $D \in \mathbb{N}$, $\mathcal{I} \subseteq [d]$, and let $G = (Q, T)$ be a d-VASS that is geometrically 2-dimensional. For every pair of configurations $p(\boldsymbol{u}), q(\boldsymbol{v}) \in Q \times \mathbb{L}_{\mathcal{I}}[D]$ with $p(\boldsymbol{u}) \xrightarrow{*}_{\mathbb{L}_{\mathcal{I}}[D]} q(\boldsymbol{v})$ there exists an LPS Λ compatible to G such that $p(\boldsymbol{u}) \xrightarrow{\Lambda} q(\boldsymbol{v})$, and $|\Lambda| \leq (|G| + D)^{O(1)}$.*

Note that Lemma A.15 follows directly from Lemma A.16 by taking $\mathcal{I} = [d]$. In the rest we will prove Lemma A.16 by induction on the size of $\mathcal{I}$.

A.5.1 The Base Case: Reachability with One Bounded Component

A run that stays in $\mathbb{L}_i[D]$ has its ith component bounded by D. In this case we can simply encode this component in the states and reduce the dimension by one. Recall that we are proving Theorem 3.4 by induction on the dimension d, so here we assume that Theorem 3.4 already holds for $(d-1)$-VASSes as the induction hypothesis.

▶ **Lemma A.17.** *Let $0 \leq B \leq C$ and $i \in [d]$, let $\mathbb{L}_i[B,C]$ denote the region $\{\boldsymbol{u} \in \mathbb{N}^d : B \leq \boldsymbol{u}(i) \leq C\}$. Let $G = (Q,T)$ be a d-VASS that is geometrically 2-dimensional. For every pair of configurations $p(\boldsymbol{u}), q(\boldsymbol{v}) \in Q \times \mathbb{L}_i[B,C]$ with $p(\boldsymbol{u}) \xrightarrow{*}_{\mathbb{L}_i[B,C]} q(\boldsymbol{v})$ there exists an LPS Λ compatible to G such that $p(\boldsymbol{u}) \xrightarrow{\Lambda} q(\boldsymbol{v})$, and $|\Lambda| \leq (|G| + (C-B))^{O(1)}$.*

Proof. Let $G^{-i} = (Q^{-i}, T^{-i})$ be the $(d-1)$-VASS obtained from G where

$$Q^{-i} = \{(q,z) : q \in Q, z \in [B,C]\}, \tag{43}$$

$$T^{-i} = \{((p,z), \boldsymbol{a}^{-i}, (q,z')) : (p, \boldsymbol{a}, q) \in T, z' - z = \boldsymbol{a}(i)\}. \tag{44}$$

For each transition $t = ((p,z), \boldsymbol{a}^{-i}, (q,z')) \in T^{-i}$ let $g(t)$ be the unique transition $(p, \boldsymbol{a}, q) \in T$ where $\boldsymbol{a}(i) = z' - z$. The mapping g is extended to linear path schemes compatible to G^{-i} as a word morphism.

Note that G^{-i} is still geometrically 2-dimensional as G is. Moreover, $|Q^{-i}| = |Q| \cdot (C - B + 1)$, and $\|T^{-i}\| \leq \|T\|$, so $|G^{-i}| \leq (|G| + C - B)^{O(1)}$. Let $p(\boldsymbol{u}), q(\boldsymbol{v}) \in Q \times \mathbb{L}_i[B,C]$ be configurations in G such that $p(\boldsymbol{u}) \xrightarrow{*}_{\mathbb{L}_i[B,C]} q(\boldsymbol{v})$ holds in G. Then $(p, \boldsymbol{u}(i))(\boldsymbol{u}^{-i}) \xrightarrow{*}_{\mathbb{N}^{d-1}} (q, \boldsymbol{v}(i))(\boldsymbol{v}^{-i})$ holds in G^{-i}. Recall we have assumed that Theorem 3.4 holds for dimension $d-1$. In particular, it holds for G^{-i}. Thus there exists an LPS Λ compatible to G^{-i} such that $(p, \boldsymbol{u}(i))(\boldsymbol{u}^{-i}) \xrightarrow{\Lambda}_{\mathbb{N}^{d-1}} (q, \boldsymbol{v}(i))(\boldsymbol{v}^{-i})$. Note that $g(\Lambda)$ is an LPS compatible to G. Since the ith coordinate is encoded in the state of G^{-i}, one easily checks that $p(\boldsymbol{u}) \xrightarrow{g(\Lambda)}_{\mathbb{N}^d} q(\boldsymbol{v})$ also holds in G.

As for the size of Λ, just note that $|g(\Lambda)| = |\Lambda| \leq |G^{-i}|^{O(1)} \leq (|G| + C - B)^{O(1)}$. ◀

The base case of Lemma A.16 where $\mathcal{I}$ contains only one index is just an easy corollary of Lemma A.17 by taking $B = 0$ and $C = D$. The general form of Lemma A.17 will be used later.

A.5.2 The Inductive Proof of Lemma A.16

We now consider $\mathbb{L}_{\mathcal{I}}[D]$-reachability for $\mathcal{I} \subseteq [d]$ with $|\mathcal{I}| > 1$, assuming that Lemma A.16 holds for all subsets $\mathcal{I}' \subsetneq \mathcal{I}$. First we identify a special case where $\mathrm{Cyc}(G)$ is orthogonal to one of the axes, which makes G essentially a $(d-1)$-VASS.

▶ **Lemma A.18.** *Let $G = (Q,T)$ be a d-VASS that is geometrically 2-dimensional. Suppose there exists $i \in [d]$ such that $i \notin \mathrm{supp}(\mathrm{Cyc}(G))$. Then for every pair of configurations $p(\boldsymbol{u}), q(\boldsymbol{v}) \in Q \times \mathbb{N}^d$ with $p(\boldsymbol{u}) \xrightarrow{*} q(\boldsymbol{v})$ there exists an LPS Λ compatible to G such that $p(\boldsymbol{u}) \xrightarrow{\Lambda} q(\boldsymbol{v})$, and $|\Lambda| \leq |G|^{O(1)}$.*

Proof. Let $p(\boldsymbol{u}), q(\boldsymbol{v}) \in Q \times \mathbb{N}^d$ be two configurations of G with $p(\boldsymbol{u}) \xrightarrow{*} q(\boldsymbol{v})$. We show that there exists an LPS Λ compatible to G such that $p(\boldsymbol{u}) \xrightarrow{\Lambda} q(\boldsymbol{v})$ and $|\Lambda| \leq |G|^{O(1)}$.

First we claim that for any configurations $r(\boldsymbol{w}) \in Q \times \mathbb{N}^d$ reachable from $p(\boldsymbol{u})$, we have $|\boldsymbol{w}(i) - \boldsymbol{u}(i)| \leq |Q| \cdot \|T\|$. Otherwise, assume there is a run $\pi : p(\boldsymbol{u}) \xrightarrow{*} r(\boldsymbol{w})$ such that

$|\boldsymbol{w}(i) - \boldsymbol{u}(i)| > |Q| \cdot \|T\|$. Assume w.l.o.g. it is the case that $\boldsymbol{w}(i) - \boldsymbol{u}(i) > |Q| \cdot \|T\|$. Write π in the following form as a sequence of sub-runs that increase the ith component like a stair:

$$\pi : p(\boldsymbol{u}) = q_0(\boldsymbol{u}_0) \xrightarrow{*} q_1(\boldsymbol{u}_1) \xrightarrow{*} \cdots \xrightarrow{*} q_k(\boldsymbol{u}_k) \xrightarrow{*} r(\boldsymbol{w}) \tag{45}$$

such that $\boldsymbol{u}_k(i) \geq \boldsymbol{w}(i) > \boldsymbol{u}(i) + |Q| \cdot \|T\|$ and for all $j \in [k]$, $\boldsymbol{u}_j(i) > \boldsymbol{u}_{j-1}(i)$ and every configuration on the sub-run from $q_{j-1}(\boldsymbol{u}_{j-1})$ to $q_j(\boldsymbol{u}_j)$ has its i-th component no greater than $\boldsymbol{u}_{j-1}(i)$. Clearly $\boldsymbol{u}_j(i) - \boldsymbol{u}_{j-1}(i) \leq \|T\|$, so $k \geq (\boldsymbol{u}_k(i) - \boldsymbol{u}(i))/\|T\| > |Q|$. By the Pigeonhole Principle there exists $j_1 < j_2$ such that $q_{j_1} = q_{j_2}$. The sub-run from $q_{j_1}(\boldsymbol{u}_{j_1})$ to $q_{j_2}(\boldsymbol{u}_{j_2})$ is a cycle which has positive effect in its i-th component, contradicting to the assumption that $i \notin \operatorname{supp}(\operatorname{Cyc}(G))$.

We have shown that $p(\boldsymbol{u}) \xrightarrow{*}_{\mathbb{L}_i[B,C]} q(\boldsymbol{v})$ where $B = \boldsymbol{u}(i) - |Q| \cdot \|T\|$ and $C = \boldsymbol{u}(i) + |Q| \cdot \|T\|$. Note that $C - B = 2 \cdot |Q| \cdot \|T\| \leq |G|^{O(1)}$. The result follows directly from Lemma A.17. ◄

Now we prove Lemma A.16.

