Reachability in Five-Dimensional Vector Addition Systems Is PSPACE-Complete

Weijun Chen
BASICS,
Shanghai Jiao Tong University
Shanghai, China
cwj2018@sjtu.edu.cn

Bo Fu
BASICS,
Shanghai Jiao Tong University
Shanghai, China
fubo0970@sjtu.edu.cn

Yuxi Fu[1]
BASICS,
Shanghai Jiao Tong University
Shanghai, China
fu-yx@cs.sjtu.edu.cn

Abstract

Vector addition systems (VAS) are a fundamental model of concurrent computation in which transitions update nonnegative integer counters by adding integer vectors, provided all counters remain nonnegative. Closely related to Petri nets, VAS are the single-state counterpart of vector addition systems with states (VASS); the latter use finitely many states to restrict the available transitions. A classical simulation of d-VASS by $(d+3)$-VAS relates their reachability problems; the complexity-theoretic picture of fixed-dimensional VAS, and VASS as well, remains however largely unclear. While reachability in 2-VASS is PSPACE-complete under binary encoding, the best known upper bound for 5-VAS was only elementary. We prove that every reachable 5-VAS instance has a reachability witness of length polynomial in its unary size. Consequently, 5-VAS reachability is NL-complete under unary encoding and PSPACE-complete under binary encoding. We also establish a polynomial factorization of the reachability relation of every d-VAS with $d \geq 3$. For 5-VAS, this strengthens Hopcroft and Pansiot's effective semilinearity of reachability sets to the full reachability relation. For every fixed $d \geq 3$, it also yields an F_d upper bound for reachability in $(d+3)$-VAS.

Key words: vector addition systems, reachability, complexity, semilinearity

[1]Corresponding author

Contents

1 Introduction **1**
- 1.1 Our Results . . . 2
- 1.2 Future Work . . . 3
- 1.3 Organization . . . 3

2 Technical Overview **4**
- 2.1 B-Encoding and Transformation Areas . . . 4
- 2.2 Polynomial Factorization of VAS . . . 5
- 2.3 Polynomial Reachability Distance in 5-VAS . . . 6

3 Preliminaries **13**

4 Generalized Transformation Areas for B-Encodings of VAS **17**
- 4.1 Steinitz Reordering and Cycle Relocation . . . 18
- 4.2 Proof of the Transformation-Area Theorem . . . 19
- 4.3 Stability of Transformation-Area Thresholds . . . 19

5 Polynomial Factorization of VAS Reachability **20**
- 5.1 Mixed LPS from Generalized Transformation Areas . . . 21
- 5.2 Linear Factorization of B-Encodings of VAS . . . 22
- 5.3 Applications of Polynomial Factorization . . . 24

6 CFZ Characterization and Basic Reachability-Distance Bounds **24**
- 6.1 Case I: Two Components with Cycle-Dimension at Most 2 . . . 25
- 6.2 Case II: A Single Encoding VAS . . . 27
- 6.3 Hybrid Sets as Sources and Targets . . . 30

7 Additive-Growth Induction for Semi-Coverable Instances **31**
- 7.1 Pumpable Source . . . 32
- 7.2 Translated Families . . . 37
- 7.3 Completing the Induction on the Number of SCCs . . . 38

8 Removing the Coverability Assumption **42**
- 8.1 Pumpable Sources and Hybrid Sets with a Nonzero Free Period . . . 42
- 8.2 Reaching a Fixed Translated Family . . . 42
- 8.3 The Final Induction . . . 43

AI Methodology **43**

Appendix **44**

References **65**

1 Introduction

Vector addition systems (VAS), introduced by Karp and Miller in the 1960s [KM69], are a fundamental model of concurrent computation. A d-dimensional VAS models d types of resources, each represented by a nonnegative integer counter, and a finite set of rules called transitions. A configuration is a nonnegative integer vector recording the values of these counters. Each transition updates a configuration by adding a fixed integer vector, and can be applied only if the resulting counters remain nonnegative. VAS are essentially equivalent in expressive power to Petri nets, widely used to model protocols, workflows, and chemical and biological processes [Esp96, Mur89, ZZ94].

The central decision problem for VAS is *reachability*: given a VAS V and two configurations $\boldsymbol{s}$ and $\boldsymbol{t}$, determine whether a *run*, i.e., a finite sequence of transitions executable from $\boldsymbol{s}$, reaches $\boldsymbol{t}$. Its decidability was established by Mayr in 1981 [May81] with subsequent refinements by Kosaraju [Kos82] and Lambert [Lam92], leading to the KLM algorithm. Its precise computational complexity remained open for decades. Leroux and Schmitz established an Ackermannian upper bound in 2019 [LS19], while Czerwiński and Orlikowski [CO21] and, independently, Leroux [Ler21] established matching lower bounds in 2021, settling reachability as Ackermann-complete.

Much of the complexity theory of reachability has been developed for a closely related model called *vector addition systems with states* (VASS), which extend VAS with finitely many states restricting the available transitions. While every VAS is a VASS with a single state, Hopcroft and Pansiot showed that every d-VASS can be simulated by a $(d+3)$-VAS [HP79]. Thus, the two models are equivalent when dimension is unfixed, but the three additional counters make the fixed-dimensional complexity comparison nontrivial. For VASS, the exact complexity of reachability under binary encoding is known only in dimensions at most two: NP-complete for 1-VASS [HKOW09] and PSPACE-complete for 2-VASS [BEF+21]. Only recently was reachability in 3-VASS shown to be elementary [CJLO25], with its upper bound later improved to EXPSPACE [CFF+26].

These advances for low-dimensional VASS make substantial use of the structure of two-dimensional reachability sets and relations. In particular, for a fixed source configuration, its *reachability set* consists of all reachable targets, whereas the *reachability relation* consists of all reachable source-target pairs. A set or relation is *semilinear* if it is a finite union of sets obtained from a base vector by adding nonnegative integer combinations of finitely many period vectors. Hopcroft and Pansiot established that the reachability sets of 2-VASS are effectively semilinear [HP79], while Leroux and Sutre proved that even the full reachability relation is effectively semilinear [LS04]. Remarkably, Hopcroft and Pansiot also proved that the reachability sets of 5-VAS are effectively semilinear [HP79]. This parallel raises two natural questions: firstly, is the reachability relation of 5-VAS semilinear, and secondly, is reachability in 5-VAS PSPACE-complete under binary encoding, as in 2-VASS? The latter problem is already PSPACE-hard by the three-counter simulation, whereas only an elementary upper bound was established recently [CFZ26].

The relationship between VAS and VASS also motivates improvements in higher dimensions. For d-VASS, reachability is known to be in F_d [FYZ24], the d-th level of the fast-growing complexity hierarchy. In particular, F_3 corresponds to tower-of-exponentials complexity, while F_ω corresponds to Ackermannian complexity [Sch16]. Our current knowledge about lower bounds is that $(2k+3)$-VASS is F_k-hard [CJL+23]. Huge gaps remain between the known upper and lower bounds. For a long time, d-VAS inherited the same upper bound, until Chen, Fu, and Zheng improved it to F_{d-2} [CFZ26] by exploiting the absence of states. In view of the three-counter simulation, one may

ask whether reachability in $(d+3)$-VAS admits the same F_d upper bound as reachability in d-VASS.

1.1 Our Results

Reachability Complexity. Our main result establishes the exact complexity of reachability in 5-VAS under both binary and unary encoding, where encoding schemes refer to how the integers specifying the transitions and the configurations are represented.

Theorem 1.1. *The reachability problem for* 5*-VAS is* PSPACE*-complete under binary encoding and* NL*-complete under unary encoding.*

These results match those for 2-VASS, whose reachability is also PSPACE-complete under binary encoding and NL-complete under unary encoding [BEF+21]. Recently, Kamiński and Lasota established PSPACE-hardness [KL27] for 3-VAS, which, together with the known upper bounds [CFZ26], settled the binary complexity of reachability in dimensions three and four. Combined with our result, this establishes the same PSPACE (binary) and NL (unary) completeness for every $d \in \{3,4,5\}$.

To prove the upper bounds in Theorem 1.1, we bound the length of shortest runs. The *reachability distance* $\ell_V(\boldsymbol{s},\boldsymbol{t})$ is the length of a shortest run from $\boldsymbol{s}$ to $\boldsymbol{t}$ in V if such a run exists. We write $\mathrm{SIZE}(V,\boldsymbol{s},\boldsymbol{t})$ for the *unary size* of the instance, which accounts for the number of transitions, the norms of their effects, and the values of the source and target configurations.

Theorem 1.2 (Polynomial Reachability Distance)**.** *There exists a constant c such that for every reachable 5-VAS instance $(V,\boldsymbol{s},\boldsymbol{t})$, one has $\ell_V(\boldsymbol{s},\boldsymbol{t}) \le \mathrm{SIZE}(V,\boldsymbol{s},\boldsymbol{t})^c$.*

A nondeterministic algorithm can guess a run of length at most M^c using $\mathcal{O}(\log M)$ space to store the current configuration and a step counter, where $M = \mathrm{SIZE}(V,\boldsymbol{s},\boldsymbol{t})$. This yields NL membership under unary encoding and PSPACE membership under binary encoding.

Factorization of Reachability Relations. Theorem 1.2 bounds the length of a witness for each reachable source-target pair, but does not describe the structure of the whole reachability relation. For every fixed $d \ge 3$, we develop a polynomial factorization of the reachability relation of every d-VAS. A *linear path scheme* (LPS) $\Lambda = \alpha_0\beta_1^*\alpha_1\cdots\beta_k^*\alpha_k$ consists of fixed paths $\alpha_0,\dots,\alpha_k$ interleaved with repeatable cycles $\beta_1,\dots,\beta_k$, providing a standard representation of reachability in low-dimensional VASS [BEF+21]. We extend them to *mixed linear path schemes* (mixed LPS), which allow different components to be interpreted under either ordinary nonnegative semantics or integer semantics.

We call a representation of the reachability relation of a VAS V of the form

$$\mathrm{Reach}_V = \bigcup_{(R_1,\dots,R_m)\in\mathcal{F}} R_1 \circ \cdots \circ R_m$$

a *polynomial factorization* if $\mathcal{F}$ is finite and each R_i is either represented by a mixed LPS or is the reachability relation of V with three counters bounded by some B. The latter is essentially captured by a $(d-3)$-VASS, obtained by encoding the bounded counters into states. The length m of each factorization, the bounds B, and all mixed-LPS parameters must be bounded by a polynomial in $\mathrm{SIZE}(V)$. Moreover, every witness for a pair in R_i lifts to a run in V of no greater length.

Theorem 1.3. *For fixed $d \ge 3$, the reachability relation of every d-VAS admits a polynomial factorization.*

For 5-VAS, the factors with three bounded counters can be reduced to 2-VASS by encoding these counters into states. Since the reachability relation of 2-VASS can be captured by LPS [BEF+21], these factors can also be represented by LPS with polynomially bounded parameters.

Corollary 1.4. *For every 5-VAS V, its reachability relation can be represented by a finite set of mixed LPS, with all parameters polynomially bounded in* $\mathrm{SIZE}(V)$.

Reachability through a mixed LPS can be described by linear inequalities over integers, with variables describing cycle repetitions and intermediate configurations. Thus, mixed LPS define semilinear relations, while retaining a path-and-cycle structure that is more convenient for VASS constructions. Our corollary therefore establishes semilinearity of the reachability relation of 5-VAS, strengthening the classical result on reachability sets with a fixed source [HP79].

Higher-Dimensional Complexity. Theorem 1.3 also yields an improved upper bound for reachability in higher-dimensional VAS. For a $(d+3)$-VAS, the factors can be simulated by VASS whose strongly connected components have at most d independent cycle effects, even though these VASS may have more than d counters. Applying the component-dimension reachability bound of [CFZ26], which builds on the improved KLM algorithm [FYZ24], yields the following result.

Theorem 1.5. *Reachability in $(d+3)$-VAS is in F_d for every fixed $d \geq 3$, even under binary encoding.*

1.2 Future Work

Taking semilinearity as a dividing line, the complexity of low-dimensional VAS reachability is not yet fully understood. In particular, the exact complexity of 2-VAS reachability remains open. It is therefore natural to ask whether 2-VAS shares the complexity of 3-, 4-, and 5-VAS: PSPACE-complete under binary encoding and NL-complete under unary encoding. A positive answer would be surprising, establishing the same complexity across four consecutive dimensions. More broadly, it would suggest that reachability complexity may increase in steps rather than with every additional counter, motivating the search for substantially better upper bounds in higher dimensions, perhaps even $\mathsf{F}_{\lceil d/2 \rceil}$ for d-VASS instead of the currently known F_d.

Beyond the semilinearity boundary, a similar comparison arises between 3-VASS and 6-VAS. Both models can have non-semilinear reachability sets [HP79]. However, their complexity analysis is quite different. Reachability in 3-VASS has recently been shown to be in EXPSPACE [CFF+26], whereas our result places 6-VAS reachability in F_3, leaving its elementary membership open.

Another direction is to investigate whether polynomial factorization can also help analyze VASS reachability. Every single-state SCC is itself a VAS, so Theorem 1.3 applies to its reachability relation, replacing it with compositions of mixed LPS and factors with three bounded counters. It would be interesting to investigate whether these local factorizations can be incorporated into existing methods, such as the KLM algorithm, to improve reachability bounds for VASS.

1.3 Organization

Section 2 provides a technical overview of our results, followed by preliminary definitions and tools in Section 3. We then establish the generalized transformation area theorem in Section 4 and prove the polynomial factorization theorem in Section 5. Sections 6, 7 and 8 are devoted to the proof of the polynomial reachability-distance bound for 5-VAS.

2 Technical Overview

We now outline the proofs of our two main technical results: the polynomial reachability-distance bound in Theorem 1.2 and the polynomial factorization theorem in Theorem 1.3. The complexity classification in Theorem 1.1 follows from the former, while the mixed LPS characterization for 5-VAS and higher-dimensional complexity results in Corollary 1.4 and Theorem 1.5 follow from the latter. We defer the proofs of these consequences to Appendix A. The two main theorems share a common tool: a generalized transformation-area theorem for VAS with two bounded coordinates encoded into states. We introduce this tool in Section 2.1, explain how it leads to a polynomial factorization of reachability in Section 2.2, and focus on the proof of Theorem 1.2 in Section 2.3.

2.1 B-Encoding and Transformation Areas

A standard technique for studying VAS is to extract bounded coordinates and encode them into states, thereby reducing the system to a lower-dimensional VASS. Although the resulting VASS has multiple states, it preserves some freedom to rearrange transitions inherited from the single-state structure of VAS. We formalize this construction as a B-encoding and then extend the classical transformation-area technique of van Leeuwen [vL74] and Hopcroft-Pansiot [HP79] to this setting.

B-Encoding. Consider a run of a 5-VAS V in which the first two coordinates remain bounded by B. We can store their values in states, while keeping the remaining three coordinates as counters. For example, a configuration $(u_1, u_2, x_1, x_2, x_3)$ is represented by the state (u_1, u_2) with counter vector (x_1, x_2, x_3). To define the transitions, write each transition of V as $\boldsymbol{a} = (\boldsymbol{b}, \boldsymbol{c}) \in \mathbb{Z}^2 \times \mathbb{Z}^3$. For every state $\boldsymbol{u} \in [0, B]^2$, the encoding contains the transition

$$\boldsymbol{u} \xrightarrow{\boldsymbol{c}} \boldsymbol{u} + \boldsymbol{b} \qquad \text{whenever} \qquad \boldsymbol{u} + \boldsymbol{b} \in [0, B]^2.$$

Thus, $\boldsymbol{b}$ updates the state, while $\boldsymbol{c}$ updates the remaining counters, subject to their nonnegativity. In this way, the resulting 3-VASS represents exactly the runs of V that respect these bounds.

The same construction applies to any d-VAS V and subset $I \subseteq [d]$. We write $[\![V]\!]_{I,B}$ for the resulting $(d - |I|)$-VASS and call it a *B-encoding*. For the rest of the discussion, we mainly consider $|I| = 2$. By permuting coordinates, we assume $I = \{1, 2\}$, i.e., the first two coordinates, and abbreviate $[\![V]\!]_{I,B}$ as $[\![V]\!]_B$, when there is no ambiguity. The *Parikh image* of a path in $[\![V]\!]_B$ counts the occurrences of transitions of V, treating their copies at different states as the same transition.

Transformation Area. Reordering the transitions of a VAS preserves their total effect, but may cause intermediate configurations to become negative. A *$\mathbb{Z}$-run* is a transition sequence interpreted without requiring counter values to remain nonnegative. Van Leeuwen [vL74] introduced *transformation areas*, where every $\mathbb{Z}$-run between two configurations can be rearranged into a valid run with the same Parikh image. He showed that such a region can be obtained by requiring all but one coordinate of the endpoints to be sufficiently large. Hopcroft and Pansiot [HP79] extended this technique to $\mathbb{Z}$-runs that are already valid on two bounded coordinates, requiring only the remaining coordinates to be sufficiently large at the endpoints. We extend this technique to B-encodings.

Let V be a $(d+2)$-VAS, and identify each configuration $\boldsymbol{p}(\boldsymbol{x})$ of $[\![V]\!]_B$ with $(\boldsymbol{p}, \boldsymbol{x}) \in \mathbb{N}^{d+2}$ in V. A set $\mathcal{A} \subseteq [0, B]^2 \times \mathbb{N}^d$ is a *generalized transformation area* for $[\![V]\!]_B$ if for every $\boldsymbol{s}, \boldsymbol{t} \in \mathcal{A}$ and every

$\mathbb{Z}$-run $s \xrightarrow{\pi}_{\mathbb{Z}} t$ in $[\![V]\!]_B$, there is a run $s \xrightarrow{\rho} t$ in V with the same Parikh image as π. We focus on generalized transformation areas of the form

$$\mathcal{G}_C = [0, B]^2 \times \{x \in \mathbb{N}^d : x \geq C \cdot \mathbf{1}\}.$$

We call C a *transformation-area threshold* for $[\![V]\!]_B$ if $\mathcal{G}_C$ is a generalized transformation area. We establish in Theorem 4.2 an explicit polynomial bound on this threshold: every $C \geq (d+2)(B+1)^2 \|T\|$ is a transformation-area threshold for $[\![V]\!]_B$, where T is the transition set of V.

The proof builds on the *cycle-relocation* technique in [HP79]. We call a state visited by a run a *peak state* if its first encoded coordinate is maximal among all states visited by the run. We choose an occurrence of a peak state q along the given $\mathbb{Z}$-run and extract simple cycles from the prefix and suffix around it. For each extracted cycle, choose a state w minimizing its second encoded coordinate and rotate the cycle to start there. We then relocate the cycle to q by translating each of its encoded states z to $q + z - w$. Moreover, all translated states lie in $[0, 2B]^2$: the first coordinate remains nonnegative by the choice of q, and the second by the choice of w. Steinitz's lemma [Ste13] then provides an ordering that keeps the remaining counters nonnegative when the endpoints are sufficiently large. Hence, our construction guarantees a run in $[\![V]\!]_{2B}$, which lifts to a run in V.

2.2 Polynomial Factorization of VAS

We then outline the proof of Theorem 1.3. Our proof repeatedly applies a factorization procedure for B-encodings. Each application produces mixed-LPS factors and factors with one additional bounded original coordinate, and only the latter require further processing. By introducing dummy coordinates with zero transition effects when necessary, three rounds suffice to obtain the desired factorization of an arbitrary d-VAS. We now explain the procedure for a single B-encoding.

2.2.1 Mixed LPS Between Configurations in $\mathcal{G}_C$

We first explain the easy case: how mixed-LPS factors arise from the generalized transformation area. Consider a $\mathbb{Z}$-run $p(x) \xrightarrow{\pi}_{\mathbb{Z}} q(y)$ in a B-encoding $[\![V]\!]_B$, with both endpoints in $\mathcal{G}_C$. It can be decomposed by removing simple cycles, leaving a short path together with the removed cycles. This allows us to capture its endpoints by a $\mathbb{Z}$-LPS Λ of polynomially bounded size under $\mathbb{Z}$-semantics. To ensure that both endpoints lie in $\mathcal{G}_C$, we add a gadget at each state $r \in \{p, q\}$, consisting of two vectors $-(r, C \cdot \mathbf{1})$ and $(r, C \cdot \mathbf{1})$ under $\mathbb{N}$-semantics. Combining these checks with Λ gives a mixed LPS whose reachability relation is contained in that of V.

However, $\mathcal{G}_C$ alone is insufficient to factor arbitrary runs. To see the difficulty, consider a 5-VAS V and its B-encoding $[\![V]\!]_B$, which is a 3-VASS. At each state of the encoding, the transformation area corresponds to the region $C \cdot \mathbf{1} + \mathbb{N}^3$. Its complement consists of three overlapping regions near the coordinate planes, each defined by

$$\mathcal{A}_i := \{x \in \mathbb{N}^3 : x(i) < C\}, \qquad i \in [3].$$

A run that stays entirely within one such region can be reduced to a 2-VASS by encoding the corresponding bounded coordinate into states. However, a run may alternate between these regions arbitrarily many times. Simply splitting a run according to which coordinate remains bounded does not yield a bound on the number of factors.

2.2.2 Factorization of Arbitrary Runs via High Positions

To bound the number of factors, we develop an iterative factorization procedure in Section 5.

Rackoff's Extraction and High Positions. The first ingredient is Rackoff's extraction lemma [LS19, CFZ26]. For a fixed-dimensional VASS U and a threshold C, there exists $H = (C + \textsc{size}(U))^{\mathcal{O}(1)}$ such that if every counter reaches at least H somewhere along a run π, then deleting cycles from π yields a run of length at most H that ends in the same state with all counters at least C.