Proof of Lemma A.16. We prove by induction on $|\mathcal{I}|$. The base case where $|\mathcal{I}| = 1$ follows from Lemma A.17 by taking $B = 0$ and $C = D$. Now assume that $\mathcal{I} = \{i\} \uplus \mathcal{I}'$ where $i \in \mathcal{I}$ is arbitrary and $\mathcal{I}' = \mathcal{I} \setminus \{i\}$ is non-empty. Fix $p(\boldsymbol{u}), q(\boldsymbol{v}) \in Q \times \mathbb{N}^d$ with $p(\boldsymbol{u}) \xrightarrow{*}_{\mathbb{L}_\mathcal{I}[D]} q(\boldsymbol{v})$, we will exhibit an LPS Λ compatible to G such that $p(\boldsymbol{u}) \xrightarrow{\Lambda} q(\boldsymbol{v})$.

Fix a run π witnessing $p(\boldsymbol{u}) \xrightarrow{*}_{\mathbb{L}_\mathcal{I}[D]} q(\boldsymbol{v})$ such that no configuration occurs more than once on π. We assume that π must visit both $\mathbb{L}_i[D]$ and $\mathbb{L}_{\mathcal{I}'}[D]$, otherwise we can directly apply the induction hypothesis. Observe that π can be written in the following form as an alternating sequence of sub-runs lie in either $\mathbb{L}_i[D]$ or $\mathbb{L}_{\mathcal{I}'}[D]$:

$$\pi : p(\boldsymbol{u}) \xrightarrow{\sigma}_{\mathbb{R}_0} q_0(\boldsymbol{u}_0) \xrightarrow{\pi_1}_{\mathbb{R}_1} q_1(\boldsymbol{u}_1) \xrightarrow{\pi_2}_{\mathbb{R}_2} \cdots \xrightarrow{\pi_n}_{\mathbb{R}_n} q_n(\boldsymbol{u}_n) \xrightarrow{\rho}_{\mathbb{R}_{n+1}} q(\boldsymbol{u}) \tag{46}$$

where
- $\mathbb{R}_h \in \{\mathbb{L}_i[D], \mathbb{L}_{\mathcal{I}'}[D]\}$ for $h = 0, \ldots, n + 1$, and $\mathbb{R}_h \neq \mathbb{R}_{h-1}$ for $h \in [n + 1]$,
- $\boldsymbol{u}_h \in \mathbb{J}_{i\ell}$ for some (not necessarily unique) $\ell \in \mathcal{I}'$ for each $h = 0, \ldots, n$, where

$$\mathbb{J}_{i\ell} := \{\boldsymbol{w} \in \mathbb{N}^d : \boldsymbol{w}(i), \boldsymbol{w}(\ell) \leq D + \|T\|\}. \tag{47}$$

(Note that we increase the bound by $\|T\|$ so that any one-step transition from $\mathbb{L}_i[D]$ to $\mathbb{L}_\ell[D]$ lies completely in $\mathbb{J}_{i\ell}$.)

Note that the prefix σ and the suffix ρ can be captured by LPSes of length $(|G| + D)^{O(1)}$ by induction hypothesis. In the following we show that the infix $\overline{\pi} := \pi_1 \pi_2 \ldots \pi_n$ can also be captured by an LPS $\overline{\Lambda}$ of length $(|G| + D)^{O(1)}$. Then the entire run π is captured by a concatenation of these three LPSes.

Let π_h be any sub-run of $\overline{\pi}$ such that $\mathbb{R}_h = \mathbb{L}_i[D]$ where $h \in [n]$. Expand π_h as

$$\pi_h : q_{h-1}(\boldsymbol{u}_{h-1}) =: s_0(\boldsymbol{v}_0) \xrightarrow{t_1}_{\mathbb{L}_i[D]} s_1(\boldsymbol{v}_1) \xrightarrow{t_2}_{\mathbb{L}_i[D]} \cdots \xrightarrow{t_m}_{\mathbb{L}_i[D]} s_m(\boldsymbol{v}_m) := q_h(\boldsymbol{u}_h), \tag{48}$$

where $t_1, \ldots, t_m \in T$ are transitions. Denote $D' := (D + \|T\|) + |Q| \cdot D \cdot \|T\|$, we classify π_h into the following types:

Type 1: *There exists $\ell \in \mathcal{I}'$ such that $\boldsymbol{v}_k(\ell) \leq D'$ for all $k = 0, \ldots, m$.*

In this case, π_h is actually an $\mathbb{L}_{\mathcal{I}'}[D']$-run.

Type 2(ℓ), where $\ell \in \mathcal{I}'$: *There exists $k \in [m]$ such that $\boldsymbol{v}_k(\ell) > D'$, and $\boldsymbol{u}_{h-1} = \boldsymbol{v}_0 \in \mathbb{J}_{i\ell}$.*

In this case $\boldsymbol{v}_k(\ell) - \boldsymbol{v}_0(\ell) > D' - (D + \|T\|) = |Q| \cdot D \cdot \|T\|$. Let k be minimal such that $\boldsymbol{v}_k(\ell) > D'$. For each $\kappa < k$ we define the index $\operatorname{next}(\kappa)$ as the minimal index κ' with $\kappa < \kappa' \leq k$ and $\boldsymbol{v}_{\kappa'}(\ell) > \boldsymbol{v}_\kappa(\ell)$. Let $\kappa_1, \kappa_2, \ldots, \kappa_M$ be the sequence of indices such

that $\kappa_1 = 0$ and $\kappa_{j+1} = \mathrm{next}(\kappa_j)$ for all $j = 1, \ldots, M-1$, and $\kappa_M = k$. We must have $M > |Q| \cdot D$. Since $\boldsymbol{v}_{\kappa_j}(i) \leq D$ for all κ_j (recall π_h stays in $\mathbb{L}_i[D]$), by the Pigeonhole Principle there exist $1 \leq \kappa_a < \kappa_b \leq k$ such that

$$s_{\kappa_a} = s_{\kappa_b}, \quad \boldsymbol{v}_{\kappa_a}(i) = \boldsymbol{v}_{\kappa_b}(i), \quad \boldsymbol{v}_{\kappa_a}(\ell) < \boldsymbol{v}_{\kappa_b}(\ell). \tag{49}$$

Therefore, the existence of a Type $2(\ell)$ sub-run π_h implies that $\mathrm{Cyc}(G)$ contains a vector $\boldsymbol{r}$ such that $i \notin \mathrm{supp}(\boldsymbol{r})$ and $\ell \in \mathrm{supp}(\boldsymbol{r})$.

Recall that the run π_h must start from a place in $\mathbb{J}_{i\ell}$ for some $\ell \in \mathcal{I}'$. So π_h is either of Type 1 or of Type $2(\ell)$ (or both, but we let Type 1 take precedence).

If all the $\mathbb{L}_i[D]$-sub-runs π_h are of Type 1, then $\overline{\pi}$ is indeed an $\mathbb{L}_{\mathcal{I}'}[D']$-run, and we are done by applying the induction hypothesis with $D' \leq (|G| + D)^{O(1)}$ in place of D. Next we assume there exist Type $2(\ell)$ sub-runs. By considering $\mathbb{L}_{\mathcal{I}'}[D']$ instead of $\mathbb{L}_{\mathcal{I}'}[D]$, we can rewrite $\overline{\pi}$ in the following form to get rid of all Type 1 sub-runs in $\mathbb{L}_i[D]$:

$$\overline{\pi} : q_0(\boldsymbol{u}_0) = q_0'(\boldsymbol{u}_0') \xrightarrow{\pi_1'}_{\mathbb{R}_1'} q_1'(\boldsymbol{u}_1') \xrightarrow{\pi_2'}_{\mathbb{R}_2'} \cdots \xrightarrow{\pi_N'}_{\mathbb{R}_N'} q_N'(\boldsymbol{u}_N') = q_n(\boldsymbol{u}_n) \tag{50}$$

where

- $\mathbb{R}_h' \in \{\mathbb{L}_i[D], \mathbb{L}_{\mathcal{I}'}[D']\}$ for $h = 0, \ldots, N+1$, and $\mathbb{R}_h' \neq \mathbb{R}_{h-1}'$ for $h \in [N+1]$,
- $\boldsymbol{u}_h' \in \mathbb{J}_{i\ell}$ for some (not necessarily unique) $\ell \in \mathcal{I}'$ for each $h = 0, \ldots, N$,
- for each $h \in [N]$ such that $\mathbb{R}_h' = \mathbb{L}_i[D]$, π_h is of Type $2(\ell)$ for some $\ell \in \mathcal{I}'$.

Let

$$\mathcal{H}_\ell = \{h \in [N] : \mathbb{R}_h' = \mathbb{L}_i[D] \text{ and } \pi_h' \text{ is of Type } 2(\ell)\} \tag{51}$$

be the set of indices of those $\mathbb{L}_i[D]$-sub-runs that are of Type $2(\ell)$. There remain two cases:

Case 1. *For all $\ell \in \mathcal{I}'$, we have $|\mathcal{H}_\ell| \leq |Q| \cdot (D + \|T\|)^2$.*
In this case, the total number of indices $h \in [N]$ such that $\mathbb{R}_h = \mathbb{L}_i[D]$ is at most $|\mathcal{I}'| \cdot |Q| \cdot (D + \|T\|)^2 \leq d \cdot |Q| \cdot (D + \|T\|)^2$ and thus $N \leq 2 \cdot d \cdot |Q|(D + \|T\|)^2 + 1 \leq (|G| + D)^{O(1)}$. Since by induction hypothesis each sub-run π_h where $h \in [N]$ can be captured by an LPS Λ_h of size $(|G| + D)^{O(1)}$, we are done by taking $\overline{\Lambda}$ as the concatenation of all Λ_h's.

Case 2. *There exists $\ell \in \mathcal{I}'$ such that $|\mathcal{H}_\ell| > |Q| \cdot (D + \|T\|)^2$.*
In this case, consider the following set of configurations:

$$\mathcal{C}_\ell := \{q_{h-1}'(\boldsymbol{u}_{h-1}') : h \in \mathcal{H}_\ell\} \tag{52}$$

which are the starting configurations of the Type $2(\ell)$ sub-runs. By definition, we have $s(\boldsymbol{v}) \in \mathbb{J}_{i\ell}$ for all $s(\boldsymbol{v}) \in \mathcal{C}_\ell$. Since $|\mathcal{C}_\ell| = |\mathcal{H}_\ell| > |Q| \cdot (D + \|T\|)^2$, by the Pigeonhole Principle there exists two configurations $s_1(\boldsymbol{v}_1), s_2(\boldsymbol{v}_2) \in \mathcal{C}_\ell$ such that

$$s_1 = s_2, \quad \boldsymbol{v}_1(i) = \boldsymbol{v}_2(i), \quad \boldsymbol{v}_1(\ell) = \boldsymbol{v}_2(\ell), \quad \boldsymbol{v}_1 \neq \boldsymbol{v}_2, \quad \text{and} \quad s_1(\boldsymbol{v}_1) \xrightarrow{*} s_2(\boldsymbol{v}_2). \tag{53}$$

This implies that $\boldsymbol{r} = \boldsymbol{s}_2 - \boldsymbol{s}_1 \in \mathrm{Cyc}(G)$ is non-zero, and $i, \ell \notin \mathrm{supp}(\boldsymbol{r})$. On the other hand, the existence of Type $2(\ell)$ sub-runs also implies that there exists $\boldsymbol{r}' \in \mathrm{Cyc}(G)$ such that $i \notin \mathrm{supp}(\boldsymbol{r}')$ and $\ell \in \mathrm{supp}(\boldsymbol{r}')$. Clearly $\boldsymbol{r}$ and $\boldsymbol{r}'$ are linearly independent. As G is geometrically 2-dimensional, we have $\mathrm{Cyc}(G) = \mathrm{span}\{\boldsymbol{r}, \boldsymbol{r}'\}$ and thus $i \notin \mathrm{supp}(\mathrm{Cyc}(G))$. This situation is already solved by Lemma A.18. ◀

A.6   Putting All Together: Proof of Theorem 3.4

Having established Lemmas A.13, A.14 and A.15, we are now at a position to prove Theorem 3.4.

▶ **Theorem 3.4.** *Let $G = (Q, T)$ be a geometrically 2-dimensional d-VASS. For every pair of configurations $p(\boldsymbol{u}), q(\boldsymbol{v}) \in Q \times \mathbb{N}^d$ with $p(\boldsymbol{u}) \xrightarrow{*} q(\boldsymbol{v})$ there exists an LPS Λ compatible to G such that $p(\boldsymbol{u}) \xrightarrow{\Lambda} q(\boldsymbol{v})$ and $|\Lambda| \leq |G|^{O(1)}$.*

Proof of Theorem 3.4. Let D be as in Lemmas A.13 and A.14. We define the following regions:

$$\mathbb{L} := \mathbb{L}[D + \|T\|], \qquad \mathbb{O} := [D, \infty)^d \tag{54}$$

where $\mathbb{L}[D + \|T\|]$ is defined in Section A.5.

Let $p(\boldsymbol{u}), q(\boldsymbol{v})$ be two configurations of G such that $p(\boldsymbol{u}) \xrightarrow{*} q(\boldsymbol{v})$. We show that there exists an LPS Λ compatible to G such that $p(\boldsymbol{u}) \xrightarrow{\Lambda} q(\boldsymbol{v})$ and $|\Lambda| \leq |G|^{O(1)}$.

Let π be a run from $p(\boldsymbol{u})$ to $q(\boldsymbol{v})$. We expand π as the following form:

$$\pi : p(\boldsymbol{u}) = q_0(\boldsymbol{u}_0) \xrightarrow{t_1} q_1(\boldsymbol{u}_1) \xrightarrow{t_2} \ldots \xrightarrow{t_n} q_n(\boldsymbol{u}_n) = q(\boldsymbol{v}) \tag{55}$$

where $t_1, \ldots, t_n \in T$. Define $K = \{k \in \{0, \ldots, n\} : \boldsymbol{u}_k \in \mathbb{L} \cap \mathbb{O}\}$. For $k \in K$ let $\mathrm{next}(k)$ be the minimal index $k' \in K$ such that $k' > k$ (if k is the maximum element in K then we set $\mathrm{next}(k) = k$), and let $\ell(k)$ be the maximal index $k' \in K$ such that $q_k = q_{k'}$. We rewrite π into the following form:

$$\begin{aligned}
\pi : q_0(\boldsymbol{u}_0) &\xrightarrow{\alpha_0}_{\mathbb{R}_0} q_{k_1}(\boldsymbol{u}_{k_1}) \xrightarrow{\gamma_1} q_{\ell(k_1)}(\boldsymbol{u}_{\ell(k_1)}) \\
&\xrightarrow{\alpha_1}_{\mathbb{R}_1} q_{k_2}(\boldsymbol{u}_{k_2}) \xrightarrow{\gamma_2} q_{\ell(k_2)}(\boldsymbol{u}_{\ell(k_2)}) \\
&\ldots \\
&\xrightarrow{\alpha_{h-1}}_{\mathbb{R}_{h-1}} q_{k_h}(\boldsymbol{u}_{k_h}) \xrightarrow{\gamma_h} q_{\ell(k_h)}(\boldsymbol{u}_{\ell(k_h)}) \xrightarrow{\alpha_h}_{\mathbb{R}_h} q_n(\boldsymbol{u}_n)
\end{aligned} \tag{56}$$

where
- $k_1 = \min K$, $k_{i+1} = \mathrm{next}(\ell(k_i))$ for all $1 \leq i < h$, $\ell(k_h) = \max K$;
- $\mathbb{R}_0, \ldots, \mathbb{R}_h \in \{\mathbb{L}, \mathbb{O}\}$.

By Lemma A.14 and A.15, each α_i can be captured by an LPS Λ_i such that $|\Lambda_i| \leq |G|^{O(1)}$. By Lemma A.13, each γ_i is also captured by an LPS Γ_i such that $|\Gamma_i| \leq |G|^{O(1)}$. Note that $h \leq |Q|$. Let

$$\Lambda := \Lambda_0 \Gamma_1 \Lambda_1 \ldots \Gamma_h \Lambda_h. \tag{57}$$

Then Λ is compatible to G and $p(\boldsymbol{u}) \xrightarrow{\Lambda} q(\boldsymbol{v})$. As for the length of $|\Lambda|$, just note that $|\Lambda| \leq (2|Q| + 1) \cdot |G|^{O(1)} \leq |G|^{O(1)}$. ◀

B   Proof of Lemma 4.2

▶ **Lemma 4.2.** *Let G be a d-VASS and $\Lambda = \alpha_0 \beta_1^+ \alpha_1 \ldots \beta_k^+ \alpha_k$ be a positive LPS from state p to q compatible to G. Then for every $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{N}^d$, $p(\boldsymbol{u}) \xrightarrow{\Lambda} q(\boldsymbol{v})$ if and only if there exists $\boldsymbol{e} \in \mathbb{N}^k$ such that $(\boldsymbol{u}, \boldsymbol{e}, \boldsymbol{v}) \models \mathcal{E}_{\mathrm{LPS}}(\Lambda)$. Moreover, for every $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{N}^d$ and every $\boldsymbol{e} \in \mathbb{N}^k$ such that $(\boldsymbol{u}, \boldsymbol{e}, \boldsymbol{v}) \models \mathcal{E}_{\mathrm{LPS}}(\Lambda)$, we have $p(\boldsymbol{u}) \xrightarrow{\alpha_0 \beta_1^{\boldsymbol{e}(1)} \alpha_1 \ldots \beta_k^{\boldsymbol{e}(k)} \alpha_k} q(\boldsymbol{v})$.*

Proof. The "only if" part is straightforward: assuming $p(\boldsymbol{u}) \xrightarrow{\Lambda} q(\boldsymbol{v})$, there exist numbers $e_1, \ldots, e_k \in \mathbb{N}_{>0}$ such that $\pi : p(\boldsymbol{u}) \xrightarrow{\alpha_0 \beta_1^{e_1} \alpha_1 \ldots \beta_k^{e_k} \alpha_k} q(\boldsymbol{v})$. Let $\boldsymbol{e} \in \mathbb{N}^k$ such that $\boldsymbol{e}(i) := e_i$, one easily sees that $(\boldsymbol{u}, \boldsymbol{e}, \boldsymbol{v}) \models \mathcal{E}_{\mathrm{LPS}}(\Lambda)$ since $\mathcal{E}_{\mathrm{LPS}}(\Lambda)$ merely asserts that a small subset of the configurations visited by the run π stay in $\mathbb{N}^d$, and that $\Delta(\pi) = \boldsymbol{v} - \boldsymbol{u}$.