We choose C to be a transformation-area threshold for $[\![V]\!]_{2B}$ and let H be given by Rackoff's extraction. Consider a run π in $[\![V]\!]_B$. We call a position of π *high* if every counter reaches at least H somewhere in both the prefix ending at that position and the suffix beginning at that position. At such a position with state $\boldsymbol{r}$, applying Rackoff's extraction to the prefix and the reversed suffix yields two short runs $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{\alpha} \boldsymbol{r}(\boldsymbol{u})$ and $\boldsymbol{r}(\boldsymbol{v}) \xrightarrow{\beta} \boldsymbol{q}(\boldsymbol{y})$ where $\boldsymbol{u}, \boldsymbol{v} \geq C \cdot \mathbf{1}$. Thus, both intermediate endpoints lie in $\mathcal{G}_C$, at the same state $\boldsymbol{r}$. If $\boldsymbol{r}$ is also a peak state, the deleted cycles can be relocated to $\boldsymbol{r}$ to obtain a $\mathbb{Z}$-run $\boldsymbol{r}(\boldsymbol{u}) \xrightarrow{*}_{\mathbb{Z}} \boldsymbol{r}(\boldsymbol{v})$ in $[\![V]\!]_{2B}$. Then the mixed-LPS construction described above allows us to represent π by a mixed LPS with polynomially bounded parameters.

High-Interval Recursion. The high positions of a run form a contiguous interval, which determines a possibly empty middle sub-run π_{mid}. The prefix and suffix outside π_{mid} each have a coordinate bounded by $H + \textsc{size}(V)$, where the additional $\textsc{size}(V)$ accounts for the boundary transitions. Otherwise, it suffices to proceed with π_{mid}, which contains no peak state of π. Thus, the first encoded coordinate of a peak state of π_{mid} is strictly smaller than that of π. Since this value lies in $[0, B]$, the procedure terminates after at most $B + 1$ steps, producing at most $2B + 2$ factors.

2.3 Polynomial Reachability Distance in 5-VAS

We now turn from the uniform relation decomposition to the polynomial length bound in Theorem 1.2, allowing the argument to use the given source $\boldsymbol{s}$ and target $\boldsymbol{t}$. Our starting point is the characterization of Chen, Fu, and Zheng [CFZ26], which reduces the polynomial reachability-distance problem for 5-VAS to two classes of instances. The characterization relies on the *cycle-dimension* of a VASS V, defined as the dimension of the space $\textsc{cyc}(V)$ spanned by the effects of its simple cycles.

The CFZ Characterization. We recall the characterization specialized to dimension five. For every reachable 5-VAS instance $(V, \boldsymbol{s}, \boldsymbol{t})$, their characterization provides a reachable VASS instance (V', s', t') of polynomial size such that $\ell_V(\boldsymbol{s}, \boldsymbol{t}) \leq \ell_{V'}(s', t')$. The new instance has one of two forms.

I. V' consists of two 5-VASS of cycle-dimension at most 2, connected by a zero-effect transition. Both components are obtained by encoding two original coordinates into states.

II. V' is a single B-encoding $[\![V]\!]_B$, where $B \leq \textsc{size}(V, \boldsymbol{s}, \boldsymbol{t})^{\mathcal{O}(1)}$.

It therefore suffices to establish polynomial reachability-distance bounds for these two classes. The first case can be handled using existing hybrid-set results for VASS of cycle-dimension at most two. The second case presents the main challenge, as its SCCs may have full three-dimensional cycle spaces. We focus on this case after explaining how to handle the first.

Case I: Hybrid-Set Representations. VASS of cycle-dimension at most two have been studied in [FYZ24, CCM+24, Zhe25, CJLO25]. These techniques were further developed in [CFF+26], yielding efficient hybrid-set representations and short reachability-distance results for geometrically 2-dimensional 3-VASS. A *hybrid set* has the form

$$H = \boldsymbol{b} + P^{\langle \boldsymbol{c}, \cdot \rangle \leq K} := \{\boldsymbol{b} + P\boldsymbol{\lambda} : \boldsymbol{\lambda} \in \mathbb{N}^m \text{ and } \langle \boldsymbol{c}, \boldsymbol{\lambda} \rangle \leq K\},$$

where $\boldsymbol{b}$ is its base, $P = \{\boldsymbol{p}_1, \ldots, \boldsymbol{p}_m\}$ is a set of nonnegative periods, and $\boldsymbol{c} \in \mathbb{N}^m$ assigns a cost to each period. The parameter K bounds the total cost of the selected period occurrences. A period is *free* if its cost is zero and *restricted* otherwise. We measure the size of a hybrid-set representation by its *representation norm* $\text{NORM}(H) := \max\{\|\boldsymbol{b}\|, \|P\|, \|\boldsymbol{c}\|\}$, which excludes K.

Hybrid-set representations are a key tool for eliminating components. They characterize the reachable configurations of a component at a selected state, allowing the corresponding periods to be added as self-loops in the remaining VASS. The only difficulty is that the added self-loops do not enforce the total-cost bound K. However, no polynomial bound on K is needed. Given a short run of length at most ℓ, if $K \geq \|\boldsymbol{c}\| \ell$, its period occurrences automatically satisfy the cost bound. Otherwise, $K < \|\boldsymbol{c}\| \ell$, and we can absorb the restricted periods into the bases, expressing H as a finite union of hybrid sets with only free periods and base norms at most $\|\boldsymbol{b}\| + K\|P\| \leq \|\boldsymbol{b}\| + \ell\|\boldsymbol{c}\|\|P\|$. The reachability of each resulting hybrid set can be faithfully simulated by an augmented VASS. Thus polynomial bounds on $\text{NORM}(H)$ suffice, independently of K.

Let V_1 and V_2 denote the two components in Case I, which may encode different coordinates. Existing hybrid-set results [CFF+26] describe the configurations reachable at the bridge through V_1. To connect with V_2, we lift a hybrid set containing the original bridge configuration to $\mathbb{N}^5$, fix the two coordinates encoded by V_2, and project onto the remaining three coordinates. This requires absorbing some period occurrences into the base, which increases the representation norm by at most a polynomial factor. Applying the reachability-distance theorem for geometrically 2-dimensional 3-VASS [CFF+26] to V_2 and V_1 yields a polynomial-length run through the bridge.

Case II: Main Challenge in B-Encodings. We now consider a reachable instance in the 3-VASS $W = [\![V]\!]_B$. Fixing the SCCs and connecting transitions (called *bridges*) visited by a witnessing run π gives a *sequential sub-VASS*

$$U = (U_1)u_1(U_2)\cdots u_{k-1}(U_k).$$

Note that these SCCs may have cycle-dimension three, and their number $k \leq (B+1)^2$ is not bounded by a constant. Throughout the analysis of Case II, we write $s = \boldsymbol{p}(\boldsymbol{x})$ and $t = \boldsymbol{q}(\boldsymbol{y})$ for the source and target configurations in U, where $\boldsymbol{p}, \boldsymbol{q} \in [0, B]^2$ are states of the encoding. Their corresponding configurations in the original VAS V are $\boldsymbol{s} = (\boldsymbol{p}, \boldsymbol{x})$ and $\boldsymbol{t} = (\boldsymbol{q}, \boldsymbol{y})$.

We begin with the basic cases, namely non-wide and coverable instances, for which polynomial-length witnesses can be obtained directly. The remaining cases are handled by *additive growth induction* over the number of SCCs k in Section 2.3.2.

2.3.1 Basic Cases: Wideness and Coverability

Recall that the cycle space $\text{CYC}(U)$ is spanned by the effects of simple cycles in U. However, since cycles can only be used with nonnegative multiplicities, the cycle space does not capture which

directions can be generated by their effects. We therefore consider the *cone* $\mathrm{cone}(U)$, generated by the effects of simple cycles with nonnegative rational coefficients. We call U *wide* if $\mathrm{cone}(U) = \mathbb{Q}^3$ and *non-wide* otherwise.

Non-Wide Decomposition. Suppose that $\mathrm{cone}(U) \neq \mathbb{Q}^3$. Geometrically, the cycle cone must lie in a closed half-space through the origin. By a separating-hyperplane argument as in [CJLO25, CFZ26], there exists a nonzero integer vector $\boldsymbol{h}$ of polynomial norm such that $\langle \boldsymbol{h}, \mathrm{EFF}(\theta) \rangle \geq 0$ for every simple cycle θ in U. Along any run from s to t, deleting simple cycles from the prefix and suffix around an intermediate configuration $\boldsymbol{r}(\boldsymbol{z})$ shows that $\langle \boldsymbol{h}, \boldsymbol{z} \rangle$ is polynomially bounded in $\mathrm{SIZE}(U, s, t)$. All intermediate configurations of such runs are therefore confined to polynomially many parallel affine planes of the form $\langle \boldsymbol{h}, \boldsymbol{z} \rangle = m$.

By recording this value in states, we obtain a polynomial-size 3-VASS whose cycle effects all lie in the two-dimensional space $\boldsymbol{h}^{\perp}$. The resulting VASS has cycle-dimension at most two, so the hybrid-set techniques from Case I yield a polynomial-length witness.

Coverability in the Wide Case. Suppose that U is wide, namely $\mathrm{cone}(U) = \mathbb{Q}^3$. Then every cycle effect can be canceled by inserting other cycles. If the source and target can be moved into $\mathcal{G}_C$, respectively, such cancellation yields a $\mathbb{Z}$-run between the new endpoints. Shortening this $\mathbb{Z}$-run and applying the transformation-area theorem then gives a polynomial-length run.

Let C be a transformation-area threshold for $[\![V]\!]_B$. We say that the source $s = \boldsymbol{p}(\boldsymbol{x})$ is *forward C-coverable* if $s \xrightarrow{*} \boldsymbol{p}(\boldsymbol{u})$ for some $\boldsymbol{u} \geq C \cdot \mathbf{1}$. Symmetrically, the target $t = \boldsymbol{q}(\boldsymbol{y})$ is *backward C-coverable* if $\boldsymbol{q}(\boldsymbol{v}) \xrightarrow{*} t$ for some $\boldsymbol{v} \geq C \cdot \mathbf{1}$. We call the instance *$C$-coverable* if both conditions hold.

When U is wide and C-coverable, we first extract short covering runs at both endpoints by Rackoff's covering bound [Rac78]. By canceling their effects, we obtain a $\mathbb{Z}$-run connecting two configurations in $\mathcal{G}_C$. Since $\mathbb{Z}$-reachability can be characterized using linear equations, Pottier's theorem [Pot91] yields a polynomial-length $\mathbb{Z}$-run. The generalized transformation-area theorem then converts it into a valid run in the original VAS without changing its length. Together with the covering runs, this establishes the polynomial reachability-distance bound in Theorem 1.2.

2.3.2 Bounded Prefixes and Additive-Growth Induction

We remove the remaining C-coverability assumptions in two stages while preserving a uniform polynomial bound. We first consider instances whose source is forward C-coverable or whose target is backward C-coverable, which we call *semi-C-coverable*. Since transformation-area thresholds are preserved under reversal, we may assume that the target is backward C-coverable. We then remove this final assumption using the same additive-growth induction. The two stages are proved in Sections 7 and 8, respectively.

Main Obstacle. Assume that $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{\pi} \boldsymbol{q}(\boldsymbol{y})$ is a run in $U = (Q, T)$. The wide coverable cases have already been handled. Rackoff's extraction implies that, if $\boldsymbol{p}(\boldsymbol{x})$ is not forward C-coverable, some coordinate $i \in [3]$ stays below a polynomial threshold D along the segment of π in the first SCC. Encoding this coordinate into states yields a 2-VASS, whose reachable configurations can be described by hybrid sets as in Case I. Following the approach of [CJLO25, CFF+26], one could adjoin the periods of a selected hybrid set to the next component and continue with fewer SCCs.

However, polynomial bounds for individual reduction steps do not suffice when they are repeatedly composed. Even a recurrence as simple as $M_{j+1} \leq M_j^c$ yields $M_k \leq M_0^{c^k}$. Since the number k of SCCs is part of the input, the resulting polynomial degree is not fixed. Thus, we need to avoid repeated polynomial composition altogether.

Bounded Coordinate Extraction. Starting from the bounded coordinate found in the first SCC, we extend the prefix across successive SCCs while keeping the same coordinate i below D. The prefix either reaches the target, in which case the entire run can be handled by a 2-VASS, or stops at the entry of the first later SCC $U_{\iota+1}$ where coordinate i becomes large along the witnessing run. In the latter case, we obtain

$$s = \boldsymbol{p}(\boldsymbol{x}) \xrightarrow{\pi_0} \boldsymbol{r}(\boldsymbol{z}) \xrightarrow{\pi_1} \boldsymbol{q}(\boldsymbol{y}) = t. \tag{1}$$

Since coordinate i remains below D throughout π_0, the entire prefix is captured by a single 2-VASS. Moreover, the stopping condition ensures that every $\boldsymbol{w} \in \mathbb{N}^3$ satisfying

$$\boldsymbol{w}(i) \geq \boldsymbol{z}(i), \qquad \boldsymbol{w}(j) \geq D \quad \text{for all } j \neq i \tag{2}$$

gives a forward C-coverable configuration $\boldsymbol{r}(\boldsymbol{w})$ in the next SCC $U_{\iota+1}$. This suggests a way to reach the transformation area. If we can modify the bounded prefix π_0 to raise the other two coordinates of its endpoint above D, while keeping coordinate i at least $\boldsymbol{z}(i)$, then the next SCC admits a run into $\mathcal{G}_C$. The challenge is to achieve this modification while preserving a connection to the target t.

Additive-Growth Induction. We carry out induction on the number k of SCCs using the partition in (1). Fix a threshold $R = D^c \geq D$, where the constant c will be chosen below. We call the prefix π_0 *R-small* if $\|\boldsymbol{z}\| \leq \|\boldsymbol{x}\| + R$, and *$R$-large* otherwise. Let $\boldsymbol{u} = (\boldsymbol{r}, \boldsymbol{z})$ be the configuration of V corresponding to the endpoint of the prefix. Each step has one of the following outcomes:

Terminal Case. We obtain a polynomial-length run to the target directly. This is already possible when U is non-wide, the source is forward C-coverable, or the bounded prefix reaches the target. We will show that every R-large prefix also leads to a terminal case.

Inductive Case. The prefix π_0 is R-small, and hence $\|\boldsymbol{u}\| \leq \|\boldsymbol{s}\| + R$. The new instance consists of the suffix $(U_{\iota+1})u_{\iota+1} \cdots u_{k-1}(U_k)$, with source $\boldsymbol{u}$ and target $\boldsymbol{t}$.

Fix $V, B, C, \boldsymbol{t}$, put $M := D + \|\boldsymbol{t}\|$, and let $L_k(N)$ bound the reachability distance for reachable sequential subinstances with at most k SCCs, source norm at most N, and backward C-coverable target $\boldsymbol{t}$. A single nondecreasing polynomial P bounds both the terminal cases and the cost of replacing the bounded prefix, giving

$$L_1(N) \leq P(M + N), \qquad L_{k+1}(N) \leq P(M + N + R) + L_k(N + R).$$

Thus $L_k(N) \leq kP(M + N + kR)$. Since $k \leq (B + 1)^2$ and $R = D^c$, where c and the degree of P are independent of k, this yields the uniform polynomial bound as required in Theorem 1.2. The full induction is given in Section 7. It remains to show that R can be chosen appropriately so that every R-large prefix falls into a terminal case.

2.3.3 Translated Families and R-Large Dichotomy

Suppose that the third coordinate remains below D along π_0, while $\boldsymbol{z}(2) > \|\boldsymbol{x}\| + R$. By (2), raising the first coordinate of the midpoint at state $\boldsymbol{r}$ to at least D, while keeping the second coordinate sufficiently large, makes it forward C-coverable. Our goal is to find such a configuration that is reachable from s and can still reach t.

Since the third coordinate remains bounded, we can encode it into states and represent π_0 by an LPS of polynomial size, in which every cycle has zero effect on the third coordinate. For sufficiently large R, the growth of the second coordinate forces some cycle β with $\textsc{eff}(\beta)(3) = 0$ and $\textsc{eff}(\beta)(2) > 0$ to be repeated many times. If $\textsc{eff}(\beta)(1) < 0$, deleting some of its repetitions raises the first coordinate at the expense of the second, as illustrated by the following example.

Example 2.1. Suppose a prefix $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{\pi_0} \boldsymbol{r}(\boldsymbol{z})$ is R-large, with $\boldsymbol{z}(3) < D$ and $\boldsymbol{z}(2) > \|\boldsymbol{x}\| + R$. Assume that π_0 contains κ repetitions of a cycle β with $\textsc{eff}(\beta) = (-2, 3, 0)$. Deleting r copies of β changes the endpoint to

$$\boldsymbol{z}^{(r)} = \boldsymbol{z} + r \cdot (2, -3, 0).$$

Both unbounded coordinates of $\boldsymbol{z}^{(r)}$ are at least D whenever

$$\max\left\{0, \left\lceil \frac{D - \boldsymbol{z}(1)}{2} \right\rceil\right\} \le r \le \min\left\{\kappa, \left\lfloor \frac{\boldsymbol{z}(2) - D}{3} \right\rfloor\right\}.$$

Taking κ sufficiently large relative to D, and then R sufficiently large relative to κ, gives a long interval of candidate endpoints. By (2), all these endpoints are forward C-coverable.

The candidate endpoints in this example differ by translations along the same direction. We extend this idea to hybrid sets by considering a family

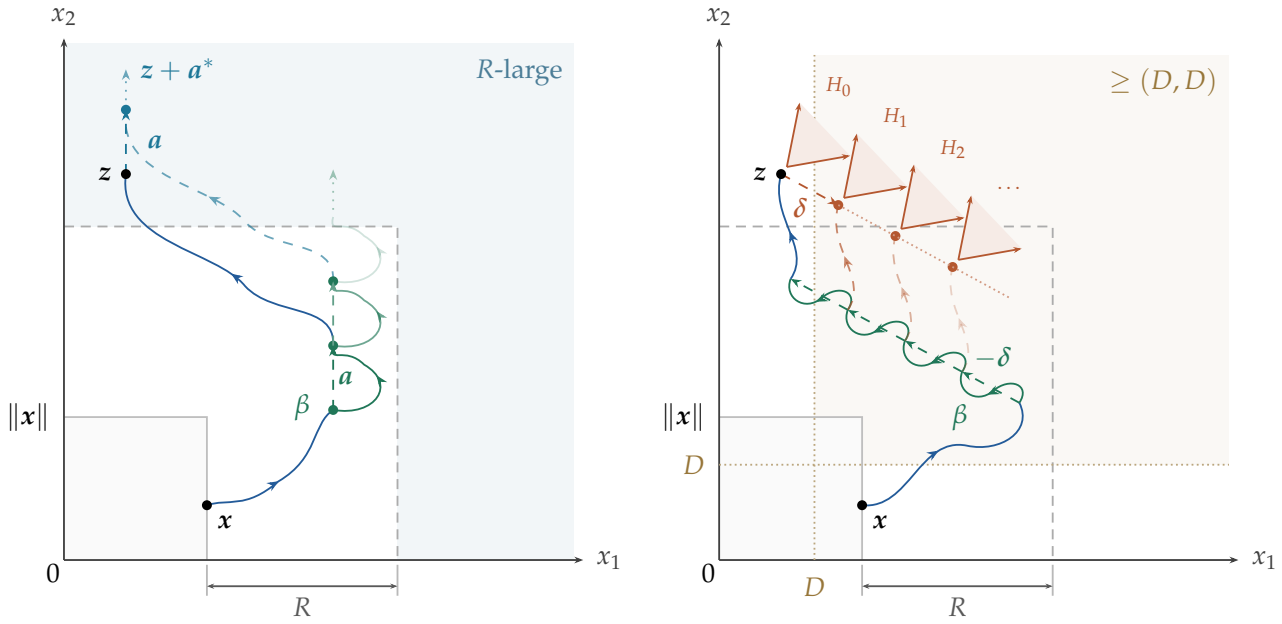
$$H_r = \boldsymbol{b} + r \cdot \boldsymbol{\delta} + P^{\langle \boldsymbol{c}, \cdot \rangle \le K},$$

where the periods P, their costs $\boldsymbol{c}$, and K remain unchanged, while the base moves along $\boldsymbol{\delta} \in \mathbb{Z}^3$. For any κ such that $\boldsymbol{b} + \kappa \cdot \boldsymbol{\delta} \ge \boldsymbol{0}$, we call $\{H_r\}_{r=0}^{\kappa}$ a *translated family*.

However, Example 2.1 relies on two properties that an R-large prefix need not satisfy. First, a frequently repeated cycle β increasing the second coordinate need not decrease the first. Second, even when $\textsc{eff}(\beta) = (-u, v, 0)$ with $u, v > 0$, deleting repetitions may violate nonnegativity along the remaining prefix. We address these difficulties through the following dichotomy.

The R-Large Dichotomy. If some cycle used in the LPS has a nonzero nonnegative effect, we can instead obtain a hybrid set with a nonzero free period. Otherwise, every used cycle β increasing the second coordinate must decrease the first, i.e., $\textsc{eff}(\beta) = (-u, v, 0)$ with $u, v > 0$. We choose the last such cycle repeated at least κ times. Since all later cycles increasing the second coordinate occur fewer than κ times, sufficiently large growth guarantees enough reserve to delete up to κ repetitions while preserving nonnegativity. Using hybrid-set representations, we obtain a translated family of reachable midpoints. These two outcomes are illustrated in Figure 1.

Consequently, for any $\kappa \ge 1$, we can choose $R \le (D + \kappa)^{\mathcal{O}(1)}$ such that the prefix yields either a hybrid set H with a nonzero free period and $\boldsymbol{z} \in H$, or a translated family $\{H_r\}_{r=0}^{\kappa}$ with $\boldsymbol{z} \in H_0 = \boldsymbol{b} + P^{\langle \boldsymbol{c}, \cdot \rangle \le K}$. Crucially, κ can be increased arbitrarily without increasing the polynomial



(a) Free-period outcome.

(b) Translated-family outcome.