Now we focus on the "if" part. Assume $(\boldsymbol{u}, \boldsymbol{e}, \boldsymbol{v}) \models \mathcal{E}_{\mathrm{LPS}}(\Lambda)$ for some $\boldsymbol{e} = (e_1, \ldots, e_k) \in \mathbb{N}_{>0}^k$. We claim that

$$p(\boldsymbol{u}) \xrightarrow{\alpha_0 \beta_1^{e_1} \alpha_1 \ldots \beta_k^{e_k} \alpha_k}_{\mathbb{N}^d} q(\boldsymbol{v}). \tag{58}$$

It is clear that $\boldsymbol{u} + \Delta(\alpha_0 \beta_1^{e_1} \alpha_1 \ldots \beta_k^{e_k} \alpha_k) = \boldsymbol{v}$, which is asserted in the fourth condition of $\mathcal{E}_{\mathrm{LPS}}(\Lambda)$. To prove that every configuration along this walk stay in $\mathbb{N}^d$, we only need to check that for every $i = 1, \ldots, k$ with $e_i > 1$, every $\ell = 1, \ldots, e_i - 1$ and every $j = 1, \ldots, |\beta_i|$,

$$\boldsymbol{u} + \Delta(\alpha_0 \beta_1^{e_1} \alpha_1 \ldots \beta_{i-1}^{e_{i-1}} \alpha_{i-1} \beta_i^{\ell}) + \Delta(\beta_i[1 \ldots j]) \geq \boldsymbol{0}. \tag{59}$$

Note that by the third condition of $\mathcal{E}_{\mathrm{LPS}}(\Lambda)$, we have

$$\Delta_0 := \boldsymbol{u} + \Delta(\alpha_0 \beta_1^{e_1} \alpha_1 \ldots \beta_{i-1}^{e_{i-1}} \alpha_{i-1}) + 0 \cdot \Delta(\beta_i) \qquad + \Delta(\beta_i[1 \ldots j]) \geq \boldsymbol{0}, \tag{60}$$

$$\Delta_1 := \boldsymbol{u} + \Delta(\alpha_0 \beta_1^{e_1} \alpha_1 \ldots \beta_{i-1}^{e_{i-1}} \alpha_{i-1}) + (e_i - 1) \cdot \Delta(\beta_i) + \Delta(\beta_i[1 \ldots j]) \geq \boldsymbol{0}. \tag{61}$$

Then

$$\begin{aligned}
&\boldsymbol{u} + \Delta(\alpha_0 \beta_1^{e_1} \alpha_1 \ldots \beta_{i-1}^{e_{i-1}} \alpha_{i-1} \beta_i^{\ell}) + \Delta(\beta_i[1 \ldots j]) \\
&= \boldsymbol{u} + \Delta(\alpha_0 \beta_1^{e_1} \alpha_1 \ldots \beta_{i-1}^{e_{i-1}} \alpha_{i-1}) + \ell \cdot \Delta(\beta_i) + \Delta(\beta_i[1 \ldots j]) \\
&= \frac{e_i - 1 - \ell}{e_i - 1} \cdot \Delta_0 + \frac{\ell}{e_i - 1} \cdot \Delta_1 \geq \boldsymbol{0}.
\end{aligned} \tag{62}$$

This proves (59). ◄

C   Proofs Omitted from Section 5

C.1   Proof of Lemma 5.6

▶ **Lemma 5.6** (cf. [15, Lem. 3.5]). *For any unsatisfiable linear KLM sequence ξ, $L_\xi = \emptyset$.*

Proof. Let $\xi = \langle p_0(\boldsymbol{x}_0) G_0 q_0(\boldsymbol{y}_0) \rangle \Lambda_1 \cdots \Lambda_k \langle p_k(\boldsymbol{x}_k) G_k q_k(\boldsymbol{y}_k) \rangle$. If $L_\xi \neq \emptyset$, then there is a run π admitted by ξ of the form

$$p_0(\boldsymbol{m}_0) \xrightarrow{\pi_0} q_0(\boldsymbol{n}_0) \xrightarrow{\rho_1} \cdots \xrightarrow{\rho_k} p_k(\boldsymbol{m}_k) \xrightarrow{\pi_k} q_k(\boldsymbol{n}_k). \tag{63}$$

Let ϕ_i be the Parikh image of π_i for each $i = 0, \ldots, k$. Since $q_{i-1}(\boldsymbol{n}_{i-1}) \xrightarrow{\rho_i} p_i(\boldsymbol{m}_i)$ is admitted by Λ_i, by Lemma 4.2 there exists vector $\boldsymbol{e}_i \in \mathbb{N}^{|\Lambda_i|_*}$ such that $(\boldsymbol{n}_{i-1}, \boldsymbol{e}_i, \boldsymbol{m}_i) \models \mathcal{E}_{\mathrm{LPS}}(\Lambda_i)$. Now the sequence $(\boldsymbol{m}_0, \phi_0, \boldsymbol{n}_0), \boldsymbol{e}_1, \ldots, \boldsymbol{e}_k, (\boldsymbol{m}_k, \phi_k, \boldsymbol{n}_k)$ is a model of $\mathcal{E}(\xi)$. ◄

C.2   Proof of Lemma 5.7

▶ **Lemma 5.7.** *Assume that $\xi = \langle p_0(\boldsymbol{x}_0) G_0 q_0(\boldsymbol{y}_0) \rangle \Lambda_1 \cdots \Lambda_k \langle p_k(\boldsymbol{x}_k) G_k q_k(\boldsymbol{y}_k) \rangle$ is satisfiable. Then for every $0 \leq i \leq k$ we have:*

- *For every $1 \leq j \leq d$, the set of values $\boldsymbol{m}_i^{\boldsymbol{h}}(j)$ where $\boldsymbol{h} \models \mathcal{E}(\xi)$ is unbounded if, and only if, there exists a model $\boldsymbol{h}_0$ of $\mathcal{E}^0(\xi)$ such that $\boldsymbol{m}_i^{\boldsymbol{h}_0}(j) > 0$.*
- *For every $t \in T_i$, the set of values $\phi_i^{\boldsymbol{h}}(t)$ where $\boldsymbol{h} \models \mathcal{E}(\xi)$ is unbounded if, and only if, there exists a model $\boldsymbol{h}_0$ of $\mathcal{E}^0(\xi)$ such that $\phi_i^{\boldsymbol{h}_0}(t) > 0$.*

- For every $1 \leq j \leq d$, the set of values $\boldsymbol{n}_i^{\boldsymbol{h}}(j)$ where $\boldsymbol{h} \models \mathcal{E}(\xi)$ is unbounded if, and only if, there exists a model $\boldsymbol{h}_0$ of $\mathcal{E}^0(\xi)$ such that $\boldsymbol{n}_i^{\boldsymbol{h}_0}(j) > 0$.

Moreover, every bounded value of $\mathcal{E}(\xi)$ is bounded by $(10|\xi|)^{12|\xi|}$.

We first recall some facts in linear algebra.

▶ **Lemma C.1** ([3, Prop. 4], see also [18]). *Let $A \in \mathbb{Z}^{m \times n}$ be an integer matrix and $\boldsymbol{b} \in \mathbb{Z}^m$ be an integer vector. Define the following two sets:*

$$\boldsymbol{X} := \{\boldsymbol{x} \in \mathbb{N}^n : A\boldsymbol{x} = \boldsymbol{b}\}, \qquad \boldsymbol{X}_0 := \{\boldsymbol{x} \in \mathbb{N}^n : A\boldsymbol{x} = \boldsymbol{0}\}. \tag{64}$$

Then every vector in $\boldsymbol{X}$ can be decomposed as the sum of a vector $\boldsymbol{x} \in \boldsymbol{X}$ and a finite sum of vectors $\boldsymbol{x}_0 \in \boldsymbol{X}_0$ such that

$$\|\boldsymbol{x}\| \leq ((n+1)\|A\| + \|\boldsymbol{b}\| + 1)^m, \qquad \|\boldsymbol{x}_0\| \leq (n\|A\| + 1)^m. \tag{65}$$

As linear inequalities can be converted to linear equations by adding new variables, we can obtain a similar result for systems of linear inequalities.