Figure 1: The R-large dichotomy, projected onto coordinates 1, 2. (a) Repeating a cycle β with nonzero nonnegative effect $\boldsymbol{a}$ produces a free period. (b) Deleting copies of a mixed-sign cycle β shifts the reachable endpoints along $\boldsymbol{\delta} = -\operatorname{EFF}(\beta)$, producing a translated family. The shapes representing $H_0, H_1, \ldots$ illustrate translations of a common hybrid set.

bounds on the periods, costs, or $\|\boldsymbol{\delta}\|$. Moreover, every member $\boldsymbol{h}$ of the resulting hybrid set or translated family is reachable from s by a short run. Thus, it suffices to find a member of polynomial norm that can reach t by a polynomial-length run. We handle the two outcomes separately.

Terminal Case: Translated Family. Suppose we have obtained such a translated family $\{H_r\}_{r=0}^{\kappa}$ with $\boldsymbol{z} \in H_0$. For sufficiently large κ, there is a sufficiently long interval of indices r for which the base $\boldsymbol{b} + r \cdot \boldsymbol{\delta}$ is forward C-coverable as a midpoint at state $\boldsymbol{r}$. The remaining difficulty is to connect one of these bases to the backward C-coverable target t. Crucially, the wide and C-coverable case requires only $\mathbb{Z}$-reachability between endpoints, which allows us to cancel the displacement caused by translation along $\boldsymbol{\delta}$ through other cycles.

Let $U[P]$ denote the VASS obtained from U by adding the periods in P as self-loops at state $\boldsymbol{r}$. Since $\boldsymbol{z} \in H_0$ reaches t in U, we have $\boldsymbol{r}(\boldsymbol{b}) \xrightarrow{*} t$ in $U[P]$. If $U[P]$ is wide, we can choose three linearly independent cycle effects $\boldsymbol{d}_1, \boldsymbol{d}_2, \boldsymbol{d}_3$ such that $-\boldsymbol{\delta} = \sum_{i \in [3]} \lambda_i \cdot \boldsymbol{d}_i$ for some nonnegative rational coefficients $\lambda_1, \lambda_2, \lambda_3$. Let $\nu := |\det(\boldsymbol{d}_1, \boldsymbol{d}_2, \boldsymbol{d}_3)|$. By Cramer's rule, $\nu\lambda_i$ is a nonnegative integer for every i, so $-\nu\boldsymbol{\delta}$ is a nonnegative integer combination of cycle effects. Thus, a translation whose index is divisible by ν can be compensated by inserting these cycles, preserving $\mathbb{Z}$-reachability to the target. Since ν is polynomially bounded, a sufficiently long interval of forward C-coverable translations contains such an index.

The non-wide case follows from the non-wide decomposition as in Section 2.3.1. Consequently, for $\kappa \geq \kappa_0 = (D + \|P\|)^{\mathcal{O}(1)}$, some $r \in [0, \kappa]$ and $\boldsymbol{h} \in H_r$ satisfy

$$\|\boldsymbol{h}\|, \ell_V\big((\boldsymbol{r}, \boldsymbol{h}), \boldsymbol{t}\big) \leq (D + \|\boldsymbol{t}\| + \operatorname{NORM}(H_0) + \|\boldsymbol{\delta}\|)^{\mathcal{O}(1)}.$$

Terminal Case: Pumpable Sources. A configuration $\boldsymbol{p}(\boldsymbol{x})$ is *pumpable* if it admits a cycle reaching $\boldsymbol{p}(\boldsymbol{x}')$ for some $\boldsymbol{x}' \geq \boldsymbol{x}$ with $\boldsymbol{x}' \neq \boldsymbol{x}$. In the other case where $\boldsymbol{z} \in H = \boldsymbol{b} + P^{\langle c, \cdot \rangle \leq K}$ and H has a free period, it can be deduced that $\boldsymbol{r}(\boldsymbol{b})$ is pumpable in $U[P]$. We show in Section 7.1 that any instance where the source is pumpable and the target is backward C-coverable is a terminal case.

We focus on the coordinates not increased by the initial pumping cycle. When the source is not forward C-coverable, bounded-coordinate extraction yields a prefix in which these coordinates remain bounded in alternating segments. If one coordinate remains bounded throughout the prefix, we can apply 2-VASS techniques. Otherwise, a hill-cutting argument yields two ordered cycles whose effects, together with the initial pumping cycle, can be combined into a strictly positive vector. An auxiliary VASS preserves the required ordering and enabling conditions, making this strictly positive vector a free period in a hybrid representation of the reachable configurations, which can then be joined to the backward C-coverable target using the coverable-case argument.

Uniform Choice of κ and R. The preceding arguments give

$$\|P\| \leq S \leq D^{\mathcal{O}(1)}, \qquad \kappa_0 \leq (D + \|P\|)^{\mathcal{O}(1)}, \qquad R \leq (D + \kappa)^{\mathcal{O}(1)}.$$

Here S bounds the period norms of all translated families extracted by the R-large dichotomy. Crucially, this bound is independent of κ, R, and the current source. We can therefore choose the parameters in the order

$$D \xrightarrow{\text{Lemma 7.14}} S \xrightarrow{\text{Lemma 7.12}} \kappa \xrightarrow{\text{Lemma 7.14}} R,$$

where each arrow represents a polynomial dependence on the previously chosen parameters. Recall also that B, C, and D are polynomially bounded in the original instance size $\text{SIZE}(V, \boldsymbol{s}, \boldsymbol{t})$. All these parameters remain polynomial in $\text{SIZE}(V, \boldsymbol{s}, \boldsymbol{t})$ and can be fixed *before the induction begins*.

2.3.4 General Case

We now establish the polynomial reachability-distance bound without any coverability assumption. Our approach is to extend the additive-growth induction developed for semi-C-coverable instances. The bounded-prefix extraction and the R-large dichotomy remain valid for arbitrary targets. Moreover, combining the pumpability argument with the semi-C-coverable result allows us to handle pumpable sources and hybrid sets with a nonzero free period without target coverability. Thus, the only remaining difficulty is an R-large prefix yielding a translated family $\mathcal{J}$.

Our construction ensures that some translation of $\mathcal{J}$ is forward C-coverable. By reversing the remaining suffix, we can therefore regard $\mathcal{J}$ as a fixed target family with a backward C-coverable translation. We then repeat the additive-growth induction from an arbitrary source, keeping $\mathcal{J}$ fixed. The main difficulty arises when an R-large prefix produces another translated family $\mathcal{H}$. Our construction guarantees that the two families are C-*joinable*: some members of H_0 and J_0 are connected by the original run, and both families admit C-coverable translations at indices at most D, forward for $\mathcal{H}$ and backward for $\mathcal{J}$. We show that C-joinability suffices to obtain a polynomial-length run between suitable members of the two families.

The same additive recurrence thus yields a polynomial reachability-distance bound for Case II of the CFZ characterization, completing the proof of Theorem 1.2.

3 Preliminaries

Let $\mathbb{N}, \mathbb{Z}, \mathbb{Q}$ denote the sets of natural numbers, integers, and rational numbers, respectively. We use, for example, $\mathbb{Q}_{\geq 0}$ for the set of rationals restricted to nonnegative values. For integers $m \leq n$, write $[m, n] = \{m, m+1, \ldots, n\}$. The interval $[1, n]$ is abbreviated as $[n]$. Vectors are written in boldface $\boldsymbol{u}, \boldsymbol{v}, \boldsymbol{x}, \boldsymbol{y}$, etc. For a vector $\boldsymbol{x} \in \mathbb{Q}^d$, we write $\boldsymbol{x}(i)$ for its i-th component, where $i \in [d]$. Its *support* is $\mathrm{supp}(\boldsymbol{x}) := \{i \in [d] : \boldsymbol{x}(i) \neq 0\}$. We compare vectors component-wise, so that $\boldsymbol{x} \leq \boldsymbol{y}$ means $\boldsymbol{x}(i) \leq \boldsymbol{y}(i)$ for all $i \in [d]$. We use $\|\boldsymbol{x}\| := \max_{i \in [d]} |\boldsymbol{x}(i)|$ for the max norm of $\boldsymbol{x}$, and $\|\boldsymbol{x}\|_1 := \sum_{i \in [d]} |\boldsymbol{x}(i)|$ for the L_1-norm of $\boldsymbol{x}$. For a finite set $X \subseteq \mathbb{Q}^d$, we write $\|X\| := \max_{\boldsymbol{x} \in X} \|\boldsymbol{x}\|$ with the understanding that $\max \emptyset := 0$. When the dimension d is clear from the context, we write $\mathbf{1} := (1, 1, \ldots, 1)$ and $\mathbf{0} := (0, 0, \ldots, 0)$. For a finite set $X \subseteq \mathbb{Q}^d$ of vectors where $X = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n\}$, the *cone* generated by X is $\mathrm{cone}(X) := \{\lambda_1 \cdot \boldsymbol{x}_1 + \cdots + \lambda_n \cdot \boldsymbol{x}_n : \lambda_1, \ldots, \lambda_n \in \mathbb{Q}_{\geq 0}\}$.

Linear Set and Hybrid Set. For a finite set $P = \{\boldsymbol{p}_1, \ldots, \boldsymbol{p}_m\} \subseteq \mathbb{N}^d \setminus \{\mathbf{0}\}$, write $P^* := \{\sum_{i=1}^m \lambda_i \cdot \boldsymbol{p}_i : \boldsymbol{\lambda} \in \mathbb{N}^m\}$, with $\emptyset^* = \{\mathbf{0}\}$. A *linear set* is a set $\boldsymbol{b} + P^*$, where $\boldsymbol{b} \in \mathbb{N}^d$ is its *base* and the vectors in P are its *periods*. A *semilinear set* is a finite union of linear sets. For a single period $\boldsymbol{a}$, write $\boldsymbol{a}^* := \{\boldsymbol{a}\}^*$.

Given $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{Q}^d$, let $\langle \boldsymbol{x}, \boldsymbol{y} \rangle$ denote the inner product of these two vectors. We use the notation for hybrid sets from [CFF+26, Section 4]. Given $\boldsymbol{c} \in \mathbb{N}^m$ and $K \in \mathbb{N}$, the associated *hybrid set* is

$$H = \boldsymbol{b} + P^{\langle \boldsymbol{c}, \cdot \rangle \leq K}, \text{ where } P^{\langle \boldsymbol{c}, \cdot \rangle \leq K} = \left\{ \sum_{i=1}^m \lambda_i \cdot \boldsymbol{p}_i : \boldsymbol{\lambda} \in \mathbb{N}^m,\ \langle \boldsymbol{c}, \boldsymbol{\lambda} \rangle \leq K \right\}.$$

We call $\boldsymbol{c}(i)$ the *cost* of $\boldsymbol{p}_i$. A period is *free* if its cost is zero and *restricted* if its cost is positive. The *representation norm* is $\mathrm{NORM}(H) := \max\{\|\boldsymbol{b}\|, \|P\|, \|\boldsymbol{c}\|\}$, where K is excluded from the norm. The case $\boldsymbol{c} = \mathbf{0}$ gives the ordinary linear set.

VASS and VAS. A d-dimensional *vector addition system with states* (d-VASS) is a pair $V = (Q, T)$, where Q is a finite set of *states* and $T \subseteq Q \times \mathbb{Z}^d \times Q$ is a finite set of *transitions*. A *configuration* is a pair $p(\boldsymbol{x})$ with $p \in Q$ and $\boldsymbol{x} \in \mathbb{N}^d$. A transition $u = (p, \boldsymbol{a}, q)$ induces a one-step run $p(\boldsymbol{x}) \xrightarrow{u} q(\boldsymbol{y})$ if $\boldsymbol{x} + \boldsymbol{a} = \boldsymbol{y}$ and both vectors are nonnegative. A *path* is a transition sequence $\pi = (p_1, \boldsymbol{a}_1, q_1) \cdots (p_k, \boldsymbol{a}_k, q_k)$ with $q_i = p_{i+1}$ for $i < k$. It is a *cycle* if it is nonempty and $q_k = p_1$, and a *simple cycle* if it repeats no state except its initial and final state. A one-transition cycle is a *self-loop*. Empty paths are allowed at every state and have length zero and effect $\mathbf{0}$. The effect and length are defined by $\mathrm{EFF}(\pi) := \sum_{i=1}^k \boldsymbol{a}_i$ and $|\pi| := k$.

Fix an enumeration $T = \{u_1, \ldots, u_m\}$, where $u_i = (p_i, \boldsymbol{a}_i, q_i)$. The *transition matrix* of V is

$$\boldsymbol{T} := \begin{pmatrix} \boldsymbol{a}_1 & \boldsymbol{a}_2 & \cdots & \boldsymbol{a}_m \end{pmatrix} \in \mathbb{Z}^{d \times m}.$$

The *Parikh image* of a path $\pi \in T^*$ is the vector $\boldsymbol{\psi}^\pi \in \mathbb{N}^m$ whose i-th component counts the number of occurrences of u_i in π. Hence, $\mathrm{EFF}(\pi) = \boldsymbol{T}\boldsymbol{\psi}^\pi$. We also write $\boldsymbol{\psi}^\pi(u)$ for the number of occurrences of a transition $u \in T$ in π when the chosen enumeration is irrelevant.

A path induces an $\mathbb{N}$-*run* (or simply a run) when all intermediate counter values remain nonnegative. For example,

$$\pi : p_1(\boldsymbol{x}_0) \xrightarrow{t_1} q_1(\boldsymbol{x}_1) \xrightarrow{t_2} \cdots \xrightarrow{t_k} q_k(\boldsymbol{x}_k),$$

where $\boldsymbol{x}_0, \ldots, \boldsymbol{x}_k \in \mathbb{N}^d$. We write $p(\boldsymbol{x}) \xrightarrow{\pi} q(\boldsymbol{y})$ if π induces a run from $p(\boldsymbol{x})$ to $q(\boldsymbol{y})$, and write $p(\boldsymbol{x}) \xrightarrow{*} q(\boldsymbol{y})$ if such a path exists. Two paths π_1 and π_2 may be concatenated as $\pi_1\pi_2$ whenever the target state of π_1 is the source state of π_2. If π is a cycle, then π^m denotes its m-fold concatenation.

A looser semantics is defined on $\mathbb{Z}$-*configurations* of the form $p(\boldsymbol{z})$, where $\boldsymbol{z} \in \mathbb{Z}^d$. A $\mathbb{Z}$-*run* or *integer run* $p(\boldsymbol{z}) \xrightarrow{\pi}_{\mathbb{Z}} q(\boldsymbol{w})$ is defined analogously to a run, with configurations replaced by $\mathbb{Z}$-configurations. We also write $p(\boldsymbol{z}) \xrightarrow{*}_{\mathbb{Z}} q(\boldsymbol{w})$ for $\mathbb{Z}$-runs. For a $\mathbb{Z}$-configuration $p(\boldsymbol{z})$, a path $\pi = t_1 \cdots t_k$ from p to q induces a $\mathbb{Z}$-run $p(\boldsymbol{z}) \xrightarrow{\pi}_{\mathbb{Z}} q(\boldsymbol{w})$, where $\boldsymbol{w} = \boldsymbol{z} + \textsc{eff}(\pi)$. We define the *drop* of π as the vector $\textsc{drop}(\pi) \in \mathbb{N}^d$ given, for every $i \in [d]$, by

$$\textsc{drop}(\pi)(i) := -\min_{j \in [0,k]} \textsc{eff}(t_1 \cdots t_j)(i).$$

Thus, $\textsc{drop}(\pi)(i)$ is the maximum amount by which the i-th counter falls below its initial value along π. Consequently, this $\mathbb{Z}$-run is a run if and only if $\boldsymbol{z} \geq \textsc{drop}(\pi)$. The *reverse* $V^{\textsc{rev}}$ replaces every transition $(p, \boldsymbol{a}, q)$ by $(q, -\boldsymbol{a}, p)$. Clearly, $s \xrightarrow{*} t$ in V if and only if $t \xrightarrow{*} s$ in $V^{\textsc{rev}}$. For a VAS, we also write $-V$ for the reversed system.

A d-dimensional *vector addition system* (d-VAS) is a d-VASS with a single state. In this case, transitions are determined by their effects, so we identify T with a finite subset of $\mathbb{Z}^d$ and omit the unique state from configurations and runs.

Size and Reachability Problems. Let V be a VASS and let s, t be two configurations of V. We call the triple (V, s, t) a *reachability instance*, with s and t referred to as its source and target configurations. The unary-encoded *size* of V is defined by

$$\textsc{size}(V) := |Q| + d \cdot |T| \cdot (\|T\| + 1) + 2,$$

where $\|T\| := \max_{(p,\boldsymbol{a},q) \in T} \|\boldsymbol{a}\|$. Note that $\textsc{size}(V) \geq 2$, so constant factors can be absorbed by increasing a polynomial exponent. The unary size of a reachability instance (V, s, t) is defined by

$$\textsc{size}(V, s, t) := \textsc{size}(V) + d \cdot \big(\|s\| + \|t\| + 1\big).$$

The norm of a configuration $p(\boldsymbol{x})$, and likewise of a $\mathbb{Z}$-configuration, is defined by $\|p(\boldsymbol{x})\| := \|\boldsymbol{x}\|$. Sometimes we also consider a pair (V, s) where the target is not fixed; its size $\textsc{size}(V, s)$ is defined similarly. For brevity, we also refer to such triples and pairs as VASS. For a VASS V, its *reachability relation* is

$$\mathrm{Reach}_V := \Big\{ \big(p(\boldsymbol{x}), q(\boldsymbol{y})\big) : p(\boldsymbol{x}) \xrightarrow{*} q(\boldsymbol{y}) \text{ in } V \Big\}.$$

We compose two relations from left to right: $(s, t) \in R_1 \circ R_2$ if there exists u such that $(s, u) \in R_1$ and $(u, t) \in R_2$. For a fixed source configuration s of V, its *reachability set* is $\mathrm{Reach}_V(s) := \Big\{ t : s \xrightarrow{*} t \text{ in } V \Big\}$. For a VAS, where states are omitted, we regard Reach_V as a subset of $\mathbb{N}^d \times \mathbb{N}^d$ and $\mathrm{Reach}_V(\boldsymbol{s})$ as a subset of $\mathbb{N}^d$. The reachability problem in VASS can be formulated as follows:

Reachability in VASS
Input: A VASS (V, s, t).
Question: Does $s \xrightarrow{*} t$ hold in V?

Complexity bounds for reachability problems are usually obtained from length bounds on runs witnessing the reachability. We write $\textsc{len}(V,s,t) := \left\{|\pi| : s \xrightarrow{\pi} t \text{ in } V\right\}$ for the set of lengths of runs from s to t. The *reachability distance* from s to t in V is defined by $\ell_V(s,t) := \min \textsc{len}(V,s,t)$, with the convention that $\ell_V(s,t) := \infty$ if t is not reachable from s. We use $\mathcal{O}(n)$ and $\mathsf{poly}(n)$ to denote unspecified linear and polynomial functions on integers, respectively.

Mixed Linear Path Schemes. A *linear path scheme* (LPS) is a regular expression

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

where each *path* α_i is a finite sequence of vectors in $\mathbb{Z}^d$, and each *cycle* β_i is a nonempty such sequence. Define $|\Lambda| := |\alpha_0\beta_1\alpha_1 \cdots \beta_k\alpha_k|$ and $\|\Lambda\| := \max\{\|\boldsymbol{a}\| : \boldsymbol{a} \text{ appears in } \Lambda\}$. A sequence captured by Λ has the form $\alpha_0\beta_1^{n_1}\alpha_1 \cdots \beta_k^{n_k}\alpha_k$ for some $n_1, \ldots, n_k \in \mathbb{N}$. For $\mathbb{X} \in \{\mathbb{N}, \mathbb{Z}\}$, write $\boldsymbol{s} \xrightarrow{\Lambda}_{\mathbb{X}} \boldsymbol{t}$ if Λ captures a sequence of vectors whose sum is $\boldsymbol{t} - \boldsymbol{s}$ and whose prefix sums, when added to $\boldsymbol{s}$, all lie in $\mathbb{X}^d$. The subscript $\mathbb{N}$ is often omitted.

A *mixed linear path scheme* (mixed LPS) is an expression

$$\Lambda = (\Lambda_1, \mathbb{X}_1) \cdots (\Lambda_m, \mathbb{X}_m),$$

where $m \geq 0$, each Λ_i is an LPS, and $\mathbb{X}_i \in \{\mathbb{N}, \mathbb{Z}\}$. When $m > 1$, we may assume that Λ_i is nonempty for all $i \in [m]$. Define $|\Lambda| := \sum_{i\in[m]} |\Lambda_i|$, and $\|\Lambda\| := \max_{i\in[m]} \|\Lambda_i\|$. We write $\boldsymbol{c}_0 \xrightarrow{\Lambda} \boldsymbol{c}_m$ if there exist nonnegative vectors $\boldsymbol{c}_0, \ldots, \boldsymbol{c}_m$ such that $\boldsymbol{c}_{i-1} \xrightarrow{\Lambda_i}_{\mathbb{X}_i} \boldsymbol{c}_i$ for all i. The counter values inside a $\mathbb{Z}$-component may be negative.

A witness w specifies the repetition counts of all cycles and the intermediate vectors between consecutive components. These choices must satisfy the semantics of each component. The length $|w|$ is the total number of vectors in the resulting sequences. Write $\mathrm{Reach}_\Lambda := \{(\boldsymbol{s}, \boldsymbol{t}) : \boldsymbol{s} \xrightarrow{\Lambda} \boldsymbol{t}\}$ and let $\ell_\Lambda(\boldsymbol{s}, \boldsymbol{t})$ be the minimum witness length, with value ∞ when there is no witness.

Cone and Cycle-Dimension. Let V be a d-VASS. Its *cycle space* $\textsc{cyc}(V) \subseteq \mathbb{Q}^d$ is the span of the effects of simple cycles in the VASS V, i.e.,

$$\textsc{cyc}(V) := \mathrm{span}\{\textsc{eff}(\pi) : \pi \text{ is a simple cycle in } V\}.$$

Its *cycle-dimension* $\dim_{\mathrm{cyc}}(V) := \dim(\textsc{cyc}(V))$ is the dimension of its cycle space. Its *component-dimension* $\dim_{\mathrm{com}}(V) := \max\{\dim_{\mathrm{cyc}}(U) : U \text{ is an SCC of } V\}$ is the maximum cycle-dimension of its SCCs. In particular, for a d-VAS V with transition set T, we have $\textsc{cyc}(V) = \mathrm{span}(T)$. We remark that the cycle-dimension is also known as the geometric dimension and has been used as a key parameter in the analysis of reachability and coverability [LS19, FYZ24, CGO+26].

B-Encoding and Sequential VASS. For a vector $\boldsymbol{x} \in \mathbb{Z}^d$ and a subset $I \subseteq [d]$, we denote by $\boldsymbol{x}\restriction_I$ the restriction of $\boldsymbol{x}$ to the coordinates in I. Let $V = (Q, T)$ be a d-VASS, $I \subseteq [d]$, and $B \in \mathbb{N}$. The $(d - |I|)$-VASS $[\![V]\!]_{I,B}$ has state set $Q' := Q \times [0, B]^I$. We write $p(\boldsymbol{u})$ and $(p, \boldsymbol{u})$ interchangeably for the state $(p, \boldsymbol{u}) \in Q'$. For every $(p, \boldsymbol{a}, q) \in T$ and $p(\boldsymbol{u}) \in Q'$, it has the transition

$$\big(p(\boldsymbol{u}), \boldsymbol{a}\restriction_{[d]\setminus I}, q(\boldsymbol{v})\big) \quad \text{precisely when} \quad \boldsymbol{v} = \boldsymbol{u} + \boldsymbol{a}\restriction_I \in [0, B]^I.$$

For $p(\boldsymbol{x})$ with $\boldsymbol{x} \restriction_I \in [0,B]^I$, set $[\![p(\boldsymbol{x})]\!]_{I,B} = (p, \boldsymbol{x} \restriction_I)(\boldsymbol{x} \restriction_{[d]\setminus I})$. In the subsequent discussion, we mainly consider the case where $|I| = 2$, and assume, w.l.o.g., that I indexes the first two coordinates of the VAS V when there is no ambiguity. In such cases, we abbreviate $[\![V]\!]_{I,B}$ as $[\![V]\!]_B$, and $[\![\boldsymbol{x}]\!]_{I,B}$ as $[\![\boldsymbol{x}]\!]_B$. We call $[\![V]\!]_B$ a *B-encoding (of a VAS)* when V is a VAS, and a *B-encoding of a VASS* otherwise.

Let $U_i = (Q_i, T_i)$, for $i \in [k]$, be strongly connected d-VASS with pairwise disjoint state sets, and, for $i \in [k-1]$, let $u_i = (q_i, \boldsymbol{a}_i, p_{i+1})$ with $q_i \in Q_i$ and $p_{i+1} \in Q_{i+1}$ be a *bridge*. The *sequential VASS*

$$U = (U_1)u_1(U_2)\cdots u_{k-1}(U_k)$$

has state set $\bigcup_i Q_i$ and transition set $\bigcup_i T_i \cup \{u_1, \ldots, u_{k-1}\}$. When considering a sequential VASS together with a designated run or $\mathbb{Z}$-run, we retain only the components traversed by that run and index them as $U_1, \ldots, U_k$. Thus its source and target states belong to U_1 and U_k, respectively.

Covering Runs and Rackoff's Extraction. We recall two results for extracting short runs reaching sufficiently large configurations. A configuration $q(\boldsymbol{z})$ is *coverable* from $p(\boldsymbol{x})$ if $p(\boldsymbol{x}) \xrightarrow{*} q(\boldsymbol{z}')$ for some $\boldsymbol{z}' \geq \boldsymbol{z}$. The classical result of Rackoff gives a bound on the length of a covering run.

Lemma 3.1 ([Rac78, Theorem 3.5])**.** *Let $V = (Q, T)$ be a d-VASS. If $q(\boldsymbol{z})$ is coverable from $p(\boldsymbol{x})$, there is a covering run ending at state q of length less than $N^{\mathcal{O}((d+1)!)}$, where $N := \max\{2, |Q|, \|T\|, \|\boldsymbol{z}\|\}$.*

We also use a related extraction result for a given run. For a run π and a threshold $D \in \mathbb{N}$, let $\mathrm{BD}(\pi, D)$ denote the set of coordinates whose values remain below D throughout π. Hence, $\mathrm{BD}(\pi, D) = \emptyset$ means that every coordinate reaches a value at least D somewhere along π, although these occurrences need not be simultaneous. The following lemma shows that, for a sufficiently large threshold D, one can instead find a short run from the same source reaching a configuration in which all coordinates are simultaneously large.

Lemma 3.2 ([CFZ26, Lemma 3.3])**.** *For every d-VASS (V, s) and number $C \in \mathbb{N}$, there exists a threshold*

$$D_0 := (C + \mathrm{SIZE}(V))^{\mathcal{O}((d+1)!)}$$

such that, if every coordinate reaches a value at least D_0 somewhere along a run π from s in V, i.e., $\mathrm{BD}(\pi, D_0) = \emptyset$, there is also a run ρ from s with $|\rho| \leq D_0$ and ending at $q(\boldsymbol{x})$ satisfying $\boldsymbol{x} \geq C \cdot \mathbf{1}$. Moreover, ρ is obtained by deleting cycles from π and therefore has the same final state as π.

Integer Programming. Let $\boldsymbol{A} \in \mathbb{Z}^{m\times n}$ be a matrix, $\boldsymbol{b} \in \mathbb{Z}^m$ be a vector, and $\boldsymbol{x}$ be a vector of variables. A *linear Diophantine system* $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$ is a set of m linear equations with n variables, whose coefficients are integers. A *minimal solution* of the system $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$ is a minimal element, under the component-wise order $\leq$, among the nonnegative integer solutions of the system. Any nonnegative integer solution can be written as $\boldsymbol{x} = \boldsymbol{x}_0 + \sum_j \lambda_j \cdot \boldsymbol{h}_j$, where $\boldsymbol{x}_0$ is minimal for $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$, the λ_j are nonnegative integers, and the $\boldsymbol{h}_j$ range over the finitely many minimal elements of $\{\boldsymbol{h} \in \mathbb{N}^n \setminus \{\mathbf{0}\} : \boldsymbol{A}\boldsymbol{h} = \mathbf{0}\}$. We write $\|\boldsymbol{A}\|_{1,\infty} := \max_{i\in[m]} \|\boldsymbol{a}_i\|_1$ and $\|\boldsymbol{A}\| := \max_{i\in[m]} \|\boldsymbol{a}_i\|$, where $\boldsymbol{a}_i$ is the i-th row of $\boldsymbol{A}$. Pottier's theorem [Pot91] is used to bound the size of minimal solutions:

Theorem 3.3 (Pottier's Theorem [Pot91, Theorem 1])**.** *Let $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$ be a linear Diophantine system, where $\boldsymbol{A} \in \mathbb{Z}^{m\times n}$ and $\boldsymbol{b} \in \mathbb{Z}^m$. Each minimal element $\boldsymbol{x}_0$ in $\{\boldsymbol{x} \in \mathbb{N}^n \setminus \{\mathbf{0}\} : \boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}\}$ satisfies*

$$\|\boldsymbol{x}_0\|_1 \leq \left(1 + \|\boldsymbol{A}\|_{1,\infty} + \|\boldsymbol{b}\|\right)^r, \tag{3}$$

where $r := \operatorname{rank}\begin{pmatrix} \boldsymbol{A} & \boldsymbol{b} \end{pmatrix}$ *is the rank of the augmented matrix* $\begin{pmatrix} \boldsymbol{A} & \boldsymbol{b} \end{pmatrix}$.

The same bound holds trivially for the zero solution, whenever it exists. Sometimes we bound $\|\boldsymbol{A}\|_{1,\infty}$ by $n \cdot \|\boldsymbol{A}\|$ to simplify (3). For a system of linear inequalities of the form $\boldsymbol{A}\boldsymbol{x} \geq \boldsymbol{b}$, we can introduce m slack variables and reduce it to a linear Diophantine system without increasing the number of equations. By Theorem 3.3, the L_1-norm of any minimal solution of the system is bounded by $\left(2 + \|\boldsymbol{A}\|_{1,\infty} + \|\boldsymbol{b}\|\right)^m$. Together with the standard decomposition of nonnegative Diophantine solutions, we obtain the following corollary.

Corollary 3.4. *Let* $\boldsymbol{A} \in \mathbb{Z}^{m\times n}$ *and* $\boldsymbol{b} \in \mathbb{Z}^m$. *The set of nonnegative integer solutions of either* $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$ *or* $\boldsymbol{A}\boldsymbol{x} \geq \boldsymbol{b}$ *is a finite union of linear sets* $\bigcup_{\boldsymbol{x}_0 \in F}(\boldsymbol{x}_0 + H^*)$, *where* $F \subseteq \mathbb{N}^n$ *is finite and* $H \subseteq \mathbb{N}^n \setminus \{\boldsymbol{0}\}$ *is a common period set independent of* $\boldsymbol{b}$. *Moreover,*

$$\|F\| \leq (2 + \|\boldsymbol{A}\|_{1,\infty} + \|\boldsymbol{b}\|)^m, \qquad \|H\| \leq (2 + \|\boldsymbol{A}\|_{1,\infty})^m.$$

4 Generalized Transformation Areas for B-Encodings of VAS

In this section, we establish a generalized transformation-area theorem for B-encodings of VAS. Classically, a transformation area is a region in which every $\mathbb{Z}$-run between two configurations can be rearranged into a run with the same Parikh image [vL74, HP79][2]. We now formally define the generalization to B-encodings outlined in Section 2.

Let V be a $(d+2)$-VAS with transition set T and $[\![V]\!]_B$ its B-encoding. We identify each configuration $\boldsymbol{p}(\boldsymbol{x})$ of $[\![V]\!]_B$ with $(\boldsymbol{p}, \boldsymbol{x}) \in \mathbb{N}^{d+2}$ in V. Throughout this section, Parikh images of paths in $[\![V]\!]_B$ count their underlying transitions in T.

Definition 4.1. A set $A \subseteq [0, B]^2 \times \mathbb{N}^d$ is a *generalized transformation area* for $[\![V]\!]_B$ if for every $\boldsymbol{s}, \boldsymbol{t} \in A$ and every $\mathbb{Z}$-run $\boldsymbol{s} \xrightarrow{\pi}_{\mathbb{Z}} \boldsymbol{t}$ in $[\![V]\!]_B$, there is a run $\boldsymbol{s} \xrightarrow{\rho} \boldsymbol{t}$ in V with $\boldsymbol{\psi}^\rho = \boldsymbol{\psi}^\pi$.

Recall that C is a transformation-area threshold for $[\![V]\!]_B$ if

$$\mathcal{G}_C = [0, B]^2 \times \{\boldsymbol{x} \in \mathbb{N}^d : \boldsymbol{x} \geq C \cdot \boldsymbol{1}\}$$

is a generalized transformation area for $[\![V]\!]_B$. The main result of this section gives an explicit polynomial threshold and, more strongly, guarantees that the new run stays within the $2B$-encoding.

Theorem 4.2 (Generalized Transformation Area Theorem)**.** *Let* V *be a* $(d+2)$*-VAS with transition set* T *and let* $B \in \mathbb{N}$. *Set*

$$C := (d+2)(B+1)^2 \|T\|.$$

Then $\mathcal{G}_C$ *is a generalized transformation area. More precisely, for every* $\boldsymbol{s}, \boldsymbol{t} \in \mathcal{G}_C$ *and every* $\mathbb{Z}$*-run* $\boldsymbol{s} \xrightarrow{\pi}_{\mathbb{Z}} \boldsymbol{t}$ *in* $[\![V]\!]_B$, *there is a run* $\boldsymbol{s} \xrightarrow{\rho} \boldsymbol{t}$ *in* $[\![V]\!]_{2B}$ *with* $\boldsymbol{\psi}^\rho = \boldsymbol{\psi}^\pi$.

The proof adapts the cycle-relocation technique of Hopcroft and Pansiot [HP79, Lemma 3.9]. We first extract and relocate simple cycles to one common state, keeping the two encoded coordinates within $[0, 2B]^2$. We then apply the nonzero-sum form of Steinitz's lemma in Corollary 4.4 to reorder their effects so that the d remaining counters stay nonnegative.

[2] We also give an explicit bound for van Leeuwen's classical transformation areas in Appendix B.

4.1 Steinitz Reordering and Cycle Relocation

We use Steinitz's lemma [Ste13] in the form of Grinberg and Sevast'yanov [GS80], which bounds the partial sums of a reordered zero-sum sequence.

Lemma 4.3 (Steinitz Lemma [GS80]). *Let $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n \in \mathbb{Q}^d$, with $\sum_{\ell=1}^{n} \boldsymbol{v}_\ell = \mathbf{0}$ and $\|\boldsymbol{v}_\ell\| \le M$. There is a permutation σ of $[n]$ such that $\left\|\sum_{\ell=1}^{k} \boldsymbol{v}_{\sigma(\ell)}\right\| \le dM$ for every $k \in [0, n]$.*

The total effect of the extracted cycles need not vanish. We also need a version of Steinitz's lemma for a nonzero total sum, which is a corollary of the above lemma.

Corollary 4.4. *Let $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n \in \mathbb{Q}^d$ have norm at most M, and let $\boldsymbol{z} := \sum_{i=1}^{n} \boldsymbol{v}_i$. There is a permutation σ such that every partial sum $\boldsymbol{p}_k = \sum_{i=1}^{k} \boldsymbol{v}_{\sigma(i)}$ satisfies*

$$\boldsymbol{p}_k \ge \min(\mathbf{0}, \boldsymbol{z}) - dM \cdot \mathbf{1},$$

where the minimum is componentwise.

Proof. The empty case is immediate. Otherwise append n auxiliary copies of $-\boldsymbol{z}/n$. Since $\|\boldsymbol{z}\| \le nM$, all $2n$ vectors have norm at most M and sum to zero. Apply Lemma 4.3, then delete the auxiliary vectors. For any prefix of the resulting order, consider the augmented prefix ending at its last original vector. If it contains m auxiliary vectors, its sum is $\boldsymbol{p}_k - (m/n) \cdot \boldsymbol{z}$, where k is the number of original vectors in the prefix. Since $(m/n) \cdot \boldsymbol{z} \ge \min(\mathbf{0}, \boldsymbol{z})$ and $\|\boldsymbol{p}_k - (m/n) \cdot \boldsymbol{z}\| \le dM$, one has the required bound $\boldsymbol{p}_k \ge \min(\mathbf{0}, \boldsymbol{z}) - dM \cdot \mathbf{1}$. The empty prefix also satisfies it. □

Cycle Relocation. A state $\boldsymbol{r}$ visited by π is a *peak state* of a path π if

$$\boldsymbol{r}(1) = \max\{\boldsymbol{p}(1) : \boldsymbol{p} \text{ appears in } \pi\}.$$

We first show that a simple cycle can be relocated to a peak state while preserving its Parikh image.

Lemma 4.5 (Cycle Relocation). *Let π be a path in $[\![V]\!]_B$. Suppose that β is a simple cycle in π and that $\boldsymbol{r}$ is a peak state of π. Then there exists a simple cycle $\widehat{\beta}$ at $\boldsymbol{r}$ in $[\![V]\!]_{2B}$ with $\boldsymbol{\psi}^{\beta} = \boldsymbol{\psi}^{\widehat{\beta}}$.*

Proof. Rotate β to start at a state $\boldsymbol{w}$ minimizing its second coordinate, and translate its encoded states by $\boldsymbol{r} - \boldsymbol{w}$. Since β has zero total effect on the encoded coordinates, every resulting state is $\boldsymbol{r} + \boldsymbol{v} - \boldsymbol{w}$ for a state $\boldsymbol{v}$ of β. Here $\boldsymbol{r}(1) \ge \boldsymbol{w}(1)$ by the peak-state condition and $\boldsymbol{v}(2) \ge \boldsymbol{w}(2)$ by the choice of $\boldsymbol{w}$. Together with $\boldsymbol{r}, \boldsymbol{v}, \boldsymbol{w} \in [0, B]^2$, this gives $0 \le \boldsymbol{r}(i) + \boldsymbol{v}(i) - \boldsymbol{w}(i) \le 2B$ for $i \in \{1, 2\}$. Thus the same underlying VAS transitions define a cycle $\widehat{\beta}$ at $\boldsymbol{r}$ in $[\![V]\!]_{2B}$. Rotation and translation preserve state distinctness and underlying transition counts, so $\widehat{\beta}$ is simple and $\boldsymbol{\psi}^{\beta} = \boldsymbol{\psi}^{\widehat{\beta}}$. □

The next corollary decomposes a path in $[\![V]\!]_B$ into two short paths and simple cycles that can be relocated to a peak state in $[\![V]\!]_{2B}$. Write $Q_B = [0, B]^2$ for the states of $[\![V]\!]_B$.

Corollary 4.6. *Let π be a path in $[\![V]\!]_B$ from $\boldsymbol{p}$ to $\boldsymbol{q}$. There exists a path $\alpha\gamma_1 \cdots \gamma_n\beta$ from $\boldsymbol{p}$ to $\boldsymbol{q}$ in $[\![V]\!]_{2B}$ such that:*

1. *The prefix α and the suffix β are paths in $[\![V]\!]_B$, with $|\alpha|, |\beta| \le (B+1)^2 - 1$.*
2. *Each γ_i is a simple cycle in $[\![V]\!]_{2B}$ of length at most $(B+1)^2$.*

3. $\boldsymbol{\psi}^{\pi} = \boldsymbol{\psi}^{\alpha} + \sum_{i=1}^{n} \boldsymbol{\psi}^{\gamma_i} + \boldsymbol{\psi}^{\beta}$.

Consequently $\alpha\gamma_{\sigma(1)} \cdots \gamma_{\sigma(n)}\beta$ is a path in $[\![V]\!]_{2B}$ for every permutation σ.

Proof. Split π at an occurrence of a peak state $\boldsymbol{r}$ and obtain two segments of π. Whenever either segment is not simple, select a simple cycle θ and use Lemma 4.5 to obtain a simple cycle $\widehat{\theta}$ at $\boldsymbol{r}$ with $\boldsymbol{\psi}^{\theta} = \boldsymbol{\psi}^{\widehat{\theta}}$. We also have $|\widehat{\theta}| = |\theta| \leq |Q_B| \leq (B+1)^2$. Delete θ from the residual path. The remaining encoded states are unchanged, so the residual path stays in $[\![V]\!]_B$ and still passes through the peak state $\boldsymbol{r}$. Hence, the same lemma applies at every step. Each deletion shortens the path, so the process ends with two simple paths α and β with $|\alpha|, |\beta| \leq (B+1)^2 - 1$.

Let $\gamma_1, \ldots, \gamma_n$ be the constructed cycles. The successive deletions give the stated Parikh identity. Finally, α ends at $\boldsymbol{r}$, all γ_i start and end there, and β starts there, so their concatenation is a path in $[\![V]\!]_{2B}$ with the cycles in any order. □

4.2 Proof of the Transformation-Area Theorem

Combining the previous lemmas, we can now prove Theorem 4.2. Write $I = \{1, 2\}$ for the encoded coordinates and $J = [d+2] \setminus I$ for the remaining coordinates.

Proof of Theorem 4.2. Assume $\boldsymbol{s} = (\boldsymbol{p}, \boldsymbol{x})$ and $\boldsymbol{t} = (\boldsymbol{r}, \boldsymbol{y})$, where $\boldsymbol{p}, \boldsymbol{r} \in Q_B$ and $\boldsymbol{x}, \boldsymbol{y} \geq C \cdot \mathbf{1}$. Put $M := (B+1)^2 \, \|T\|$, so $C = (d+2)M$. Apply Corollary 4.6 to obtain $\alpha, \gamma_1, \ldots, \gamma_n, \beta$ from π. For a permutation σ to be chosen, consider

$$\rho = \alpha\gamma_{\sigma(1)} \cdots \gamma_{\sigma(n)}\beta.$$

This path has the same Parikh image as π. Let $\boldsymbol{u} = \boldsymbol{x} + \textsc{eff}(\alpha) \restriction_J$ and $\boldsymbol{v} = \boldsymbol{y} - \textsc{eff}(\beta) \restriction_J$ be the J-coordinates immediately after α and immediately before β, respectively. Every prefix of α and every suffix of β has effect of norm at most M on J. Thus α and β remain nonnegative on J, and

$$\boldsymbol{u}, \boldsymbol{v} \geq (C - M) \cdot \mathbf{1} = (d+1)M \cdot \mathbf{1}.$$

Let $\boldsymbol{a}_i = \textsc{eff}(\gamma_i) \restriction_J$ for $i \in [n]$. Since $|\gamma_i| \leq (B+1)^2$, we have $\|\boldsymbol{a}_i\| \leq M$, and the Parikh identity gives $\sum_i \boldsymbol{a}_i = \boldsymbol{v} - \boldsymbol{u}$. By Corollary 4.4, there is a permutation σ such that every boundary between the reordered middle blocks has its J-coordinates at least

$$\boldsymbol{u} + \min(\mathbf{0}, \boldsymbol{v} - \boldsymbol{u}) - dM \cdot \mathbf{1} = \min(\boldsymbol{u}, \boldsymbol{v}) - dM \cdot \mathbf{1} \geq M \cdot \mathbf{1}.$$

A prefix of any γ_i decreases each coordinate by at most M. Hence ρ remains nonnegative on J.

It remains to check the two encoded coordinates. By Corollary 4.6, ρ stays in $[0, 2B]^2$ on I when started at $\boldsymbol{p}$. Finally, $\boldsymbol{\psi}^{\rho} = \boldsymbol{\psi}^{\pi}$ gives the target $\boldsymbol{t}$. Thus ρ is the required run in $[\![V]\!]_{2B}$. □

4.3 Stability of Transformation-Area Thresholds

Later arguments may reverse the system or add transitions with zero effect on I and nonnegative or nonpositive effect on J. Corollary 4.7 allows us to use the same threshold under these operations.