▶ **Corollary C.2.** *Let $A \in \mathbb{Z}^{m \times n}$ be an integer matrix and $\boldsymbol{b} \in \mathbb{Z}^m$ be an integer vector. Define the following two sets:*

$$\boldsymbol{X} := \{\boldsymbol{x} \in \mathbb{N}^n : A\boldsymbol{x} \leq \boldsymbol{b}\}, \qquad \boldsymbol{X}_0 := \{\boldsymbol{x} \in \mathbb{N}^n : A\boldsymbol{x} \leq \boldsymbol{0}\}. \tag{66}$$

Then every vector in $\boldsymbol{X}$ can be decomposed as the sum of a vector $\boldsymbol{x} \in \boldsymbol{X}$ and a finite sum of vectors $\boldsymbol{x}_0 \in \boldsymbol{X}_0$ such that

$$\|\boldsymbol{x}\| \leq ((n+m+1)\|A\| + \|\boldsymbol{b}\| + 1)^m, \qquad \|\boldsymbol{x}_0\| \leq ((n+m)\|A\| + 1)^m. \tag{67}$$

Proof. The corollary holds trivially when $A = 0$, so we assume $A \neq 0$ here, especially $\|A\| \geq 1$. Let's define the following sets where $I_m \in \mathbb{Z}^{m \times m}$ is the identity matrix:

$$\boldsymbol{X}^+ := \left\{ (\boldsymbol{x}, \boldsymbol{y}) \in \mathbb{N}^n \times \mathbb{N}^m : \begin{bmatrix} A & I_m \end{bmatrix} \begin{bmatrix} \boldsymbol{x} \\ \boldsymbol{y} \end{bmatrix} = \boldsymbol{b} \right\}, \tag{68}$$

$$\boldsymbol{X}_0^+ := \left\{ (\boldsymbol{x}, \boldsymbol{y}) \in \mathbb{N}^n \times \mathbb{N}^m : \begin{bmatrix} A & I_m \end{bmatrix} \begin{bmatrix} \boldsymbol{x} \\ \boldsymbol{y} \end{bmatrix} = \boldsymbol{0} \right\}. \tag{69}$$

Let $\boldsymbol{x} \in \boldsymbol{X}$, there exists $\boldsymbol{y} \in \mathbb{N}^m$ with $(\boldsymbol{x}, \boldsymbol{y}) \in \boldsymbol{X}^+$. By Lemma C.1, there exists vectors $(\boldsymbol{x}', \boldsymbol{y}') \in \boldsymbol{X}^+$ and $(\boldsymbol{x}_1, \boldsymbol{y}_1), \ldots, (\boldsymbol{x}_k, \boldsymbol{y}_k) \in \boldsymbol{X}_0^+$ such that

$$\boldsymbol{x} = \boldsymbol{x}' + \sum_{i=1}^k \boldsymbol{x}_i, \quad \boldsymbol{y} = \boldsymbol{y}' + \sum_{i=1}^k \boldsymbol{y}_i, \tag{70}$$

where $\|\boldsymbol{x}'\| \leq ((n+m+1)\|A\| + \|\boldsymbol{b}\| + 1)^m$, and $\|\boldsymbol{x}_i\| \leq ((n+m)\|A\| + 1)^m$ for all $i = 1, \ldots, k$. Clearly $\boldsymbol{x}' \in \boldsymbol{X}$ and $\boldsymbol{x}_i \in \boldsymbol{X}_0$ for all $i = 1, \ldots, k$. This proves the desired result. ◀

The following lemma is just an application of Corollary C.2 to the characteristic systems of linear KLM sequences:

▶ **Lemma C.3.** *Let ξ be a linear KLM sequence. Then any model of $\mathcal{E}(\xi)$ can be decomposed into the sum of a model $\boldsymbol{h}$ of $\mathcal{E}(\xi)$ and a finite sum of models $\boldsymbol{h}_0$ of $\mathcal{E}^0(\xi)$, such that*

$$\|\boldsymbol{h}\|, \|\boldsymbol{h}_0\| \leq (10|\xi|)^{12|\xi|}. \tag{71}$$

Proof. Suppose $\xi = \langle p_0(\boldsymbol{x}_0) G_0 q_0(\boldsymbol{y}_0) \rangle \Lambda_1 \cdots \Lambda_k \langle p_k(\boldsymbol{x}_k) G_k q_k(\boldsymbol{y}_k) \rangle$ where $G_i = (Q_i, T_i)$ for each $i = 0, \ldots, k$. In order to apply Corollary C.2 to $\mathcal{E}(\xi)$, we need to bound the following four parameters:

- The number of variables in $\mathcal{E}(\xi)$ is

$$2d \cdot (k+1) + \sum_{i=0}^{k} |T_i| + \sum_{i=1}^{k} |\Lambda_i|_* \leq 2|\xi|. \tag{72}$$

- The number of inequalities in $\mathcal{E}(\xi)$ is bounded by

$$6d \cdot (k+1) + 2\sum_{i=0}^{k} |T_i| + \sum_{i=1}^{k} (|\Lambda_i|_* + d \cdot (2|\Lambda_i| + 2)) \leq 6|\xi|. \tag{73}$$

To see this, note that each Kirchoff system K_{G_i, p_i, q_i} consists of $|T_i|$ equations; each $\mathcal{E}_{\mathrm{LPS}}(\Lambda_i)$ consists of at most $|\Lambda_i|_* + 2d \cdot |\Lambda_i|$ inequalities and one equation; and the constraints $\boldsymbol{m}_i \sqsubseteq \boldsymbol{x}_i$, $\boldsymbol{n}_i \sqsubseteq \boldsymbol{y}_i$, and $\boldsymbol{n}_i = \boldsymbol{m}_i + \Delta(\phi_i)$ each consists of d equations. Also remember that one equation corresponds to two inequalities.

- The maximum absolute value of the coefficients in $\mathcal{E}(\xi)$ is bounded by

$$\max \left\{ 1, \max_{0 \leq i \leq k} \|T_i\|, \max_{1 \leq i \leq k} \|\Lambda_i\| \right\} \leq |\xi|. \tag{74}$$

Note that the coefficients in K_{G_i, p_i, q_i} and in constraints $\boldsymbol{m}_i \sqsubseteq \boldsymbol{x}_i$, $\boldsymbol{n}_i \sqsubseteq \boldsymbol{y}_i$ are just ± 1. The coefficients in $\mathcal{E}_{\mathrm{LPS}}(\Lambda_i)$ are bounded by $\|\Lambda_i\|$, and that in the constraints $\boldsymbol{n}_i = \boldsymbol{m}_i + \Delta(\phi_i)$ are bounded by $\|T_i\|$.

- The maximum absolute value of the constant terms in $\mathcal{E}(\xi)$ is bounded by

$$\max \left\{ 1, \max_{0 \leq i \leq k} \|\boldsymbol{x}_i\|, \max_{0 \leq i \leq k} \|\boldsymbol{y}_i\|, \max_{1 \leq i \leq k} \|\Lambda_i\| \right\} \leq |\xi|. \tag{75}$$

Note that the constraints $\boldsymbol{n}_i = \boldsymbol{m}_i + \Delta(\phi_i)$ contain no constant terms. The constant terms in constraints $\boldsymbol{m}_i \sqsubseteq \boldsymbol{x}_i$, $\boldsymbol{n}_i \sqsubseteq \boldsymbol{y}_i$ are the non-ω components of $\boldsymbol{x}_i$ and $\boldsymbol{y}_i$, and that in K_{G_i, p_i, q_i} are just ± 1. Finally, the constant terms in $\mathcal{E}_{\mathrm{LPS}}(\Lambda_i)$ are bounded by $\|\Lambda_i\|$.

Now by Corollary C.2 we can bound $\|\boldsymbol{h}\|$ and $\|\boldsymbol{h}_0\|$ by

$$\|\boldsymbol{h}\|, \|\boldsymbol{h}_0\| \leq ((2|\xi| + 6|\xi| + 1) \cdot |\xi| + |\xi| + 1)^{6|\xi|} \leq (10|\xi|)^{12|\xi|}. \qquad \blacktriangleleft$$

Lemma 5.7 now follows immediately from Lemma C.3.