Corollary 4.7. *Let V be a $(d+2)$-VAS and $B \in \mathbb{N}$. Let C be any transformation-area threshold for $[\![V]\!]_B$. The same C is a transformation-area threshold for $[\![V^{\mathrm{REV}}]\!]_B$ and for $[\![V']\!]_B$ whenever V' is obtained from V by adjoining transitions in $\{(0,0)\} \times (\mathbb{N}^d \cup (-\mathbb{N}^d))$.*

Proof. For reversal, we reverse a $\mathbb{Z}$-run in $[\![V^{\mathrm{REV}}]\!]_B$ to be a $\mathbb{Z}$-run in $[\![V]\!]_B$, apply the transformation-area property for $[\![V]\!]_B$, and reverse the resulting run. For the second assertion, consider a $\mathbb{Z}$-run π from $\boldsymbol{s}$ to $\boldsymbol{t}$ in $[\![V']\!]_B$ with endpoints in $\mathcal{G}_C$. Partition the added occurrences into nonnegative and nonpositive ones, assigning every occurrence of a zero-effect transition to the former group. Let $\boldsymbol{a}, \boldsymbol{b} \in \{(0,0)\} \times \mathbb{N}^d$ be their respective total effect and negated total effect. The remaining path η is still a path in $[\![V]\!]_B$, which induces a $\mathbb{Z}$-run $\boldsymbol{s} + \boldsymbol{a} \xrightarrow{\eta}_{\mathbb{Z}} \boldsymbol{t} + \boldsymbol{b}$. Since $\boldsymbol{s} + \boldsymbol{a}, \boldsymbol{t} + \boldsymbol{b} \in \mathcal{G}_C$, the transformation-area property gives a run in V between these two configurations with the same Parikh image as η. The deleted nonnegative and nonpositive transitions give a run from $\boldsymbol{s}$ to $\boldsymbol{s} + \boldsymbol{a}$ and a run from $\boldsymbol{t} + \boldsymbol{b}$ to $\boldsymbol{t}$, respectively. Concatenating these three runs gives a run from $\boldsymbol{s}$ to $\boldsymbol{t}$ in V' with the same Parikh image as π. □

5 Polynomial Factorization of VAS Reachability

We now establish the polynomial factorization of VAS reachability stated in Theorem 1.3. The proof uses the following linear factorization of B-encodings, with their reachability relations henceforth viewed on configurations of V by restoring the counter values stored in their states.

Recall that a mixed LPS consists of LPS components interpreted under $\mathbb{N}$- or $\mathbb{Z}$-semantics. For a VASS or a mixed LPS V_i, we write Reach_{V_i} for its reachability relation and $\ell_{V_i}(s,t)$ for its minimum witness length. In particular, for a VASS this is the ordinary reachability distance.

Lemma 5.1. *For every fixed $d \geq 1$, there exists a polynomial P such that, for every $(d+2)$-VAS V, every $I \subseteq [d+2]$ with $|I| = 2$, and every $B \in \mathbb{N}$, there is a finite family $\mathcal{F}$ of sequences $(R_1, \ldots, R_m)$ satisfying*

$$\mathrm{Reach}_{[\![V]\!]_{I,B}} \subseteq \bigcup_{(R_1,\ldots,R_m)\in\mathcal{F}} R_1 \circ \cdots \circ R_m \subseteq \mathrm{Reach}_V .$$

Every sequence in $\mathcal{F}$ is nonempty and has length at most $2B+2$. Each factor is $R_i = \mathrm{Reach}_{V_i} \subseteq \mathrm{Reach}_V$, where V_i has one of the following forms:

1. *$V_i = [\![V]\!]_{I\cup\{j\},D}$ for some $j \in [d+2] \setminus I$ and $D \leq P(B + \mathrm{SIZE}(V))$;*
2. *$V_i = \Lambda$ for some mixed LPS Λ, with $\|\Lambda\|, |\Lambda| \leq P(B + \mathrm{SIZE}(V))$.*

Moreover, for every factor R_i and every $(\boldsymbol{x}, \boldsymbol{y}) \in R_i$,

$$\ell_V(\boldsymbol{x}, \boldsymbol{y}) \leq \ell_{V_i}(\boldsymbol{x}, \boldsymbol{y}).$$

Proof sketch of Theorem 1.3. By adding dummy counters, any encoding of V with fewer than two bounded coordinates can be regarded as one with exactly two bounded coordinates, so Lemma 5.1 applies. Starting from V, we repeatedly apply Lemma 5.1 to B-encoding factors with fewer than three bounded original coordinates, leaving mixed-LPS factors unchanged. Each application replaces such a factor by a finite union of compositions of mixed LPS and B-encoding factors with one additional bounded original coordinate. After three levels, every remaining B-encoding factor has three bounded coordinates, and the procedure terminates.

Since each application produces compositions of polynomial length with polynomially bounded parameters, the fixed recursion depth ensures that the resulting factorization has the required polynomial bounds. Every reachable pair in V is covered by a resulting composition $R_1 \circ \cdots \circ R_m$.

Conversely, for any resulting composition, witnesses for its factors can be realized as runs of V, which concatenate to a run between the same endpoints. □

The full proof is given in Appendix C. The rest of this section proves Lemma 5.1 and presents applications of the polynomial factorization.

5.1 Mixed LPS from Generalized Transformation Areas

To apply Theorem 4.2, we represent $\mathbb{Z}$-reachability in $[\![V]\!]_B$ by mixed LPS whose endpoints are restricted to $\mathcal{G}_C$. We enforce this restriction using a checking gadget under $\mathbb{N}$-semantics.

Endpoint Checks. Let V be a $(d+2)$-VAS with transition set T, and let C be a transformation-area threshold for $[\![V]\!]_B$. For each $\boldsymbol{p} \in [0,B]^2$, define $\boldsymbol{g}_{\boldsymbol{p}}^C := (\boldsymbol{p}, C \cdot \mathbf{1}) \in \mathbb{N}^{d+2}$ and a checking path

$$\chi_{\boldsymbol{p}}^C := (-\boldsymbol{g}_{\boldsymbol{p}}^C)(+\boldsymbol{g}_{\boldsymbol{p}}^C).$$

Under $\mathbb{N}$-semantics, this gadget is enabled from $\boldsymbol{z} \in \mathbb{N}^{d+2}$ exactly when $\boldsymbol{z} \geq \boldsymbol{g}_{\boldsymbol{p}}^C$. Given a $\mathbb{Z}$-LPS over T, we place $\chi_{\boldsymbol{p}}^C$ and $\chi_{\boldsymbol{q}}^C$ at its two ends under $\mathbb{N}$-semantics. The checks ensure that, after subtracting a common nonnegative offset on the first two coordinates, the endpoints of the middle $\mathbb{Z}$-run lie in $\mathcal{G}_C$. Thus Theorem 4.2 applies to every witness of the resulting mixed LPS. This yields the following lemma. In what follows, we only consider mixed LPS over $T \cup \{\pm \boldsymbol{g}_{\boldsymbol{p}}^C : \boldsymbol{p} \in [0,B]^2\}$.

Lemma 5.2. *If $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{*}_{\mathbb{Z}} \boldsymbol{q}(\boldsymbol{y})$ in $[\![V]\!]_B$ with $\boldsymbol{p}(\boldsymbol{x}), \boldsymbol{q}(\boldsymbol{y}) \in \mathcal{G}_C$, then some mixed LPS Λ satisfies*

$$(\boldsymbol{p}, \boldsymbol{x}) \xrightarrow{\Lambda} (\boldsymbol{q}, \boldsymbol{y}), \qquad \mathrm{Reach}_\Lambda \subseteq \mathrm{Reach}_V, \qquad \|\Lambda\|, |\Lambda| \leq \mathsf{poly}(B + C + \mathrm{SIZE}(V)).$$

Every $(\boldsymbol{s}, \boldsymbol{t}) \in \mathrm{Reach}_\Lambda$ satisfies $\ell_V(\boldsymbol{s}, \boldsymbol{t}) \leq \ell_\Lambda(\boldsymbol{s}, \boldsymbol{t})$.

Proof. The $\mathbb{Z}$-reachability relation from $\boldsymbol{p}$ to $\boldsymbol{q}$ in $W = [\![V]\!]_B$ is represented by finitely many LPSs of length at most $\mathsf{poly}(\mathrm{SIZE}(W))$ [BEF+21, Lemma 4.10 and Proposition 4.11]. These LPSs are obtained from paths and cycles in W by recording their transition effects. Choose such an LPS Γ with $\boldsymbol{x} \xrightarrow{\Gamma}_{\mathbb{Z}} \boldsymbol{y}$, retaining the underlying transition of W for each vector occurrence. Replacing each occurrence by its corresponding transition in V gives an LPS Γ' over $\mathbb{Z}^{d+2}$. Set

$$\Lambda := (\chi_{\boldsymbol{p}}^C, \mathbb{N})(\Gamma', \mathbb{Z})(\chi_{\boldsymbol{q}}^C, \mathbb{N}).$$

Since the given endpoints lie in $\mathcal{G}_C$, both checking gadgets are enabled, and hence $(\boldsymbol{p}, \boldsymbol{x}) \xrightarrow{\Lambda} (\boldsymbol{q}, \boldsymbol{y})$.

Consider any witness w of Λ from $\boldsymbol{s}$ to $\boldsymbol{t}$ satisfying $|w| = \ell_\Lambda(\boldsymbol{s}, \boldsymbol{t})$, with middle $\mathbb{Z}$-run ρ. The gadgets ensure that ρ starts at $\boldsymbol{s} = (\boldsymbol{p} + \boldsymbol{\eta}, \boldsymbol{x}')$ and ends at $\boldsymbol{t} = (\boldsymbol{q} + \boldsymbol{\eta}', \boldsymbol{y}')$, where $\boldsymbol{\eta}, \boldsymbol{\eta}' \in \mathbb{N}^2$ and $\boldsymbol{x}', \boldsymbol{y}' \geq C \cdot \mathbf{1}$. Projecting ρ onto the last d coordinates gives a sequence captured by Γ. By construction, this sequence induces a $\mathbb{Z}$-run from $\boldsymbol{p}(\boldsymbol{x}')$ to $\boldsymbol{q}(\boldsymbol{y}')$ in $[\![V]\!]_B$. By Definition 4.1, this run can be rearranged into a run in V from $(\boldsymbol{p}, \boldsymbol{x}')$ to $(\boldsymbol{q}, \boldsymbol{y}')$. Monotonicity then gives a run ρ' from $(\boldsymbol{p} + \boldsymbol{\eta}, \boldsymbol{x}')$ to $(\boldsymbol{q} + \boldsymbol{\eta}, \boldsymbol{y}')$. It remains to show that $\boldsymbol{\eta} = \boldsymbol{\eta}'$. All the preceding operations preserve the Parikh image over V, so $\boldsymbol{\psi}^{\rho'} = \boldsymbol{\psi}^{\rho}$ and

$$(\boldsymbol{q} + \boldsymbol{\eta}') - (\boldsymbol{p} + \boldsymbol{\eta}) = \mathrm{EFF}(\rho) \restriction_{\{1,2\}} = \mathrm{EFF}(\rho') \restriction_{\{1,2\}} = \boldsymbol{q} - \boldsymbol{p},$$

which implies $\boldsymbol{\eta}' = \boldsymbol{\eta}$. We also have $|\rho'| = |\rho| \leq |w|$. This proves that $(\boldsymbol{s}, \boldsymbol{t}) \in \mathrm{Reach}_V$ and $\ell_V(\boldsymbol{s}, \boldsymbol{t}) \leq |\rho'| \leq |w| = \ell_\Lambda(\boldsymbol{s}, \boldsymbol{t})$. Since $|\Lambda| \leq |\Gamma| + 4$ and $\|\Lambda\| \leq \max\{\|\Gamma\|, B, C\}$, the parameters of Λ are polynomial in $B + C + \mathrm{SIZE}(V)$. □

5.2 Linear Factorization of B-Encodings of VAS

Assume $I = \{1, 2\}$ and put $W := [\![V]\!]_B$. Set $C := (d+2)(2B+1)^2 \|T\|$, where T is the transition set of V. All LPS below are over $T \cup \{\pm \boldsymbol{g}_{\boldsymbol{p}}^C : \boldsymbol{p} \in [0, B]^2\}$ as in Section 5.1.

By Lemma 3.2, we can choose $H \geq C$, polynomial in $B + \text{SIZE}(V)$, with the following property for both W and its reverse. If every counter reaches at least H somewhere along a run, then deleting cycles yields a run of length at most H that ends in the same state with all counters at least C.

5.2.1 Normalization at High Peak Positions

A *position* of a run π is an occurrence of a configuration along it. A position is *high* if every counter reaches at least H both in the prefix ending there and in the suffix starting there. If π visits a peak state at a high position, Rackoff's extraction gives short runs covering $C \cdot \mathbf{1}$ at that state, yielding configurations in $\mathcal{G}_C$. Cycle relocation then collects the removed cycles into a middle $\mathbb{Z}$-run.

Lemma 5.3. *Let $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{\pi} \boldsymbol{q}(\boldsymbol{y})$ be a run in W that visits a peak state at a high position. Then some mixed LPS Λ satisfies $(\boldsymbol{p}, \boldsymbol{x}) \xrightarrow{\Lambda} (\boldsymbol{q}, \boldsymbol{y})$, $\text{Reach}_\Lambda \subseteq \text{Reach}_V$, and $\|\Lambda\|, |\Lambda| \leq \text{poly}(B + \text{SIZE}(V))$. Every $(\boldsymbol{u}, \boldsymbol{v}) \in \text{Reach}_\Lambda$ satisfies $\ell_V(\boldsymbol{u}, \boldsymbol{v}) \leq \ell_\Lambda(\boldsymbol{u}, \boldsymbol{v})$.*

Proof. Choose such a high position, let $\boldsymbol{p}^*$ be its state, and split $\pi = \pi^- \pi^+$ there. By the definition of H and applying Lemma 3.2 to π^- and to the reverse of π^+, we obtain

$$\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{\alpha} \boldsymbol{p}^*(\boldsymbol{z}), \qquad \boldsymbol{p}^*(\boldsymbol{u}) \xrightarrow{\delta} \boldsymbol{q}(\boldsymbol{y}), \qquad |\alpha|, |\delta| \leq H, \qquad \boldsymbol{z}, \boldsymbol{u} \geq C \cdot \mathbf{1}.$$

Both runs are obtained by deleting cycles from the corresponding segments. Decompose the deleted cycles into simple cycles. These cycles have total effect $\boldsymbol{u} - \boldsymbol{z}$. Since $\boldsymbol{p}^*$ is a peak state of π, the construction in Lemma 4.5 rotates and relocates them to $\boldsymbol{p}^*$ in $[\![V]\!]_{2B}$. Concatenating them gives a $\mathbb{Z}$-run $\boldsymbol{p}^*(\boldsymbol{z}) \xrightarrow{\tau}_{\mathbb{Z}} \boldsymbol{p}^*(\boldsymbol{u})$.

Apply Lemma 5.2 with encoding bound $2B$ and threshold C to obtain a mixed LPS Θ, and set

$$\Lambda := (\alpha, \mathbb{N})\Theta(\delta, \mathbb{N}).$$

Then $(\boldsymbol{p}, \boldsymbol{x}) \xrightarrow{\Lambda} (\boldsymbol{q}, \boldsymbol{y})$ holds. For every witness of Λ, the α- and δ-components are runs of V, and the Θ-component yields a run of no greater length by Lemma 5.2. This proves containment and the distance inequality. The bounds on α, δ, C and Θ give the parameter bounds. □

5.2.2 Recursive Factorization

We first isolate the interval of high positions, leaving a prefix and suffix that each have a bounded counter. If this interval contains a peak state, the preceding lemma yields a mixed-LPS factor. Otherwise, we either obtain bounded-counter factors directly or recurse on the middle subrun, whose maximum first encoded coordinate is strictly smaller.

Lemma 5.4. *Every run π in W can be partitioned into $\pi = \alpha\pi'\beta$ such that:*

1. *Every nonempty path among α, β has a counter that stays below $H + \|T\|$.*
2. *Every position of π' is high in π, unless π' is empty.*

Proof. Number the positions of π from 0 to ℓ. If some counter stays below H throughout π, take $\alpha = \pi$ and let π', β be empty at its final configuration. Otherwise, let a_j, b_j be the first and last positions where counter j reaches at least H, and set $a := \max_j a_j$ and $b := \min_j b_j$. The high positions are precisely $[a, b]$ when $a \le b$, and there are none otherwise. Choose j_- and j_+ attaining a and b, respectively. Let α be the subrun from position 0 to a, let π' be empty if $a > b$ and the subrun from a to b otherwise, and let β be the subrun from $\max\{a, b\}$ to ℓ. By construction, every position of π' is high in π whenever π' is nonempty. Moreover, counter j_- stays below H before position a, while counter j_+ stays below H after position b. Since each transition changes a counter by at most $\|T\|$, every nonempty path among α, β has a counter bounded by $H - 1 + \|T\|$, as required. □

We now combine the preceding two lemmas to prove the linear factorization.

Proof of Lemma 5.1. Fix a run π in W. For a run ρ in W, let $\mathrm{peak}(\rho) \in [0, B]$ be the maximum first coordinate of its states. If π is empty, it is captured by an empty $\mathbb{N}$-LPS, so we assume that π is nonempty. Put $D := \max\{B, H - 1 + \|T\|\}$ and set $\rho_0 := \pi$. We recursively process ρ_i as follows.

At stage i, apply Lemma 5.4 to obtain $\rho_i = \alpha_i \rho_{i+1} \beta_i$, where ρ_{i+1} spans the high interval of ρ_i, or is empty if no high position exists. We distinguish two cases:

- If the high interval contains a peak state of ρ_i, then Lemma 5.3 gives a mixed LPS Λ containing the endpoints of ρ_i in its reachability relation, with
$$\|\Lambda\|, |\Lambda| \le \mathsf{poly}(B + \textsc{size}(V)), \qquad \mathrm{Reach}_\Lambda \subseteq \mathrm{Reach}_V,$$
and $\ell_V(\boldsymbol{x}, \boldsymbol{y}) \le \ell_\Lambda(\boldsymbol{x}, \boldsymbol{y})$ for all $(\boldsymbol{x}, \boldsymbol{y}) \in \mathrm{Reach}_\Lambda$. We record Reach_Λ as one factor and terminate.
- Otherwise, each nonempty path among α_i, β_i has some counter j bounded by D, and is therefore captured by $[\![V]\!]_{I \cup \{j\}, D}$. We record the corresponding factor
$$R = \mathrm{Reach}_{[\![V]\!]_{I \cup \{j\}, D}} \subseteq \mathrm{Reach}_V,$$
which satisfies $\ell_V(\boldsymbol{x}, \boldsymbol{y}) \le \ell_{[\![V]\!]_{I \cup \{j\}, D}}(\boldsymbol{x}, \boldsymbol{y})$ for all $(\boldsymbol{x}, \boldsymbol{y}) \in R$. If ρ_{i+1} is empty, we terminate. Otherwise, we continue with ρ_{i+1}.

At every nonterminal stage, the high interval contains no peak state of ρ_i, so
$$\mathrm{peak}(\rho_{i+1}) < \mathrm{peak}(\rho_i).$$
Since $\mathrm{peak}(\rho_i) \in [0, B]$, there are at most B nonterminal stages. Each such stage contributes at most two bounded-counter factors, while the terminal stage contributes at most two factors. Thus the resulting composition has at most $2B + 2$ factors, ordered according to their corresponding subruns. The factors produced above have uniformly bounded parameters, including at most five components per mixed LPS, so they belong to a finite family $\mathcal{R}$.

Include in $\mathcal{R}$ the identity relation represented by the empty $\mathbb{N}$-LPS, and choose a common polynomial P for the parameter bounds. Let $\mathcal{F}$ contain all nonempty sequences of at most $2B + 2$ relations from $\mathcal{R}$. The recursive construction gives coverage of $\mathrm{Reach}_{[\![V]\!]_{I,B}}$, and transitivity of Reach_V gives the other inclusion. □

The first inclusion in Lemma 5.1 may be strict: cycle relocation uses the enlarged encoding $[\![V]\!]_{2B}$, and reconstruction in V need not preserve the original bound B. This suffices for the polynomial factorization, since every resulting factor is contained in Reach_V.

5.3 Applications of Polynomial Factorization

Mixed LPS for 5-VAS Reachability. In a 5-VAS, encoding the three bounded counters of a factor into states leaves a 2-VASS of polynomial size. By [BEF+21, Theorem 3.1], its reachability relation is covered by finitely many $\mathbb{N}$-LPS with polynomially bounded parameters. The decoded LPS cover the factor, and each defines a subrelation of Reach_V. Concatenating these LPS with the mixed-LPS factors proves Corollary 1.4. See Appendix A.2 for the full proof.

Fixing which cycles occur in a mixed-LPS witness yields a system $\mathcal{E}_\Lambda$ of linear inequalities over nonnegative integers, with the source and target vectors as variables. The system has at most $\mathrm{poly}(\mathrm{SIZE}(V))$ variables and inequalities, and its coefficients and constants have norm at most $\mathrm{poly}(\mathrm{SIZE}(V))$. By Corollary 3.4, its solution set is a finite union of linear sets whose bases and periods have norm at most $2^{\mathrm{poly}(\mathrm{SIZE}(V))}$. Projecting onto the source and target coordinates and taking the union over all mixed LPS and cycle choices yields

$$\mathrm{Reach}_V = \bigcup_{i\in\mathcal{I}} (\boldsymbol{b}_i + P_i^*), \qquad \|\boldsymbol{b}_i\|, \|P_i\| \le 2^{\mathrm{poly}(\mathrm{SIZE}(V))}.$$

Since the family is effectively enumerable, Reach_V is effectively semilinear.

Now let $M := \mathrm{SIZE}(V, \boldsymbol{s}, \boldsymbol{t})$, and let n be the binary input length. Fixing $\boldsymbol{s}, \boldsymbol{t}$ in $\mathcal{E}_\Lambda$ gives a system $\mathcal{E}_{\Lambda,\boldsymbol{s},\boldsymbol{t}}$ with polynomially bounded size and coefficients in M. By Corollary 3.4, a feasible system has a solution of binary length polynomial in M, yielding a NEXPTIME algorithm since M is at most exponential in n. Moreover, the cycle repetition counts are at most $2^{\mathrm{poly}(M)}$, and every witness yields a run in V of no greater length, so $\ell_V(\boldsymbol{s}, \boldsymbol{t}) \le 2^{\mathrm{poly}(M)}$. The PSPACE upper bound uses the stronger polynomial reachability-distance bound of Theorem 1.2, proved in the following sections.

Reachability in $(d+3)$-VAS Is in F_d. For a $(d+3)$-VAS, encode the three bounded counters of each bounded-counter factor into states while preserving all original counter effects. Every cycle returns to the same encoded state, so it has zero effect on these three counters and the resulting gadget has cycle-dimension at most d. The bound of [CFZ26, Theorem 2.1] concerns the maximum cycle-dimension of an SCC, so the mixed-LPS factors must also be converted into VASS gadgets with controlled cycle-dimension. To build a VASS gadget, give each path and cycle of the LPS its own copy and handle its $\mathbb{N}$- and $\mathbb{Z}$-components differently:

- an $\mathbb{N}$-component needs no further change, and
- a $\mathbb{Z}$-component is padded by two self-loops, one raising $(+\mathbf{1})$ and one restoring $(-\mathbf{1})$ all original counters, with an auxiliary counter forcing the two repetition counts to be equal.

The auxiliary counter adds a counter, but each SCC of the resulting gadget has at most one simple cycle, so its cycle-dimension is at most 1. Connecting the gadgets in sequence keeps every SCC inside a single gadget, so the compiled VASS has component-dimension at most d, and the F_d bound of [CFZ26, Theorem 2.1] applies. Since every witness becomes a run of V of no greater length, the bound transfers back and proves Theorem 1.5 for $d \ge 3$. The full proof is in Appendix A.3.

6 CFZ Characterization and Basic Reachability-Distance Bounds

We now turn to the proof of Theorem 1.2 starting from a decomposition of 5-VAS reachability into two classes of instances. Recall that for a d-VASS $V = (Q, T)$, the encoding $[\![V]\!]_{I,B}$ has state set

$Q' = Q \times [0, B]^I$ and removes the coordinates in I. We also use $\widehat{[\![V]\!]}_{I,B}$, with the same state set but retaining all d coordinates. For every $(p, \boldsymbol{a}, q) \in T$ and $p(\boldsymbol{u}) \in Q'$, it has the transition

$$\big(p(\boldsymbol{u}), \boldsymbol{a}, q(\boldsymbol{v})\big) \quad \text{precisely when} \quad \boldsymbol{v} = \boldsymbol{u} + \boldsymbol{a} \restriction_I \in [0, B]^I.$$

Retaining the original counters allows two components encoding different coordinate sets to share the same configuration across a zero-effect transition. From a configuration $p(\boldsymbol{u})(\boldsymbol{x})$ satisfying $\boldsymbol{x} \restriction_I = \boldsymbol{u}$, every reachable configuration also preserves this equality.

The CFZ Characterization. Based on the concrete construction underlying [CFZ26, Proposition 5.1], we use the following formulation for $d = 5$. The construction is recalled in Appendix D.1.

Theorem 6.1 (CFZ Characterization [CFZ26, Proposition 5.1]). *For every 5-VAS $(V, \boldsymbol{s}, \boldsymbol{t})$ with $\boldsymbol{s} \xrightarrow{*} \boldsymbol{t}$, there exists $B \leq \textsc{size}(V, \boldsymbol{s}, \boldsymbol{t})^{\mathcal{O}(1)}$ and a VASS (V', s', t') such that $s' \xrightarrow{*} t'$ in V' and $\ell_V(\boldsymbol{s}, \boldsymbol{t}) \leq \ell_{V'}(s', t')$. Moreover, one of the following alternatives holds:*

I. V' consists of two disjoint 5-VASS $\widehat{[\![V'']\!]}_{I_1,B}, \widehat{[\![V'']\!]}_{I_2,B}$, connected by a zero-effect transition from the first to the second. Here, V'' is a 5-VASS of size $\textsc{size}(V, \boldsymbol{s}, \boldsymbol{t})^{\mathcal{O}(1)}$, and for each $i = 1, 2$, we have $|I_i| = 2$ and $\dim_{\text{cyc}}(\widehat{[\![V'']\!]}_{I_i,B}) \leq 2$. The source and target configurations of V' are $s' = (p, \boldsymbol{s} \restriction_{I_1})(\boldsymbol{s})$ and $t' = (q, \boldsymbol{t} \restriction_{I_2})(\boldsymbol{t})$, where p and q are two states in V''.

II. (V', s', t') is a single B-encoding VAS $([\![V]\!]_B, [\![\boldsymbol{s}]\!]_B, [\![\boldsymbol{t}]\!]_B)$ for some $B \leq \textsc{size}(V, \boldsymbol{s}, \boldsymbol{t})^{\mathcal{O}(1)}$.

The constant hidden in the $\mathcal{O}(1)$ exponent above is independent of the input. It suffices to establish polynomial reachability-distance bounds for these two classes.

6.1 Case I: Two Components with Cycle-Dimension at Most 2

We now treat the two cases separately, beginning with Case I. For clarity, we fix a 5-VASS V'', $I_1, I_2 \subseteq [5]$ containing two elements, and $B \in \mathbb{N}$, and consider the reachable instance (V', s', t') obtained from V'' by the construction in Case I. In this part, we establish the following proposition.

Proposition 6.2. *Every reachable instance (V', s', t') in Case I satisfies*

$$\ell_{V'}(s', t') \leq (B + \textsc{size}(V', s', t'))^{\mathcal{O}(1)}.$$

It is crucial to observe that both $\widehat{[\![V'']\!]}_{I_1,B}$ and $\widehat{[\![V'']\!]}_{I_2,B}$ have cycle-dimension at most 2. Removing the coordinates already recorded in the states yields a 3-VASS $[\![V'']\!]_{I_i,B}$ with cycle-dimension at most two. We use two existing results for such systems.

For a d-VASS V and a state q, write $\text{Reach}_V^q(s) := \{\boldsymbol{x} : q(\boldsymbol{x}) \in \text{Reach}_V(s)\}$ for the slice of $\text{Reach}_V(s)$ at state q. More generally, for $X \subseteq \mathbb{N}^d$, we define $\text{Reach}_V^q(p(X)) := \bigcup_{\boldsymbol{x} \in X} \text{Reach}_V^q(p(\boldsymbol{x}))$, namely, the set of vectors reachable at q from some configuration $p(\boldsymbol{x})$ with $\boldsymbol{x} \in X$. The first result represents configurations reachable from a hybrid-set source by a finite union of hybrid sets with polynomially bounded representation norms.

Theorem 6.3 (Efficient Hybrid-Set Representations [CFF+26, Theorem 4.2]). *Let V be a 3-VASS with $\dim_{\text{cyc}}(V) \leq 2$. Let $H = \boldsymbol{b} + P^{\langle \boldsymbol{c}, \cdot \rangle \leq K}$ be a hybrid set whose periods satisfy $P \subseteq \textsc{cyc}(V)$. For any states p, q, the set $\text{Reach}_V^q(p(H))$ is a finite union of hybrid sets H' satisfying*

$$\textsc{norm}(H') \leq \mathsf{poly}(\textsc{size}(V) + \textsc{norm}(H)).$$

To bound $\ell_V(s,t)$, we also need the following reachability-distance bound for 3-VASS of cycle-dimension at most 2 with a hybrid source.

Lemma 6.4 ([CFF+26, Lemma 4.3]). *Let V be a 3-VASS with $\dim_{\mathrm{cyc}}(V) \le 2$, t be a configuration, and H be a hybrid set. If $p(\boldsymbol{x}) \xrightarrow{*} t$ for some $\boldsymbol{x} \in H$, there exists $\boldsymbol{x}' \in H$ with*

$$\|\boldsymbol{x}'\|, \ell_V(p(\boldsymbol{x}'),t) \le \mathsf{poly}(\mathrm{SIZE}(V,t) + \mathrm{NORM}(H)).$$

Together, these results allow us to describe the reachable midpoints of the first component by hybrid sets and select a midpoint admitting a short run through the second. Since the two components may encode different coordinates when $I_1 \neq I_2$, we must ensure that the selected midpoint is compatible with both encodings.

Proof of Proposition 6.2. Set $M := B + \mathrm{SIZE}(V',s',t')$. For $i = 1,2$, write $J_i = [5] \setminus I_i$ and $V_i := [\![V'']\!]_{I_i,B}$ for the 3-dimensional projection of the i-th component. Every cycle in V_i lifts to a cycle in the original component, with zero effect on I_i. Hence $\dim_{\mathrm{cyc}}(V_i) \le 2$. Let s_1, t_2 be the projections of s', t' into V_1, V_2, respectively. Write $(r_1(\boldsymbol{u}_1), \boldsymbol{0}, r_2(\boldsymbol{u}_2))$ for the zero-effect bridge, where $\boldsymbol{u}_i \in [0,B]^{I_i}$. It is easy to see that, for any run in (V',s',t'), the configurations at the two ends of the bridge share the same vector $\boldsymbol{w}$, with $\boldsymbol{w}\restriction_{I_1} = \boldsymbol{u}_1$ and $\boldsymbol{w}\restriction_{I_2} = \boldsymbol{u}_2$.

Apply Theorem 6.3 to V_1, with the singleton source specified by s_1. This set satisfies the cycle-space hypothesis. Choose a hybrid set in the resulting union at state $r_1(\boldsymbol{u}_1)$ that contains $\boldsymbol{w}\restriction_{J_1}$. By inserting $\boldsymbol{u}_1$ into its base on I_1 and zeros into its periods on I_1, we obtain a 5-dimensional hybrid set

$$H = \boldsymbol{b} + P^{\langle \boldsymbol{c}, \cdot \rangle \le K}, \qquad \mathrm{NORM}(H) \le \mathsf{poly}(M),$$

containing $\boldsymbol{w}$, all of whose vectors are reachable at the bridge through the first component. Then we impose the second restriction $\boldsymbol{x}\restriction_{I_2} = \boldsymbol{u}_2$ before projecting to J_2.

Fix a representation $\boldsymbol{w} = \boldsymbol{b} + \sum_{\boldsymbol{p}\in P} \lambda_{\boldsymbol{p}} \cdot \boldsymbol{p}$ with $\sum_{\boldsymbol{p}\in P} \boldsymbol{c}(\boldsymbol{p})\lambda_{\boldsymbol{p}} \le K$, where $\boldsymbol{c}(\boldsymbol{p})$ denotes the cost of period $\boldsymbol{p}$. Partition P into $P_0 = \{\boldsymbol{p} \in P : \boldsymbol{p}\restriction_{I_2} = \boldsymbol{0}\}$ and $P_1 = P \setminus P_0$. Each period in P_1 contributes at least 1 to the sum of the coordinates in I_2. Since all bases and periods are nonnegative, we have

$$\sum_{\boldsymbol{p}\in P_1} \lambda_{\boldsymbol{p}} \le \sum_{j\in I_2} \boldsymbol{w}(j) \le 2B.$$

Absorb these boundedly many period occurrences into the base by setting $\boldsymbol{b}_0 := \boldsymbol{b} + \sum_{\boldsymbol{p}\in P_1} \lambda_{\boldsymbol{p}} \cdot \boldsymbol{p}$ and $K_0 := K - \sum_{\boldsymbol{p}\in P_1} \boldsymbol{c}(\boldsymbol{p})\lambda_{\boldsymbol{p}} \ge 0$. Then the resulting hybrid set

$$H_0 := \boldsymbol{b}_0 + P_0^{\langle \boldsymbol{c}\restriction_{P_0}, \cdot \rangle \le K_0} \subseteq H, \qquad \mathrm{NORM}(H_0) \le (2B+1)\,\mathrm{NORM}(H) \le \mathsf{poly}(M).$$

The set H_0 contains $\boldsymbol{w}$, and every $\boldsymbol{x} \in H_0$ satisfies $\boldsymbol{x}\restriction_{I_2} = \boldsymbol{u}_2$. Its projection $H_2 := \{\boldsymbol{x}\restriction_{J_2} : \boldsymbol{x} \in H_0\}$ is therefore a hybrid set of no larger norm. The original suffix shows that some vector in H_2 reaches t_2 from state $r_2(\boldsymbol{u}_2)$ in V_2. By Lemma 6.4, there are $\boldsymbol{z} \in H_2$ and a run in V_2 satisfying

$$r_2(\boldsymbol{u}_2)(\boldsymbol{z}) \xrightarrow{\rho_2} t_2, \qquad \|\boldsymbol{z}\|, |\rho_2| \le \mathsf{poly}(M).$$

Insert $\boldsymbol{u}_2$ into $\boldsymbol{z}$ to obtain its unique lift $\boldsymbol{w}' \in H_0$. Then $\|\boldsymbol{w}'\| \le \max\{B, \|\boldsymbol{z}\|\} \le \mathsf{poly}(M)$.

Since $\boldsymbol{w}' \in H_0 \subseteq H$, the configuration $r_1(\boldsymbol{u}_1)(\boldsymbol{w}'\restriction_{J_1})$ is reachable from s_1 in V_1. Apply Lemma 6.4 once more, now with the singleton source specified by s_1 and this fixed target. A singleton leaves no

choice of source vector, so this gives a run ρ_1 from s_1 to that target with $|\rho_1| \le \mathrm{poly}(M + \|\boldsymbol{w}'\|) \le \mathrm{poly}(M)$. Both runs lift to their respective five-dimensional components without changing their lengths. They meet at the same vector $\boldsymbol{w}'$, so the zero-effect bridge joins them. Consequently,

$$\ell_{V'}(s', t') \le |\rho_1| + 1 + |\rho_2| \le M^{\mathcal{O}(1)},$$

which completes the proof. □

6.2 Case II: A Single Encoding VAS

We now consider Case II. Let $W = [\![V]\!]_B$ be the B-encoding of a 5-VAS V, and fix a run $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{\pi} \boldsymbol{q}(\boldsymbol{y})$ in W, where $\boldsymbol{s} = (\boldsymbol{p}, \boldsymbol{x})$ and $\boldsymbol{t} = (\boldsymbol{q}, \boldsymbol{y})$. The SCCs traversed by π, together with their connecting transitions, form a sequential sub-VASS $U = (U_1)u_1(U_2)\cdots u_{k-1}(U_k)$. Since W has $(B+1)^2$ states, we have $k \le (B+1)^2$. Our goal is to establish a polynomial bound on $\ell_V(\boldsymbol{s}, \boldsymbol{t})$.

Proposition 6.5. *Let V be a 5-VAS and let $(U, \boldsymbol{p}(\boldsymbol{x}), \boldsymbol{q}(\boldsymbol{y}))$ be a reachable instance specified above. Then*

$$\ell_V(\boldsymbol{s}, \boldsymbol{t}) \le \big(B + \mathrm{SIZE}(V, \boldsymbol{s}, \boldsymbol{t})\big)^{\mathcal{O}(1)}. \tag{4}$$

Together with Theorem 6.1 and Proposition 6.2, Proposition 6.5 implies Theorem 1.2. The formal deduction is given in Appendix D.2. We first establish polynomial reachability-distance bounds for two basic cases: non-wide instances and C-coverable instances as specified in Section 2. We formalize the related notions below.

Definition 6.6 (Wideness)**.** Let U be a d-VASS. Let $E(U)$ be the set of effects of simple cycles in U, and let $\mathrm{cone}(U) := \mathrm{cone}(E(U))$. A d-VASS U is called *wide* if $\mathrm{cone}(U) = \mathbb{Q}^d$.

This notion differs from the wideness conditions of [CJLO25, CFF+26], which use positive-orthant containment $\mathbb{Q}^3_{\ge 0} \subseteq \mathrm{cone}(U)$. Definition 6.6 considers the full space $\mathbb{Q}^3$ because it permits compensation in every direction.

Definition 6.7 (C-Coverability)**.** Let $(U, p(\boldsymbol{x}), q(\boldsymbol{y}))$ be a d-VASS. For $C \in \mathbb{N}$, the source $p(\boldsymbol{x})$ is *forward C-coverable* in U if $p(\boldsymbol{x}) \xrightarrow{*} p(\boldsymbol{u})$ for some $\boldsymbol{u} \ge C \cdot \mathbf{1}$. The target $q(\boldsymbol{y})$ is *backward C-coverable* if $q(\boldsymbol{v}) \xrightarrow{*} q(\boldsymbol{y})$ for some $\boldsymbol{v} \ge C \cdot \mathbf{1}$. The VASS $(U, p(\boldsymbol{x}), q(\boldsymbol{y}))$ is *C-coverable* if both conditions hold.

6.2.1 Non-Wide Case

The non-wide case follows the standard argument of [CJLO25, Lemma 9]. If $\dim_{\mathrm{cyc}}(U) \le 2$, the preceding low-dimensional result applies directly. Otherwise a proper cycle cone is contained in a half-space with a nonzero normal. The inner product with this normal is bounded along a given run. Recording that inner product in the state reduces the entire system to cycle-dimension at most two. We state the reduction below and defer its proof to Appendix D.3.

Lemma 6.8. *Let (U, s, t) be a reachable 3-VASS with $\mathrm{cone}(U) \ne \mathbb{Q}^3$. There exists a reachable 3-VASS $(\widetilde{U}, \widetilde{s}, \widetilde{t})$ of size at most $\mathrm{SIZE}(U, s, t)^{\mathcal{O}(1)}$ such that $\dim_{\mathrm{cyc}}(\widetilde{U}) \le 2$ and $\ell_U(s, t) \le \ell_{\widetilde{U}}(\widetilde{s}, \widetilde{t})$.*

Apply Lemma 6.4 from the previous subsection to $\widetilde{U}$ with a singleton $\{\widetilde{s}\}$. It gives a run from $\widetilde{s}$ to $\widetilde{t}$ with length at most $\mathrm{poly}(\mathrm{SIZE}(\widetilde{U}, \widetilde{s}, \widetilde{t}))$, hence the following.

Corollary 6.9. *Let (U, s, t) be a reachable 3-VASS with $\mathrm{cone}(U) \ne \mathbb{Q}^3$. Its reachability distance satisfies*

$$\ell_U(s, t) \le \mathrm{SIZE}(U, s, t)^{\mathcal{O}(1)}.$$

6.2.2 Wide and Coverable Case

We now consider the wide case, where $\mathrm{cone}(U) = \mathbb{Q}^3$. Combined with the transformation-area theorem, this gives a reachability-distance bound under C-coverability. We state the result more generally for B-encodings of $(d+2)$-VAS.

Theorem 6.10 (Wide and Coverable Case)**.** *Let V be a $(d+2)$-VAS, $B \in \mathbb{N}$, and $(U, \boldsymbol{p}(\boldsymbol{x}), \boldsymbol{q}(\boldsymbol{y}))$ be a sequential sub-VASS of $[\![V]\!]_B$. Assume that the source and target lie in the first and last components of U, respectively. If $(U, \boldsymbol{p}(\boldsymbol{x}), \boldsymbol{q}(\boldsymbol{y}))$ is wide and C-coverable and $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{*}_{\mathbb{Z}} \boldsymbol{q}(\boldsymbol{y})$ in U, then*

$$\ell_V(\boldsymbol{s}, \boldsymbol{t}) \leq \big(B + C + \mathrm{SIZE}(V, \boldsymbol{s}, \boldsymbol{t})\big)^{(d+1)^{\mathcal{O}(d)}}, \tag{5}$$

where C is a transformation-area threshold for $[\![V]\!]_B$, $\boldsymbol{s} = (\boldsymbol{p}, \boldsymbol{x})$ and $\boldsymbol{t} = (\boldsymbol{q}, \boldsymbol{y})$.

For $d = 3$, the threshold supplied by Theorem 4.2 is polynomial in $B + \mathrm{SIZE}(V)$, giving the required polynomial reachability-distance bound.

Short Runs Between Configurations in $\mathcal{G}_C$. The transformation-area theorem converts a $\mathbb{Z}$-run between configurations in $\mathcal{G}_C$ into a run of the same length, but does not bound that length. We therefore use the following shortening result for $\mathbb{Z}$-runs, whose proof is given in Appendix D.4.

Lemma 6.11. *Let $U = (Q, T)$ be a d-VASS. If $p(\boldsymbol{x}) \xrightarrow{*}_{\mathbb{Z}} q(\boldsymbol{y})$ in U, there is such a $\mathbb{Z}$-run of length at most $(\mathrm{SIZE}(U) + \|\boldsymbol{x}\| + \|\boldsymbol{y}\|)^{\mathsf{poly}(d)}$.*

By Lemma 6.11, every $\mathbb{Z}$-run in $[\![V]\!]_B$ can be shortened to polynomial length in the size of the instance for fixed d. If its endpoints lie in $\mathcal{G}_C$, Theorem 4.2 then converts the shortened $\mathbb{Z}$-run into a run in V with the same Parikh image, and hence the same length. This gives the following corollary.

Corollary 6.12. *Let V be a $(d+2)$-VAS, $B \in \mathbb{N}$, and U a sub-VASS of $[\![V]\!]_B$. Let C be a transformation-area threshold for $[\![V]\!]_B$. If $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{*}_{\mathbb{Z}} \boldsymbol{q}(\boldsymbol{y})$ in U and $\boldsymbol{x}, \boldsymbol{y} \geq C \cdot \mathbf{1}$, then*

$$\ell_V\big((\boldsymbol{p}, \boldsymbol{x}), (\boldsymbol{q}, \boldsymbol{y})\big) \leq \big(\mathrm{SIZE}(U) + \|\boldsymbol{x}\| + \|\boldsymbol{y}\|\big)^{\mathsf{poly}(d)}.$$

Compatibility of the Covering Configurations. By Corollary 6.12, we can obtain a short run once the endpoints of a $\mathbb{Z}$-run lie in $\mathcal{G}_C$. For the C-coverable case, we use Lemma 3.1 to extract short covering runs. However, the resulting configurations in $\mathcal{G}_C$ need not be reachable by the original $\mathbb{Z}$-run. The following lemma uses wideness to provide a $\mathbb{Z}$-run by inserting suitable cycles.