C.3 Proof of Lemma 5.8

▶ **Lemma 5.8** ([15, Lem. 4.2]). *For any linear KLM sequence ξ which is not strongly connected, one can compute in time $\exp(|\xi|)$ a finite set Ξ of strongly connected linear KLM sequences such that $L_\xi = \bigcup_{\xi' \in \Xi} L_{\xi'}$ and that $\mathrm{rank}(\xi') \leq_{lex} \mathrm{rank}(\xi)$ and $|\xi'| \leq (2d+1)|\xi|$ for every $\xi' \in \Xi$.*

Proof. For each tuple $\zeta = \langle p(\boldsymbol{x}) G q(\boldsymbol{y}) \rangle$ occurring in ξ where G is not strongly connected, we replace it by all possible sequences of the form

$$\zeta' = \langle p(\boldsymbol{x}) G_0 r_0(\boldsymbol{\omega}) \rangle t_1 \langle s_1(\boldsymbol{\omega}) G_1 r_1(\boldsymbol{\omega}) \rangle t_2 \ldots \langle s_{n-1}(\boldsymbol{\omega}) G_{n-1} r_{n-1}(\boldsymbol{\omega}) \rangle t_n \langle s_n(\boldsymbol{\omega}) G_n q(\boldsymbol{y}) \rangle \tag{76}$$

where $G_0, \ldots, G_n$ are distinct strongly connected components of G connected by transitions $t_1, \ldots, t_n$. Clearly the action language is preserved, and the rank does not increase. As for the size, note that $n \leq |\zeta|$ and thus

$$|\zeta'| = \sum_{i=0}^{n} |G_i| + d \cdot (\|\boldsymbol{x}\| + \|\boldsymbol{y}\| + 1) + d \cdot n + \sum_{i=1}^{n} d \cdot (\|t_i\| + 1) \leq |\zeta| + 2nd \leq (2d+1)|\zeta|. \quad (77)$$

This bounds the amplification in sizes. ◀

C.4 Proof of Lemma 5.9

▶ **Lemma 5.9.** *Let ξ be a strongly connected linear KLM sequence. Whether ξ is pure is in* PSPACE. *If ξ is not pure, one can compute in space* poly($|\xi|$) *a finite set Ξ of pure linear KLM sequences such that $\bigcup_{\xi' \in \Xi} L_{\xi'} \subseteq L_\xi$ and $\bigcup_{\xi' \in \Xi} L_{\xi'} \neq \emptyset$ whenever $L_\xi \neq \emptyset$, and such that* rank(ξ') = rank(ξ) *and $|\xi'| \leq |\xi|^{O(1)}$ for all $\xi' \in \Xi$.*

Proof. Since ξ is strongly connected, by Corollary 5.4, to decide whether ξ is pure it is enough to check whether there is a non-trivial tuple $\langle p(\boldsymbol{x})Gq(\boldsymbol{y})\rangle$ in ξ such that G is geometrically 2-dimensional, i.e. $\dim(\mathrm{Cyc}(G)) \leq 2$. Observe that $\mathrm{Cyc}(G)$ is indeed spanned by the effects of all simple cycles in G, thus its dimension can be computed by enumerating the simple cycles, which can be done in PSPACE.

Suppose ξ is not pure, then for each non-trivial tuple $\langle p(\boldsymbol{x})Gq(\boldsymbol{y})\rangle$ occurring in ξ such that rank(G) = $\boldsymbol{0}$, we replace it by all possible sequences of the form

$$\langle p(\boldsymbol{x})\rangle \Lambda \langle q(\boldsymbol{y})\rangle \quad (78)$$

where Λ is a positive LPS compatible to G from state p to q satisfying $|\Lambda| \leq |G|^{O(1)}$ where the $O(1)$-term comes from Theorem 3.3. Let the resulting set of linear KLM sequences be Ξ, we claim that Ξ satisfies the lemma. Clearly every sequence ξ' in Ξ is pure. Note that rank(ξ') = rank(ξ) since we only replace rank-$\boldsymbol{0}$ tuples by rank-$\boldsymbol{0}$ sequences. Also note that as $|\Lambda| \leq |G|^{O(1)}$ and $\|\Lambda\| \leq |G|$ (because Λ is compatible to G), we have $|\xi'| \leq |\xi|^{O(1)}$ for all $\xi' \in \Xi$. Now we focus on the action languages.

Since we replace rank-$\boldsymbol{0}$ VASSes by positive LPSes compatible to them, every run admitted by the LPS must be admitted by the original VASS. Thus the inclusion $\bigcup_{\xi' \in \Xi} L_{\xi'} \subseteq L_\xi$ is clear. Let ξ be given by $\langle p_0(\boldsymbol{x}_0)G_0q_0(\boldsymbol{y}_0)\rangle \Lambda_1 \cdots \Lambda_k \langle p_k(\boldsymbol{x}_k)G_kq_k(\boldsymbol{y}_k)\rangle$. If $L_\xi \neq \emptyset$, consider a run $\pi = \pi_0\rho_1\pi_1 \ldots \rho_k\pi_k$ admitted by ξ of the form

$$p_0(\boldsymbol{m}_0) \xrightarrow{\pi_0} q_0(\boldsymbol{n}_0) \xrightarrow{\rho_1} p_1(\boldsymbol{m}_1) \xrightarrow{\pi_1} q_1(\boldsymbol{n}_1) \xrightarrow{\rho_2} \cdots \xrightarrow{\rho_k} p_k(\boldsymbol{m}_k) \xrightarrow{\pi_k} q_k(\boldsymbol{n}_k) \quad (79)$$

where $\boldsymbol{m}_i \sqsubseteq \boldsymbol{x}_i$ and $\boldsymbol{n}_i \sqsubseteq \boldsymbol{y}_i$. We will exhibit a run $\overline{\pi}$ admitted by some $\xi' \in \Xi$. For each $i = 0, \ldots, k$, if rank(G_i) $\neq \boldsymbol{0}$ or if the KLM tuple ξ_i is trivial, we let $\overline{\pi_i} = \pi_i$. Otherwise, rank(G_i) = $\boldsymbol{0}$ and G_i is geometrically 2-dimensional by Corollary 5.4. Since $p_i(\boldsymbol{m}_i) \xrightarrow{\pi_i} q_i(\boldsymbol{n}_i)$ holds in G_i, by Theorem 3.3 there is a positive LPS Λ_i compatible to G_i with $|\Lambda_i| \leq |G|^{O(1)}$ such that $p_i(\boldsymbol{m}_i) \xrightarrow{\Lambda_i} q_i(\boldsymbol{n}_i)$. In this case we let $\overline{\pi_i}$ be a path admitted by Λ_i so that $p_i(\boldsymbol{m}_i) \xrightarrow{\overline{\pi_i}} q_i(\boldsymbol{n}_i)$. Now observe that $\overline{\pi} := \overline{\pi_0}\rho_1\overline{\pi_1} \ldots \rho_k\overline{\pi_k}$ is admitted by some $\xi' \in \Xi$, witnessing $\bigcup_{\xi' \in \Xi} L_{\xi'} \neq \emptyset$. ◀

C.5 Proof of Lemma 5.11

▶ **Lemma 5.11.** *From any linear KLM sequence ξ, one can compute in time $\exp(g(|\xi|))$ a finite set* clean(ξ) *of clean linear KLM sequences such that $\bigcup_{\xi' \in \mathrm{clean}(\xi)} L_{\xi'} \subseteq L_\xi$ and*

$\bigcup_{\xi' \in \text{clean}(\xi)} L_{\xi'} \neq \emptyset$ *whenever* $L_\xi \neq \emptyset$. *Moreover, for every* $\xi' \in \text{clean}(\xi)$ *we have* $\text{rank}(\xi') \leq_{lex}$ $\text{rank}(\xi)$ *and* $|\xi'| \leq g(|\xi|)$ *where* $g(x) = x^{x^{O(1)}}$.

Proof. We first apply Lemma 5.8 to ξ to obtain a set of strongly connected linear KLM sequences with the action language preserved. Then we apply Lemma 5.9 to make them pure. The action language is partially preserved. These linear KLM sequences are then saturated using Lemma 5.10, which preserves action language. Finally, as unsatisfiable sequences have empty action languages by Lemma 5.6, we can safely remove them. The ranks of the linear KLM sequences never increased. As for the sizes, note that Lemma 5.8 increases $|\xi|$ linearly, while Lemma 5.9 increases $|\xi|$ by a polynomial function. After that, Lemma 5.10 produces sequences of size bounded by $\text{poly}(|\xi|)^{\text{poly}(|\xi|)} \leq |\xi|^{\text{poly}(|\xi|)}$. ◀

C.6 Proof of Lemma 5.12

▶ **Lemma 5.12.** *Let* ξ' *be a pure linear KLM sequence. For any linear KLM sequence* ξ' *with* $\text{rank}_{full}(\xi') <_{lex} \text{rank}_{full}(\xi)$, *we have* $\text{rank}(\xi') <_{lex} \text{rank}(\xi)$.

Proof. Let $\text{rank}_{\text{full}}(\xi) = (r_d, \dots, r_0)$. By the definition of pure sequences and Lemma 5.3, we must have $r_2 = r_1 = r_0 = 0$. Now let ξ' be a linear KLM sequence with $\text{rank}_{\text{full}}(\xi') <_{\text{lex}}$ $\text{rank}_{\text{full}}(\xi)$. Suppose $\text{rank}_{\text{full}}(\xi') = (r'_d, \dots, r'_0)$, then there exists $3 \leq i \leq d$ with $r'_i < r_i$ and $r'_j = r_j$ for all $i < j \leq d$. This implies $\text{rank}(\xi') <_{\text{lex}} \text{rank}(\xi)$. ◀

C.7 Proof of Lemma 5.16

▶ **Lemma 5.16.** *From any clean linear KLM sequences* ξ *that is not normal, one can compute in time* $\exp(h(\xi))$ *a finite set* $\text{dec}(\xi)$ *of clean linear KLM sequence such that* $\bigcup_{\xi' \in \text{dec}(\xi)} L_{\xi'} \subseteq$ L_ξ *and* $\bigcup_{\xi' \in \text{dec}(\xi)} L_{\xi'} \neq \emptyset$ *whenever* $L_\xi \neq \emptyset$. *Moreover, for every* $\xi' \in \text{dec}(\xi)$ *we have* $\text{rank}(\xi') <_{lex} \text{rank}(\xi)$ *and* $|\xi'| \leq h(\xi)$ *where* $h(x) = x^{x^{x^{O(1)}}}$.