Lemma 6.13. *Let $U = (U_1)u_1 \cdots u_{k-1}(U_k)$ be wide, with p in U_1 and q in U_k. Suppose that*

$$p(\boldsymbol{x}) \xrightarrow{\alpha}_{\mathbb{Z}} p(\boldsymbol{u}), \qquad p(\boldsymbol{x}) \xrightarrow{*}_{\mathbb{Z}} q(\boldsymbol{y}), \qquad q(\boldsymbol{v}) \xrightarrow{\beta}_{\mathbb{Z}} q(\boldsymbol{y})$$

in U. Then $p(\boldsymbol{u}) \xrightarrow{*}_{\mathbb{Z}} q(\boldsymbol{v})$ in U.

Proof. Let $E := E(U)$ be the set of effects of simple cycles in U, and let E^* denote the set of their nonnegative integer combinations. We claim that E^* is closed under negation. Indeed, for every $\boldsymbol{e} \in E$, wideness gives $-\boldsymbol{e} \in \mathrm{cone}(E)$. Since $\mathrm{cone}(E)$ consists of nonnegative rational combinations of vectors in E, there exists an integer $m \geq 1$ such that $-m \cdot \boldsymbol{e} \in E^*$. Consequently, $-\boldsymbol{e} = (-m \cdot \boldsymbol{e}) + (m-1) \cdot \boldsymbol{e} \in E^*$.

Fix a $\mathbb{Z}$-run $p(\boldsymbol{x}) \xrightarrow{\pi}_{\mathbb{Z}} q(\boldsymbol{y})$ in U. In each component U_i, choose a cycle ω_i based at a state visited by π and traversing every state of U_i. Strong connectivity provides such a cycle. For a one-state component we may instead take the empty path. Inserting these paths into π gives a path τ. Let

$$\boldsymbol{a} := \mathrm{EFF}(\alpha) = \boldsymbol{u} - \boldsymbol{x}, \qquad \boldsymbol{b} := \mathrm{EFF}(\beta) = \boldsymbol{y} - \boldsymbol{v}, \qquad \boldsymbol{h} := \sum_{i=1}^{k} \mathrm{EFF}(\omega_i).$$

The paths α, β, ω_i are cycles or empty paths, so their effects belong to E^*. Hence $-\boldsymbol{a} - \boldsymbol{b} - \boldsymbol{h} \in E^*$, which implies that there exists a multiset of simple cycles with effect $-\boldsymbol{a} - \boldsymbol{b} - \boldsymbol{h}$.

Insert these cycles at visits of τ to their initial states, which exist because τ visits every state of U. This yields a path from p to q, whose effect is

$$(\boldsymbol{y} - \boldsymbol{x}) + \boldsymbol{h} - \boldsymbol{a} - \boldsymbol{b} - \boldsymbol{h} = \boldsymbol{v} - \boldsymbol{u}.$$

Hence, there is a $\mathbb{Z}$-run from $p(\boldsymbol{u})$ to $q(\boldsymbol{v})$ in U. □

Proof of Theorem 6.10. Let T be the transition set of V, $N := (B+1)^2$, and $R := \max\{\|\boldsymbol{x}\|, \|\boldsymbol{y}\|\}$. Applying Lemma 3.1 to the source covering run in U and to the target covering run in the reversed system, we obtain two short covering runs

$$\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{\alpha} \boldsymbol{p}(\boldsymbol{u}), \qquad \boldsymbol{q}(\boldsymbol{v}) \xrightarrow{\beta} \boldsymbol{q}(\boldsymbol{y}),$$

with $\boldsymbol{u}, \boldsymbol{v} \geq C \cdot \mathbf{1}$ and $|\alpha|, |\beta| \leq L := \big(\mathrm{SIZE}(U) + C\big)^{\mathcal{O}((d+1)!)}$. Consequently, $\|\boldsymbol{u}\|, \|\boldsymbol{v}\| \leq R + L\,\|T\|$. By Lemma 6.13, there is a $\mathbb{Z}$-run π in U giving

$$\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{\alpha} \boldsymbol{p}(\boldsymbol{u}) \xrightarrow{\pi}_{\mathbb{Z}} \boldsymbol{q}(\boldsymbol{v}) \xrightarrow{\beta} \boldsymbol{q}(\boldsymbol{y}).$$

By Corollary 6.12, the $\mathbb{Z}$-run $\boldsymbol{p}(\boldsymbol{u}) \xrightarrow{\pi}_{\mathbb{Z}} \boldsymbol{q}(\boldsymbol{v})$ can be replaced by a run in V of length at most $M \leq (\mathrm{SIZE}(U) + \|\boldsymbol{u}\| + \|\boldsymbol{v}\|)^{\mathsf{poly}(d)}$. The two covering runs also lift to V, so concatenation gives

$$\ell_V(\boldsymbol{s}, \boldsymbol{t}) \leq 2L + M \leq 2L + (\mathrm{SIZE}(U) + 2R + 2L\,\|T\|)^{\mathsf{poly}(d)}.$$

Finally, U has at most N states and $N|T|$ transitions, all of norm at most $\|T\|$, whence

$$\mathrm{SIZE}(U) \leq N\,\mathrm{SIZE}(V).$$

Combining this with the bounds on N and C yields (5). □

Theorem 6.10 shows that, in the wide and C-coverable case, even $\mathbb{Z}$-reachability suffices to obtain a short run. Without wideness, Lemma 6.8 only shortens an *existing run*. For a reachable sequential instance, first restrict U to the components between its source and target. The forward and backward covering runs stay in their respective components and remain available in this restriction. Since every run in U lifts to V with the same length and $\mathrm{SIZE}(U) \leq (B+1)^2\,\mathrm{SIZE}(V)$, combining the two results gives the following bound for all reachable C-coverable instances.

Corollary 6.14 (Coverable Case)**.** *Let V be a 5-VAS, $B \in \mathbb{N}$, and $(U, \boldsymbol{p}(\boldsymbol{x}), \boldsymbol{q}(\boldsymbol{y}))$ be a sequential sub-VASS of $[\![V]\!]_B$ such that $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{*} \boldsymbol{q}(\boldsymbol{y})$ in U. If $(U, \boldsymbol{p}(\boldsymbol{x}), \boldsymbol{q}(\boldsymbol{y}))$ is C-coverable, then*

$$\ell_V(\boldsymbol{s}, \boldsymbol{t}) \leq \big(B + C + \mathrm{SIZE}(V, \boldsymbol{s}, \boldsymbol{t})\big)^{\mathcal{O}(1)},$$

where C is a transformation-area threshold for $[\![V]\!]_B$, $\boldsymbol{s} = (\boldsymbol{p}, \boldsymbol{x})$, and $\boldsymbol{t} = (\boldsymbol{q}, \boldsymbol{y})$.

6.3 Hybrid Sets as Sources and Targets

The preceding results give polynomial reachability-distance bounds for fixed configurations. In the reductions used later, however, the source and target may range over hybrid sets. We handle such endpoints by reducing the problem to reachability between fixed configurations in an augmented VASS. The idea is to add the periods of the source hybrid set as positive self-loops at the source state, and those of the target hybrid set as negative self-loops at the target state.

Formally, for a VASS $U = (Q, T)$ with specified source and target states $\boldsymbol{p}, \boldsymbol{q}$ and period sets $P_1, P_2 \subseteq \mathbb{N}^d$, define the augmented system:

$$U[P_1, P_2] := \big(Q, T \cup \{(\boldsymbol{p}, \boldsymbol{a}, \boldsymbol{p}) : \boldsymbol{a} \in P_1\} \cup \{(\boldsymbol{q}, -\boldsymbol{a}, \boldsymbol{q}) : \boldsymbol{a} \in P_2\}\big),$$

and abbreviate $U[P, \varnothing]$ to $U[P]$. We first extend C-coverability to hybrid sets using only their free periods, whose multiplicities are unrestricted.

Definition 6.15. Let $H = \boldsymbol{b} + P^{\langle \boldsymbol{c}, \cdot \rangle \le K}$ be a hybrid set at state $\boldsymbol{p}$ of a VASS U, and let P_0 be its free periods. We say that H is *forward C-coverable in U* if $\boldsymbol{p}(\boldsymbol{b})$ is forward C-coverable in $U[P_0]$. Similarly, H at target state $\boldsymbol{q}$ is *backward C-coverable in U* if $\boldsymbol{q}(\boldsymbol{b})$ is backward C-coverable in $U[\varnothing, P_0]$.

Consider a forward C-coverable hybrid set H_1 and a backward C-coverable hybrid set H_2 with bases $\boldsymbol{b}_1, \boldsymbol{b}_2$ and period sets P_1, P_2, respectively. If $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{*} \boldsymbol{q}(\boldsymbol{y})$ in U for some $\boldsymbol{x} \in H_1$ and $\boldsymbol{y} \in H_2$, then $\boldsymbol{p}(\boldsymbol{b}_1) \xrightarrow{*} \boldsymbol{q}(\boldsymbol{b}_2)$ in $U[P_1, P_2]$. Indeed, the added self-loops first generate $\boldsymbol{x}$ from $\boldsymbol{b}_1$ and then reduce $\boldsymbol{y}$ to $\boldsymbol{b}_2$. Moreover, the new instance is C-coverable. Thus, Corollary 6.14 applies. The added periods may violate the cost bounds of H_1 or H_2. However, this can happen only when $K_i < \|\boldsymbol{c}_i\| L$, where L is the polynomial length bound for the augmented run. In this case, we absorb the restricted periods of a fixed connected member into the base. Since each restricted period has positive integer cost, their total multiplicity is at most K. This gives the following lemma.

Lemma 6.16. *Let $H = \boldsymbol{b} + P^{\langle \boldsymbol{c}, \cdot \rangle \le K}$ be a hybrid set, and write $P = P_0 \sqcup P_+$ for its free and restricted periods. Fix a representation $\boldsymbol{z} = \boldsymbol{b} + \sum_{\boldsymbol{a} \in P} \eta_{\boldsymbol{a}} \cdot \boldsymbol{a}$, where $\eta_{\boldsymbol{a}} \in \mathbb{N}$ for all $\boldsymbol{a} \in P$ and $\sum_{\boldsymbol{a} \in P} \boldsymbol{c}(\boldsymbol{a}) \eta_{\boldsymbol{a}} \le K$. Define*

$$\widehat{\boldsymbol{b}} := \boldsymbol{b} + \sum_{\boldsymbol{a} \in P_+} \eta_{\boldsymbol{a}} \cdot \boldsymbol{a}, \qquad \widehat{H} := \widehat{\boldsymbol{b}} + P_0^*.$$

Then $\boldsymbol{z} \in \widehat{H} \subseteq H$ and $\left\|\widehat{\boldsymbol{b}}\right\| \le \|\boldsymbol{b}\| + K \|P\|$.

Proof. Every restricted cost is a positive integer, so $\sum_{\boldsymbol{a} \in P_+} \eta_{\boldsymbol{a}} \le \sum_{\boldsymbol{a} \in P} \boldsymbol{c}(\boldsymbol{a}) \eta_{\boldsymbol{a}} \le K$, which gives the norm bound. Fixing the restricted coefficients gives $\boldsymbol{z} \in \widehat{H} \subseteq H$. □

Moreover, this operation preserves C-coverability. Since $\widehat{\boldsymbol{b}} \ge \boldsymbol{b}$ and the free periods are unchanged, the covering runs can be shifted by $\widehat{\boldsymbol{b}} - \boldsymbol{b}$ using monotonicity. The same argument applies to backward coverability. We use these properties to handle the cost constraints when applying the reachability bounds for fixed configurations to augmented systems. The resulting lemma is stated below, and its proof is deferred to Appendix D.5.

Lemma 6.17. *Let V be a 5-VAS, $B \in \mathbb{N}$, and C a transformation-area threshold for $[\![V]\!]_B$. Let H_1, H_2 be hybrid sets at states $\boldsymbol{p}, \boldsymbol{q}$ of a sequential sub-VASS U of $[\![V]\!]_B$. Suppose that there exist $\boldsymbol{x} \in H_1$ and $\boldsymbol{y} \in H_2$ such that $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{*} \boldsymbol{q}(\boldsymbol{y})$ in U. Assume that H_1 is forward C-coverable in U and H_2 is backward C-coverable in U. Then there exist $\boldsymbol{x}' \in H_1$ and $\boldsymbol{y}' \in H_2$ such that*

$$\|\boldsymbol{x}'\|, \ \|\boldsymbol{y}'\|, \ell_V((\boldsymbol{p}, \boldsymbol{x}'), (\boldsymbol{q}, \boldsymbol{y}')) \le \big(B + C + \mathrm{SIZE}(V) + \mathrm{NORM}(H_1) + \mathrm{NORM}(H_2)\big)^{\mathcal{O}(1)}.$$

7 Additive-Growth Induction for Semi-Coverable Instances

We have established polynomial reachability-distance bounds for C-coverable instances in Section 6. We now show that coverability at just one endpoint suffices. We call an instance *semi-C-coverable* if its source is forward C-coverable or its target is backward C-coverable.

Theorem 7.1 (Semi-Coverable Case)**.** *Let V be a 5-VAS, $B \in \mathbb{N}$, and $(U, \boldsymbol{p}(\boldsymbol{x}), \boldsymbol{q}(\boldsymbol{y}))$ be a sequential sub-VASS of $[\![V]\!]_B$ with $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{*} \boldsymbol{q}(\boldsymbol{y})$ in U. If $(U, \boldsymbol{p}(\boldsymbol{x}), \boldsymbol{q}(\boldsymbol{y}))$ is semi-C-coverable, then*

$$\ell_V(\boldsymbol{s}, \boldsymbol{t}) \leq \big(B + C + \textsc{size}(V, \boldsymbol{s}, \boldsymbol{t})\big)^{\mathcal{O}(1)}, \tag{6}$$

where C is a transformation-area threshold for $[\![V]\!]_B$, $\boldsymbol{s} = (\boldsymbol{p}, \boldsymbol{x})$ and $\boldsymbol{t} = (\boldsymbol{q}, \boldsymbol{y})$.

We prove the theorem when the target is backward C-coverable, without assuming coverability of the source. The case of a forward C-coverable source follows by reversal. The proof proceeds by induction on the number k of SCCs. We aim to ensure that each step increases the source norm by an additive polynomial bound, so that the final polynomial degree does not depend on k. By Corollary 4.7, the same transformation-area threshold C remains valid under reversal and when adjoining the nonnegative or nonpositive period transitions during the induction.

A Maximal Bounded-Coordinate Prefix. If the source is not forward C-coverable, Rackoff's extraction implies that some coordinate remains below a polynomial threshold along the first SCC segment of a witnessing run. We follow this coordinate through successive SCCs and stop at the entry of the first later SCC in which it reaches the threshold, or at the target if no such SCC exists. In the former case, any configuration at this entry state becomes forward C-coverable within that SCC if the other two coordinates are sufficiently large and the tracked coordinate is not decreased.

Lemma 7.2. *Let $U = (U_1)u_1(U_2)\cdots u_{k-1}(U_k)$ be a sequential 3-VASS, and fix a run $p(\boldsymbol{x}) \xrightarrow{\pi} q(\boldsymbol{y})$ whose source is not forward C-coverable. Then there exists an integer $D \leq (C + \textsc{size}(U))^{\mathcal{O}(1)}$ such that one can find a coordinate $i \in [3]$ and a prefix $p(\boldsymbol{x}) \xrightarrow{\pi_0} r(\boldsymbol{z})$ with $i \in \textsc{bd}(\pi_0, D)$ satisfying one of the following:*

1. *$r(\boldsymbol{z}) = q(\boldsymbol{y})$;*
2. *$r(\boldsymbol{z})$ is the entry configuration of a later component U_h, and every $\boldsymbol{w} \in \mathbb{N}^3$ satisfying*

$$\boldsymbol{w}(i) \geq \boldsymbol{z}(i), \qquad \boldsymbol{w}(j) \geq D \quad \text{for all } j \neq i$$

   has $r(\boldsymbol{w})$ forward C-coverable in U_h.

Moreover, π_0 is captured by a 2-VASS of size at most $D\,\textsc{size}(U)$.

The proof uses Rackoff's bounds to establish Item 2 at the first component where the selected coordinate ceases to remain bounded. We defer the details to Appendix E.1.

Additive-Growth Induction and Prefix Classification. Throughout this section, fix a 5-VAS V, an encoding bound $B \in \mathbb{N}$, and a transformation-area threshold C for $W := [\![V]\!]_B$. Let $U = (U_1)u_1(U_2)\cdots u_{k-1}(U_k)$ be a reachable sequential sub-VASS of W whose source s and target t belong to U_1 and U_k, respectively. Their corresponding configurations in V are $\boldsymbol{s} = (\boldsymbol{p}, \boldsymbol{x})$ and

$\boldsymbol{t} = (\boldsymbol{q}, \boldsymbol{y})$. For the proof of Theorem 7.1, assume that t is backward C-coverable in U. Write $U = (Q, T)$. Since $\textsc{size}(W) \leq (B+1)^2 \textsc{size}(V)$, the threshold in Lemma 7.2 can be chosen uniformly over all sequential subinstances of W. By enlarging it if necessary, we fix

$$B + C + \textsc{size}(V) \leq D \leq (B + C + \textsc{size}(V))^{\mathcal{O}(1)}. \tag{7}$$

If s is forward C-coverable, Corollary 6.14 gives the required polynomial bound. Otherwise, Lemma 7.2 gives a partition

$$\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{\pi_0} \boldsymbol{r}(\boldsymbol{z}) \xrightarrow{\pi_1} \boldsymbol{q}(\boldsymbol{y}),$$

where some coordinate i remains below D throughout π_0. The prefix is captured by a 2-VASS of size at most $D\,\textsc{size}(U)$, so Lemma 6.4 gives a polynomial-length replacement in terms of the system size and endpoint norms. If π_1 is empty, lifting this replacement to V proves the desired bound. Otherwise, the suffix lies in a sequential subinstance with fewer SCCs.

Fix a structural polynomial threshold $R \geq D$, to be chosen in Section 7.3. We call π_0 *R-small* if $\|\boldsymbol{z}\| \leq \|\boldsymbol{x}\| + R$. In this case, induction applies to the unchanged suffix, with the source norm increased by at most R. Otherwise, we call π_0 *R-large*. Since coordinate i remains below $D \leq R$, this means that

$$\boldsymbol{z}(j) > \|\boldsymbol{x}\| + R \qquad \text{for some } j \neq i.$$

For a suitable choice of R, every R-large prefix yields either a hybrid set with a nonzero free period or a translated family. The tools for these two outcomes are developed in Sections 7.1 and 7.2, respectively. In Section 7.3, we establish this dichotomy and use it to complete the induction.

7.1 Pumpable Source

A configuration $p(\boldsymbol{x})$ is *pumpable* in U if an enabled cycle returns to $p(\boldsymbol{x} + \boldsymbol{a})$ for some $\boldsymbol{a} \geq \boldsymbol{0}$ with $\boldsymbol{a} \neq \boldsymbol{0}$. Such a cycle can be repeated indefinitely, but it may leave some coordinates unchanged. Thus, pumpability does not necessarily imply forward C-coverability. We show that it nevertheless suffices for a polynomial reachability-distance bound when the target is backward C-coverable.

Lemma 7.3. *Let U be a sequential sub-VASS of $[\![V]\!]_B$ and let $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{*} \boldsymbol{q}(\boldsymbol{y})$ in U, with $\boldsymbol{s} = (\boldsymbol{p}, \boldsymbol{x})$ and $\boldsymbol{t} = (\boldsymbol{q}, \boldsymbol{y})$. If $\boldsymbol{p}(\boldsymbol{x})$ is pumpable in U and $\boldsymbol{q}(\boldsymbol{y})$ is backward C-coverable, then*

$$\ell_V(\boldsymbol{s}, \boldsymbol{t}) \leq \big(B + C + \textsc{size}(V, \boldsymbol{s}, \boldsymbol{t})\big)^{\mathcal{O}(1)}. \tag{8}$$

The proof starts with a short pumping cycle and considers the coordinates left unchanged by its effect. Using bounded-coordinate extraction, we either find a prefix with one of these coordinates bounded throughout or obtain two ordered cycles whose effects, together with the initial pumping cycle, yield a strictly positive free period. In particular, the positive free period ensures forward C-coverability, allowing us to apply Lemma 6.17 to the backward C-coverable target.

D-Alternating Prefixes. We first strengthen Lemma 7.2 using the effect $\boldsymbol{a}_0$ of the initial pumping cycle. By temporarily adding $\boldsymbol{a}_0$ as a self-loop at each component entry, we ensure that every coordinate in $\operatorname{supp}(\boldsymbol{a}_0)$ becomes large. Thus, whenever forward C-coverability fails, the bounded coordinate extracted by Lemma 7.2 must lie outside $\operatorname{supp}(\boldsymbol{a}_0)$. Since this bounded coordinate may change between components, we introduce the following alternating condition.

Definition 7.4. For $D \geq 1$ and $I \subseteq [3]$, we say that a run π is *D-alternating on I* if it admits a partition $\pi = \pi_1 \cdots \pi_m$, where $m \geq 0$, such that for $S_j := \textsc{bd}(\pi_j, D) \cap I$,

$$S_j \neq \varnothing \text{ for all } j \in [m], \qquad S_j \cap S_{j+1} = \varnothing \text{ for all } j \in [m-1]. \tag{9}$$

When $|I| = 1$, a nonempty D-alternating run keeps that coordinate bounded throughout. When $|I| = 2$, the bounded coordinates alternate between consecutive segments.