Proof. If a clean linear KLM sequence ξ is not normal, it is either bounded, unpumpable, or not rigid. We apply one of Lemmas 5.13, 5.14, and 5.15 to decompose ξ into a set Ξ of linear KLM sequences such that $\bigcup_{\xi' \in \Xi} L_{\xi'} = L_\xi$, and $\text{rank}(\xi') <_{\text{lex}} \text{rank}(\xi)$ and $|\xi'| \leq |\xi|^{O(|\xi|)}$ for every $\xi' \in \Xi$. Then apply Lemma 5.11 to clean each sequence $\xi' \in \Xi$. We are done by setting $\text{dec}(\xi) := \bigcup_{\xi' \in \Xi} \text{clean}(\xi')$. ◀

C.8 Proof of Lemma 5.17

▶ **Lemma 5.17.** *Let* ξ *be a normal linear KLM sequence, then there is a word* $\sigma \in L_\xi$ *whose length is bounded by* $|\sigma| \leq \ell(|\xi|)$ *where* $\ell(x) \leq x^{O(x)}$.

The proof of 5.17 is basically just a repetition of that of [15, Lem. 4.19], except that we need to take additional care of the part of linear path schemes. Fix a normal linear KLM sequence $\xi = \langle p_0(\boldsymbol{x}_0) G_0 q_0(\boldsymbol{y}_0) \rangle \Lambda_1 \cdots \Lambda_k \langle p_k(\boldsymbol{x}_k) G_k q_k(\boldsymbol{y}_k) \rangle$ throughout this section, where $G_i = (Q_i, T_i)$ for each $i = 0, \dots, k$. First we need a bound on the models of $\mathcal{E}(\xi)$ and $\mathcal{E}^0(\xi)$.

▷ **Claim C.4.** There is a model $\boldsymbol{h}_0 \models \mathcal{E}^0(\xi)$ with $\|\boldsymbol{h}_0\| \leq |\xi|^{O(|\xi|)}$ such that for every $i = 0, \dots, k$, every $j \in [d]$ and every $t \in T_i$,

- $\boldsymbol{m}_i^{\boldsymbol{h}_0}(j) > 0$ whenever $\boldsymbol{x}_i(j) = \omega$,
- $\boldsymbol{n}_i^{\boldsymbol{h}_0}(j) > 0$ whenever $\boldsymbol{y}_i(j) = \omega$,
- $\phi_i^{\boldsymbol{h}_0}(t) > 0$.

Proof. Since ξ is saturated and unbounded, we are done by taking $\boldsymbol{h}_0$ as a sum of at most $2d(k+1) + \sum_{i=0}^{k} |T_i| \leq 2|\xi|$ models of $\mathcal{E}^0(\xi)$ each of whose norm is bounded by $|\xi|^{O(|\xi|)}$ by Lemma 5.7. ◁

▷ **Claim C.5.** There is a model $\boldsymbol{h} \models \mathcal{E}(\xi)$ with $\|\boldsymbol{h}\| \leq |\xi|^{O(|\xi|)}$ such that $\phi_i^{\boldsymbol{h}}(t) > 0$ for every $i = 0, \ldots, k$ and every $t \in T_i$.

Proof. By Lemma C.3 there is a model $\boldsymbol{h}' \models \mathcal{E}(\xi)$ with $\|\boldsymbol{h}'\| \leq |\xi|^{O(|\xi|)}$. We are done by taking $\boldsymbol{h} = \boldsymbol{h}' + \boldsymbol{h}_0$ where $\boldsymbol{h}_0$ is given by Claim C.4. ◁

We also need to borrow the following result from [15], which bounds the lengths of the runs that "pump" each input and output configurations in ξ.

▷ **Claim C.6** ([15, Claim C.4]). For each $i = 0, \ldots, k$ there are paths u_i, v_i of G_i with $|u_i|, |v_i| \leq |\xi|^{O(1)}$ and vectors $\boldsymbol{x}_i', \boldsymbol{y}_i' \in \mathbb{N}_\omega^d$ with $\boldsymbol{x}_i' \geq \boldsymbol{x}_i, \boldsymbol{y}_i' \geq \boldsymbol{y}_i$ such that
- $p_i(\boldsymbol{x}_i) \xrightarrow{u_i} p_i(\boldsymbol{x}_i')$ and $\boldsymbol{x}_i'(j) > \boldsymbol{x}_i(j)$ whenever $\text{FACC}_{G_i, p_i}(\boldsymbol{x}_i)(j) = \omega \neq \boldsymbol{x}_i(j)$ for every $j \in [d]$,
- $q_i(\boldsymbol{y}_i') \xrightarrow{v_i} q_i(\boldsymbol{y}_i)$ and $\boldsymbol{y}_i'(j) > \boldsymbol{y}_i(j)$ whenever $\text{BACC}_{G_i, q_i}(\boldsymbol{y}_i)(j) = \omega \neq \boldsymbol{y}_i(j)$ for every $j \in [d]$.

Now we can prove Lemma 5.17.

Proof of Lemma 5.17. Let $\boldsymbol{h}$ and $\boldsymbol{h}_0$ be the models of $\mathcal{E}(\xi)$ and $\mathcal{E}^0(\xi)$ given by Claim C.5 and C.4. For each $i = 0, \ldots, k$ let the paths u_i, v_i and the vectors $\boldsymbol{x}_i', \boldsymbol{y}_i'$ be given by Claim C.6. Denote by ψ_{u_i} and ψ_{v_i} the Parikh image of u_i and v_i, respectively.

Let $\phi_i := r_i \phi_i^{\boldsymbol{h}_0} - (\psi_{u_i} + \psi_{v_i})$ where we choose $r_i := \|T_i\| \cdot (|u_i| + |v_i|) + 1$. By Claim C.6 we have $r_i \leq |\xi|^{O(|\xi|)}$. Note that $\phi_i(t) > 0$ for every $t \in T_i$. Now observe that $\phi_i \models K_{G_i, p_i, q_i}^0$. Since G_i is strongly connected, by Euler's Lemma there is a path w_i in G_i whose Parikh image is ϕ_i. We have $|w_i| = \sum_{t \in T_i} \phi_i(t) \leq r_i |\xi|^{O(|\xi|)} \leq |\xi|^{O(|\xi|)}$. Also note that

$$\Delta(w_i) = \Delta(\phi_i) = r_i \Delta(\phi_i^{\boldsymbol{h}}) - (\Delta(u_i) + \Delta(v_i)) = r_i(\boldsymbol{n}_i^{\boldsymbol{h}_0} - \boldsymbol{m}_i^{\boldsymbol{h}_0}) - (\Delta(u_i) + \Delta(v_i)). \quad (80)$$

Let $\boldsymbol{m}_i^0 := r_i \boldsymbol{m}_i^{\boldsymbol{h}_0} + \Delta(u_i)$ and $\boldsymbol{n}_i^0 := r_i \boldsymbol{n}_i^{\boldsymbol{h}_0} - \Delta(v_i)$, then $\Delta(w_i) = \boldsymbol{n}_i^0 - \boldsymbol{m}_i^0$. Now we prove that $\boldsymbol{m}_i^0, \boldsymbol{n}_i^0 \geq \boldsymbol{0}$ and $\boldsymbol{m}_i^0(j), \boldsymbol{n}_i^0(j) > 0$ for every coordinate j not fixed by G_i. For coordinates j fixed by G_i we know that $\Delta(u_i)(j) = \Delta(v_i)(j) = 0$ since u_i, v_i are cycles, so $\boldsymbol{m}_i^0(j) = r_i \boldsymbol{m}_i^{\boldsymbol{h}_0}(j) \geq 0$ and $\boldsymbol{n}_i^0(j) = r_i \boldsymbol{n}_i^{\boldsymbol{h}_0}(j) \geq 0$. For coordinates j not fixed by G_i, if $\boldsymbol{x}_i(j) \in \mathbb{N}$ then $\text{FACC}_{G_i, p_i}(\boldsymbol{x}_i)(j) = \omega$ and $\Delta(u_i)(j) > 0$, so $\boldsymbol{m}_i^0(j) \geq \Delta(u_i)(j) > 0$; if $\boldsymbol{x}_i(j) = \omega$ then $\boldsymbol{m}_i^{\boldsymbol{h}_0}(j) > 0$ and we have $\boldsymbol{m}_i^0(j) \geq r - \Delta(u_i)(j) > 0$ since $r > \|T_i\| \cdot |u_i| \geq \Delta(u_i)(j)$. Symmetrically we can prove $\boldsymbol{n}_i^0(j) > 0$ for j not fixed by G_i.

Since u_i is a run from $\boldsymbol{x}_i = \boldsymbol{m}_i^{\boldsymbol{h}} + \omega \boldsymbol{m}_i^{\boldsymbol{h}_0}$, and $\Delta(u)_i(j) \leq \|T_i\| \cdot |u_i| < r_i$ for every $j \in [d]$, we deduce that u_i can be fired from the configuration $p_i(\boldsymbol{m}_i^{\boldsymbol{h}} + r_i \boldsymbol{m}_i^{\boldsymbol{h}_0})$. Then we have

$$p_i(\boldsymbol{m}_i^{\boldsymbol{h}} + r_i \boldsymbol{m}_i^{\boldsymbol{h}_0}) \xrightarrow{u_i} p_i(\boldsymbol{m}_i^{\boldsymbol{h}} + \boldsymbol{m}_i^0). \quad (81)$$

Let $s_i := 1 + \|T_i\| \cdot \max\{|w_i|, |T_i| \cdot \|\boldsymbol{h}\|\} \leq |\xi|^{O(|\xi|)}$, by monotony we have

$$p_i(\boldsymbol{m}_i^{\boldsymbol{h}} + s_i r_i \boldsymbol{m}_i^{\boldsymbol{h}_0}) \xrightarrow{u_i^{s_i}} p_i(\boldsymbol{m}_i^{\boldsymbol{h}} + s_i \boldsymbol{m}_i^0). \quad (82)$$

Symmetrically, we have

$$q_i(\boldsymbol{n}_i^{\boldsymbol{h}} + s_i \boldsymbol{n}_i^0) \xrightarrow{u_i^{s_i}} q_i(\boldsymbol{n}_i^{\boldsymbol{h}} + s_i r_i \boldsymbol{n}_i^{\boldsymbol{h}_0}). \quad (83)$$

Note that $s_i > |w_i| \cdot \|T_i\| \geq \Delta(w_i)(j)$ for all $j \in [d]$ and that $\boldsymbol{m}_i^0(j), \boldsymbol{n}_i^0(j) > 0$ for every coordinate j not fixed by G_i, we deduce

$$p_i(\boldsymbol{m}_i^h + s_i\boldsymbol{m}_i^0) \xrightarrow{w_i^{s_i}} p_i(\boldsymbol{m}_i^h + s_i\boldsymbol{n}_i^0). \tag{84}$$

Since $\phi_i^h \models K_{G_i,p_i,q_i}$ and $\phi_i^h(t) > 0$ for every $t \in T_i$, again by Euler's Lemma we know there exists a path σ_i from p to q whose Parikh image is ϕ_i^h. Observe that

$$p_i(\mathrm{FACC}_{G_i,p_i}(\boldsymbol{x}_i)) \xrightarrow{\sigma_i} q_i(\mathrm{BACC}_{G_i,q_i}(\boldsymbol{y}_i)). \tag{85}$$

Also note that $\mathrm{FACC}_{G_i,p_i}(\boldsymbol{x}_i) = \boldsymbol{m}_i^h + \omega\boldsymbol{n}_i^0$ and $\mathrm{BACC}_{G_i,q_i}(\boldsymbol{y}_i) = \boldsymbol{n}_i^h + \omega\boldsymbol{n}_i^0$. Since $|\sigma_i| = \sum_{t \in T_i} \phi_i^h(t) \leq |T_i| \|\boldsymbol{h}\|$, we have $s_i > |\sigma| \cdot \|T_i\| \geq \Delta(\sigma_i)(j)$ for all $j \in [d]$. Thus

$$p_i(\boldsymbol{m}_i^h + s_i\boldsymbol{n}_i^0) \xrightarrow{\sigma_i} q_i(\boldsymbol{n}_i^h + s_i\boldsymbol{n}_i^0). \tag{86}$$

Let $s := \max_{0 \leq i \leq k} s_i$ and $r := \max_{0 \leq i \leq k} r_i$, one can verify that the above discussion is still valid with r_i, s_i replaced by r, s. Combining (82), (84), (86) and (83), let $\pi_i := u_i^s w_i^s \sigma_i v_i^s$, we have

$$p_i(\boldsymbol{m}_i^h + sr\boldsymbol{m}_i^{h_0}) \xrightarrow{\pi_i} q_i(\boldsymbol{n}_i^h + sr\boldsymbol{n}_i^{h_0}) \tag{87}$$

holds for all $i = 0, \ldots, k$. Observe that $|\pi_i| \leq |\xi|^{O(|\xi|)}$.

Let $\boldsymbol{h}' := \boldsymbol{h} + sr\boldsymbol{h}_0$, then $\boldsymbol{h}' \models \mathcal{E}(\xi)$, especially $(\boldsymbol{n}_{i-1}^{h'}, \boldsymbol{e}_i^{h'}, \boldsymbol{m}_i^{h'}) \models \mathcal{E}_{\mathrm{LPS}}(\Lambda_i)$ for each $i = 1, \ldots, k$. By Lemma 4.2 there exists a path ρ_i admitted by Λ_i such that

$$q_{i-1}(\boldsymbol{n}_{i-1}^h + sr\boldsymbol{n}_{i-1}^{h_0}) \xrightarrow{\rho_i} p_i(\boldsymbol{m}_i^h + sr\boldsymbol{m}_i^{h_0}). \tag{88}$$

Note that $|\rho_i| \leq |\Lambda_i| + |\Lambda_i|_* \cdot |\Lambda_i| \cdot \left\|\boldsymbol{e}_i^{h'}\right\| \leq |\xi|^{O(|\xi|)}$.

Combining (87) and (88), we know that the run $\pi := \pi_0\rho_1\pi_1\ldots\rho_k\pi_k$ is admitted by ξ, thus $[\![\pi]\!] \in L_\xi \neq \emptyset$. Observe that $|\pi| \leq (2k+1) \cdot |\xi|^{O(|\xi|)} \leq |\xi|^{O(|\xi|)}$. ◀

D Proof of Lemma 6.3

▶ **Lemma 6.3** (cf. [21, Lem. A.5]). *Let $h : \mathbb{N} \to \mathbb{N}$ be a monotone inflationary function, let $a, b, c \geq 1$ and $x_0 \geq 0$ be natural numbers. If for all $x \geq x_0$, $h(x) \leq H^{\omega^b \cdot c}(x)$, then $h^{\omega^a}(x) \leq H^{\omega^{b+a}}((c+1)x)$ for all $x \geq \max\{2c, x_0\}$.*

Proof. We need to use the fact that if h is a monotone inflationary function, then so are the functions h^α in the Hardy hierarchy, which can be proved by induction on α. Let $G := H^{\omega^b \cdot c}$, first we show the following inequality holds for all $x \geq 2c$:

$$H^{\omega^{b+a}}((c+1)x) \geq (c+1)G^{\omega^a}(x). \tag{89}$$

Proof. We proceed by induction on a. For $a = 1$:

$$
\begin{aligned}
(c+1)G^\omega(x) &= (c+1)G^{x+1}(x) \\
&= (c+1)H^{\omega^b \cdot c(x+1)}(x) \\
&\leq H^{\omega^b \cdot c}(H^{\omega^b \cdot c(x+1)}(x)) && \text{since } H^{\omega^b \cdot c}(x) \geq H^{\omega \cdot c}(x) \geq 2^c \cdot x \\
&= H^{\omega^b \cdot c(x+2)}(x) \\
&\leq H^{\omega^b \cdot ((c+1)x+1)}(x) && \text{by } x \geq 2c \\
&\leq H^{\omega^b \cdot ((c+1)x+1)}((c+1)x) && \text{by monotony} \\
&= H^{\omega^{b+1}}((c+1)x).
\end{aligned}
$$

Next let's consider the case $a + 1$, assuming (89) holds for a. We first prove the following inequality holds for all $y \geq 2c$ by induction on j:

$$H^{\omega^{b+a} \cdot j}((c+1)y) \geq (c+1)G^{\omega^a \cdot j}(y). \tag{90}$$

The base case $j = 0$ holds trivially. For the induction step, note that

$$
\begin{aligned}
H^{\omega^{b+a} \cdot (j+1)}((c+1)y) &= H^{\omega^{b+a}}(H^{\omega^{b+a} \cdot j}((c+1)y)) \\
&\geq H^{\omega^{b+a}}((c+1)G^{\omega^a \cdot j}(y)) && \text{by induction hypothesis of (90)} \\
&\geq (c+1)G^{\omega^a}(G^{\omega^a \cdot j}(y)) && \text{by induction hypothesis of (89)} \\
&= (c+1)G^{\omega^a \cdot (j+1)}(y).
\end{aligned}
$$

Now we can prove (89) for $a + 1$:

$$
\begin{aligned}
H^{\omega^{b+a+1}}((c+1)x) &= H^{\omega^{b+a} \cdot ((c+1)x+1)}((c+1)x) \\
&\geq H^{\omega^{b+a} \cdot (x+1)}((c+1)x) \\
&\geq (c+1)G^{\omega^a \cdot (x+1)}(x) && \text{by (90)} \\
&= (c+1)G^{\omega^{a+1}}(x).
\end{aligned}
$$

This completes the proof of (89). $\triangleleft$

Now we return to the proof of the lemma. For all $x \geq \max\{2c, x_0\}$ we have

$$
\begin{aligned}
h^{\omega^a}(x) &\leq G^{\omega^a}(x) && \text{since } h(x) \leq G(x) \text{ for } x \geq x_0 \\
&\leq (c+1)G^{\omega^a}(x) \\
&\leq H^{\omega^{b+a}}((c+1)x) && \text{by (89).}
\end{aligned}
$$

This completes the proof. $\blacktriangleleft$