Lemma 7.5. *Let $U = (U_1)u_1 \cdots u_{k-1}(U_k)$ be a sequential 3-VASS, let $p(\boldsymbol{x}) \xrightarrow{\pi} t$ be a run in U, and let $\boldsymbol{a}_0 \in \mathbb{N}^3 \setminus \{\mathbf{0}\}$ be the effect of an enabled cycle at $p(\boldsymbol{x})$. There exist*

$$D \leq (C + \textsc{size}(U) + \|\boldsymbol{a}_0\|)^{\mathcal{O}(1)} \tag{10}$$

and a possibly empty prefix $p(\boldsymbol{x}) \xrightarrow{\pi_*} r_*(\boldsymbol{z}_*)$ that is D-alternating on $[3] \setminus \operatorname{supp}(\boldsymbol{a}_0)$. Either $r_*(\boldsymbol{z}_*) = t$, or it is the entry of some component U_* such that $r_*(\boldsymbol{z}_*)$ is forward C-coverable in $U_*[\{\boldsymbol{a}_0\}]$.

The proof repeatedly applies Lemma 7.2 after adding $\boldsymbol{a}_0$ as a self-loop at each SCC entry until the target is reached or a forward C-coverable entry is found. Merging consecutive segments with a common bounded coordinate gives the required alternation. See Appendix E.2 for its proof.

7.1.1 Extracting Two Ordered Cycles for a Positive Free Period

By Lemma 7.5, the coordinates left unchanged by the initial pumping cycle are D-alternating along a prefix. After permuting coordinates so that the pumping cycle increases coordinate 3, this prefix is also D-alternating on $\{1,2\}$: if only one coordinate is left unchanged, the partition has at most one segment. The following lemma shows that either some coordinate remains bounded throughout the prefix under a larger polynomial threshold, or there are two ordered cycles, each increasing one coordinate while leaving the other unchanged. In the latter case, repeating the initial pumping cycle sufficiently many times makes the combined effect positive on coordinate 3 as well, yielding a strictly positive free period.

Lemma 7.6. *Let $U = (Q, T)$ be a 3-VASS and let $p(\boldsymbol{x}) \xrightarrow{\pi} t$ be a run D-alternating on $\{1,2\}$. There exist*

$$D_1, L \leq \left(D + \textsc{size}(U) + \|\boldsymbol{x}\|\right)^{\mathcal{O}(1)}$$

such that either $\textsc{bd}(\pi, D_1) \neq \varnothing$, or, after exchanging the first two counters, there is a partition

$$p(\boldsymbol{x}) \xrightarrow{\rho_1} q_1(\boldsymbol{z}_1) \xrightarrow{\rho_2} q_2(\boldsymbol{z}_2) \xrightarrow{\rho} t \tag{11}$$

and enabled cycles θ_1, θ_2 at $q_1(\boldsymbol{z}_1)$ and $q_2(\boldsymbol{z}_2)$, respectively, with $|\theta_\iota| \leq L$, $\textsc{eff}(\theta_\iota)(\iota) > 0$ and $\textsc{eff}(\theta_\iota)(3-\iota) = 0$ for $\iota \in \{1,2\}$. Moreover, $\{1,2\} \subseteq \textsc{bd}(\rho_1, D_1)$ and $2 \in \textsc{bd}(\rho_1\rho_2, D_1)$.

We first isolate a hill-cutting argument. It extracts a short enabled cycle when one coordinate remains bounded while another increases sufficiently. The large increase produces many positions where the growing coordinate reaches a new maximum. There are at most $D|Q|$ pairs consisting of a state and the bounded-coordinate value. Two records must therefore share such a pair while the growing coordinate increases, yielding a cycle that can be shortened using Rackoff's covering bound. The proof is deferred to Appendix E.3.

Lemma 7.7. *Let $U = (Q, T)$ be a 3-VASS, and let $i, j \in [3]$ be distinct coordinates. Suppose $p(\boldsymbol{x}) \xrightarrow{\pi} q(\boldsymbol{y})$ is a run with $D \geq 1$ satisfying $j \in \mathrm{BD}(\pi, D)$ and*

$$\boldsymbol{y}(i) \geq D' := \boldsymbol{x}(i) + D|Q|\,\|T\| + 1.$$

Then there exist a configuration $r(\boldsymbol{z})$ along π and a cycle θ fireable at $r(\boldsymbol{z})$ such that

$$\mathrm{EFF}(\theta)(i) > 0, \qquad \mathrm{EFF}(\theta)(j) = 0, \qquad |\theta| \leq (D + \mathrm{SIZE}(U) + \boldsymbol{x}(i))^{\mathcal{O}(1)}.$$

Proof of Lemma 7.6. Let $D_1 := \max\{\|\boldsymbol{x}\|, D\} + D|Q|\,\|T\| + 1$. We call the first position where coordinate i has value at least D_1 its D_1-*crossing*. If either coordinate has no D_1-crossing, the first conclusion holds with $L := 1$. Otherwise, fix a partition witnessing D-alternation.

For $i \in \{1, 2\}$, let $r_i(\boldsymbol{u}_i)$ be the source of the segment containing the incoming transition of its D_1-crossing. This segment bounds coordinate $3 - i$ below D. If it is not the first segment, its predecessor bounds coordinate i below D by alternation. Otherwise, its source is the initial configuration. Hence $\boldsymbol{u}_i(i) \leq \max\{\|\boldsymbol{x}\|, D\}$, so Lemma 7.7 applies to the segment prefix ending at the crossing. Its construction gives a configuration $q_i(\boldsymbol{z}_i)$ before the D_1-crossing and an enabled cycle θ_i there with

$$\mathrm{EFF}(\theta_i)(i) > 0, \qquad \mathrm{EFF}(\theta_i)(3 - i) = 0.$$

The bound on $\boldsymbol{u}_i(i)$ gives a common length bound $L \leq (D + \mathrm{SIZE}(U) + \|\boldsymbol{x}\|)^{\mathcal{O}(1)}$.

The two crossings lie in distinct segments, since each segment keeps one coordinate below D. Exchange coordinates so that the D_1-crossing of coordinate 1 comes first. Then $q_1(\boldsymbol{z}_1)$ precedes $q_2(\boldsymbol{z}_2)$, giving (11). The first cut precedes both D_1-crossings, and the second precedes that of coordinate 2. Consequently, $\{1, 2\} \subseteq \mathrm{BD}(\rho_1, D_1)$ and $2 \in \mathrm{BD}(\rho_1\rho_2, D_1)$, as required. □

7.1.2 Drop Tests and Hybrid Set Representation

In the two-cycle case of Lemma 7.6, coordinate 2 remains bounded along $\rho_1\rho_2$. By recording this coordinate in states, we can describe the reachability set at q_2 by a finite union of hybrid sets. However, this alone does not allow us to obtain a strictly positive free period: a run reaching an arbitrary member of such a hybrid set need not visit q_1 with θ_1 enabled or have θ_2 enabled. We therefore construct an auxiliary 3-VASS $U^{(q_1,q_2)}$ that captures exactly the runs visiting q_1 before reaching q_2, with both cycles enabled at the corresponding configurations. At each state q_i, subtracting and then restoring $\mathrm{DROP}(\theta_i)$ checks that θ_i is enabled without changing the counter values.

Lemma 7.8. *Let $U = (Q, T)$ be a 3-VASS, $p(\boldsymbol{x}) \xrightarrow{\rho_1} q_1(\boldsymbol{z}_1) \xrightarrow{\rho_2} q_2(\boldsymbol{z}_2)$ be a run in U, and $\theta_0, \theta_1, \theta_2$ be three cycles enabled at $p(\boldsymbol{x})$, $q_1(\boldsymbol{z}_1)$, and $q_2(\boldsymbol{z}_2)$, respectively, with lengths at most L and effects*

$$\mathrm{EFF}(\theta_0) = (0, 0, a), \qquad \mathrm{EFF}(\theta_1) = (u, 0, b), \qquad \mathrm{EFF}(\theta_2) = (0, v, c),$$

where $a, u, v \geq 1$. There exists a 3-VASS $U^{(q_1,q_2)}$ of size polynomial in $\mathrm{SIZE}(U) + L$ with initial state $p^{(0)}$ and final state f, satisfying the following properties.

1. *$p^{(0)}(\boldsymbol{x}) \xrightarrow{\pi} f(\boldsymbol{z}_2)$ in $U^{(q_1,q_2)}$ for some run π with $\mathrm{BD}(\rho_1\rho_2, D) \subseteq \mathrm{BD}(\pi, D)$ for every $D \geq 0$.*
2. *There exists $\boldsymbol{a}_+ \geq \mathbf{1}$ with $\|\boldsymbol{a}_+\| \leq (L + \mathrm{SIZE}(U))^{\mathcal{O}(1)}$ such that*

$$\mathrm{Reach}^{f}_{U^{(q_1,q_2)}}(p^{(0)}(\boldsymbol{x})) + \boldsymbol{a}_+^* \subseteq \mathrm{Reach}^{q_2}_{U}(p(\boldsymbol{x})).$$

Moreover, if $f(\boldsymbol{h})$ is reachable in $U^{(q_1,q_2)}$ from $p^{(0)}(\boldsymbol{x})$, then for every $k \in \mathbb{N}$,

$$\ell_U\big(p(\boldsymbol{x}), q_2(\boldsymbol{h} + k \cdot \boldsymbol{a}_+)\big) \le \ell_{U^{(q_1,q_2)}}\big(p^{(0)}(\boldsymbol{x}), f(\boldsymbol{h})\big) + k(L + \mathrm{SIZE}(U))^{\mathcal{O}(1)}.$$

Proof. Let $\boldsymbol{a}_0 := \mathrm{EFF}(\theta_0)$, $\boldsymbol{e}_j := \mathrm{EFF}(\theta_j)$, and $\boldsymbol{d}_j := \mathrm{DROP}(\theta_j)$ for $j = 1, 2$. Then $\|\boldsymbol{a}_0\|, \|\boldsymbol{e}_j\|, \|\boldsymbol{d}_j\| \le L\|T\|$ and $\boldsymbol{d}_j \le \boldsymbol{z}_j$. Take two disjoint copies $U^{(0)}, U^{(1)}$ of U, with copied states $q^{(0)}, q^{(1)}$, and three fresh states $\hat{q}_1, \hat{q}_2, f$. Besides the copied transitions, add only the following transitions $t_j^{\pm}$, of effects $\pm\boldsymbol{d}_j$, respectively, for $j = 1, 2$:

$$q_1^{(0)} \xrightarrow{t_1^-} \hat{q}_1 \xrightarrow{t_1^+} q_1^{(1)}, \qquad q_2^{(1)} \xrightarrow{t_2^-} \hat{q}_2 \xrightarrow{t_2^+} f.$$

Denote the resulting system by $U^{(q_1,q_2)}$. It has $2|Q| + 3$ states, $2|T| + 4$ transitions, and transition norm at most $L\|T\|$, so its size is polynomial in $\mathrm{SIZE}(U) + L$. Each test decreases and then restores the vector, and is enabled precisely when its source vector is at least $\boldsymbol{d}_j$. Thus $\rho_1\rho_2$ lifts to a run π reaching $f(\boldsymbol{z}_2)$ with the same coordinatewise maxima, proving the first property.

Every run $p^{(0)}(\boldsymbol{x}) \xrightarrow{\gamma} f(\boldsymbol{h})$ traverses the two tests in order. Deleting their four transitions gives a run in U of the form

$$p(\boldsymbol{x}) \xrightarrow{\eta_1} q_1(\boldsymbol{w}_1) \xrightarrow{\eta_2} q_2(\boldsymbol{h}), \qquad \boldsymbol{w}_1 \ge \boldsymbol{d}_1, \quad \boldsymbol{h} \ge \boldsymbol{d}_2.$$

Let

$$\lambda := 2L\|T\| + 1, \qquad \boldsymbol{a}_1 := \lambda \cdot \boldsymbol{a}_0 + \boldsymbol{e}_1, \qquad \boldsymbol{a}_+ := \boldsymbol{a}_1 + \boldsymbol{e}_2.$$

The prescribed effects and $|b|, |c| \le L\|T\|$ give $\boldsymbol{a}_1 \ge \boldsymbol{0}$, $\boldsymbol{a}_+ \ge \boldsymbol{1}$, and

$$\|\boldsymbol{a}_+\| \le (2L\|T\| + 3)L\|T\| \le (L + \mathrm{SIZE}(U))^{\mathcal{O}(1)}.$$

For $k \in \mathbb{N}$, consider $\theta_0^{k\lambda}\eta_1\theta_1^k\eta_2\theta_2^k$. The copies of θ_0 and the translated η_1, η_2 are enabled by monotonicity. Before the $(j+1)$-st copies of θ_1 and θ_2, respectively, where $0 \le j < k$, we have

$$\begin{aligned} \boldsymbol{w}_1 + k\lambda \cdot \boldsymbol{a}_0 + j \cdot \boldsymbol{e}_1 &= \boldsymbol{w}_1 + (k-j)\lambda \cdot \boldsymbol{a}_0 + j \cdot \boldsymbol{a}_1 \ge \boldsymbol{d}_1, \\ \boldsymbol{h} + k \cdot \boldsymbol{a}_1 + j \cdot \boldsymbol{e}_2 &= \boldsymbol{h} + (k-j) \cdot \boldsymbol{a}_1 + j \cdot \boldsymbol{a}_+ \ge \boldsymbol{d}_2. \end{aligned}$$

Hence every cycle copy is enabled, and the run ends at $q_2(\boldsymbol{h} + k \cdot \boldsymbol{a}_+)$. Its length is at most

$$|\gamma| + k\big(\lambda|\theta_0| + |\theta_1| + |\theta_2|\big) \le |\gamma| + k(2L\|T\| + 3)L.$$

Taking γ shortest proves the distance bound and the reachability inclusion. □

The two properties in Lemma 7.8 serve different purposes. The first preserves the boundedness of the original run, allowing us to record a bounded coordinate in states without losing $\boldsymbol{z}_2$ from the reachability set. The second property applies to every vector reachable at f in the auxiliary VASS: adding any multiple of $\boldsymbol{a}_+$ yields a vector reachable at q_2 in U, with a polynomial bound on the additional run length. Thus, after obtaining a hybrid set G containing $\boldsymbol{z}_2$ from the auxiliary reachability set, we can add $\boldsymbol{a}_+$ as a free period.

Together with Theorem 6.3 and Lemma 6.4, this yields a hybrid set with polynomially bounded representation norm and short prefix witnesses. The complete proof is given in Appendix E.4.

Lemma 7.9. *Let $p(\boldsymbol{x}) \xrightarrow{\pi} t = q(\boldsymbol{y})$ be a run in a sequential 3-VASS $U = (Q, T)$, with $p(\boldsymbol{x})$ pumpable, and let $C \in \mathbb{N}$. There exist a partition of π*

$$p(\boldsymbol{x}) \xrightarrow{\tau} r(\boldsymbol{z}) \xrightarrow{\sigma} t,$$

and a hybrid set H containing $\boldsymbol{z}$, with $\textsc{norm}(H) \leq (C + \textsc{size}(U, p(\boldsymbol{x})))^{\mathcal{O}(1)}$, such that every $r(\boldsymbol{h})$ with $\boldsymbol{h} \in H$ has a prefix witness $p(\boldsymbol{x}) \xrightarrow{\rho} r(\boldsymbol{h})$ of length

$$|\rho| \leq (C + \textsc{size}(U, p(\boldsymbol{x}), r(\boldsymbol{h})))^{\mathcal{O}(1)}.$$

Either σ is empty, or H is forward C-coverable in the component containing r.

7.1.3 Joining the Backward Coverable Target

Now we can show that the reachability distance from a pumpable source to a backward C-coverable target in $[\![V]\!]_B$ is polynomially bounded in V.

Proof of Lemma 7.3. Let $M := C + B + \textsc{size}(V, \boldsymbol{s}, \boldsymbol{t})$. Since $\textsc{size}(U) \leq (B+1)^2 \textsc{size}(V)$, applying Lemma 7.9 to a witnessing run $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{\pi} \boldsymbol{q}(\boldsymbol{y})$ gives a partition $\pi = \tau\sigma$ at $\boldsymbol{r}(\boldsymbol{z})$ and a hybrid set H containing $\boldsymbol{z}$, with

$$\textsc{norm}(H) \leq M^{\mathcal{O}(1)}, \qquad \ell_U\big(\boldsymbol{p}(\boldsymbol{x}), \boldsymbol{r}(\boldsymbol{h})\big) \leq (M + \|\boldsymbol{h}\|)^{\mathcal{O}(1)} \quad \text{for all } \boldsymbol{h} \in H. \tag{12}$$

If $\sigma = \varepsilon$, take $\boldsymbol{h} := \boldsymbol{y} \in H$ in (12) and lift the resulting run to V, obtaining the required bound.

Otherwise, let U' be the suffix of U starting with the component containing $\boldsymbol{r}$, including the original bridges. By Lemma 7.9, H is forward C-coverable in this component and hence in U'. The backward covering run at $\boldsymbol{q}(\boldsymbol{y})$ stays in its component, so $\{\boldsymbol{y}\}$ is backward C-coverable in U'. Since $\boldsymbol{z} \in H$ and $\boldsymbol{r}(\boldsymbol{z}) \xrightarrow{\sigma} \boldsymbol{q}(\boldsymbol{y})$ in U', Lemma 6.17 applied to H and $\{\boldsymbol{y}\}$ gives $\boldsymbol{h} \in H$ with

$$\|\boldsymbol{h}\|, \ \ell_V\big((\boldsymbol{r}, \boldsymbol{h}), \boldsymbol{t}\big) \leq M^{\mathcal{O}(1)}.$$

By (12), we have

$$\ell_V(\boldsymbol{s}, \boldsymbol{t}) \leq \ell_U\big(\boldsymbol{p}(\boldsymbol{x}), \boldsymbol{r}(\boldsymbol{h})\big) + \ell_V\big((\boldsymbol{r}, \boldsymbol{h}), \boldsymbol{t}\big) \leq M^{\mathcal{O}(1)},$$

which is the bound in (8). □

Hybrid Sources with a Nonzero Free Period. The preceding result assumes a fixed pumpable source, whereas the R-large case may yield a hybrid set representation. If this set has a nonzero free period, we can adjoin its periods as self-loops to make the base pumpable and then apply Lemma 7.3. By removing the added transitions and handling the restricted-period costs as in Lemma 6.17, we obtain the following result. The proof can be found in Appendix E.5.

Lemma 7.10. *Let U be a sequential sub-VASS of $[\![V]\!]_B$, let C be a transformation-area threshold for $[\![V]\!]_B$, and let H be a hybrid set at state $\boldsymbol{r}$ with a nonzero free period. Suppose that some member of H reaches a backward C-coverable target $\boldsymbol{q}(\boldsymbol{y})$ in U. Then there exists $\boldsymbol{h} \in H$ such that*

$$\|\boldsymbol{h}\|, \ell_V\big((\boldsymbol{r}, \boldsymbol{h}), \boldsymbol{t}\big) \leq \big(B + C + \textsc{size}(V, \boldsymbol{t}) + \textsc{norm}(H)\big)^{\mathcal{O}(1)}, \text{ where } \boldsymbol{t} = (\boldsymbol{q}, \boldsymbol{y}).$$

7.2 Translated Families

The preceding results handle hybrid sources with a nonzero free period. We now consider another terminal case, based on hybrid sets whose bases vary along a common direction while their periods and cost constraints remain unchanged.

Definition 7.11 (Translated Family)**.** For a permutation σ of $[3]$, a *translated family* $\mathcal{H} = \{H_r\}_{r=0}^{\kappa}$ consists of hybrid sets $H_r = \boldsymbol{b} + r \cdot \boldsymbol{\delta} + P^{\langle \boldsymbol{c}, \cdot \rangle \le K}$, where $\boldsymbol{\delta} \in \mathbb{Z}^3$, $\boldsymbol{\delta}(\sigma(3)) = 0$, $\boldsymbol{p}(\sigma(3)) = 0$ for every $\boldsymbol{p} \in P$, and $\boldsymbol{b} + r \cdot \boldsymbol{\delta} \ge \mathbf{0}$ for each integer $r \in [0, \kappa]$.

We next consider reachability between two translated families. The key is to combine a known run between their unshifted members with C-coverability at suitable translation indices. Recall that C is the fixed transformation-area threshold for $[\![V]\!]_B$, and D is the polynomial threshold fixed in (7). Let $\mathcal{H} = \{H_r\}_{r=0}^{\kappa}$ and $\mathcal{J} = \{J_r\}_{r=0}^{\kappa}$ be two translated families, where

$$H_r = \boldsymbol{b} + r \cdot \boldsymbol{\delta} + P^{\langle \boldsymbol{c}, \cdot \rangle \le K}, \qquad J_r = \boldsymbol{b}' + r \cdot \boldsymbol{\delta}' + P'^{\langle \boldsymbol{c}', \cdot \rangle \le K'}.$$

Let $\boldsymbol{p}$ and $\boldsymbol{q}$ be two states of a sequential 3-VASS U with $\boldsymbol{p}$ in the first SCC and $\boldsymbol{q}$ in the last SCC. We call $(U, \boldsymbol{p}(\mathcal{H}), \boldsymbol{q}(\mathcal{J}))$ *C-joinable* if the following conditions hold:

1. There exist $\boldsymbol{x} \in H_0$ and $\boldsymbol{x}' \in J_0$ such that $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{*} \boldsymbol{q}(\boldsymbol{x}')$ in U.
2. There exist integers $0 \le r, r' \le \min\{D, \kappa\}$ such that H_r is forward C-coverable and $J_{r'}$ is backward C-coverable in U.

The following lemma shows that C-joinability allows us to select members of the two translated families connected by a short run in the original VAS.

Lemma 7.12. *Let $\mathcal{H} = \{H_r\}_{r=0}^{\kappa}$ and $\mathcal{J} = \{J_r\}_{r=0}^{\kappa}$ be two translated families at states $\boldsymbol{p}$ and $\boldsymbol{q}$ of U, where*

$$H_r = \boldsymbol{b} + r \cdot \boldsymbol{\delta} + P^{\langle \boldsymbol{c}, \cdot \rangle \le K}, \qquad J_\ell = \boldsymbol{b}' + \ell \cdot \boldsymbol{\delta}' + P'^{\langle \boldsymbol{c}', \cdot \rangle \le K'}.$$

Suppose that $(U, \boldsymbol{p}(\mathcal{H}), \boldsymbol{q}(\mathcal{J}))$ is C-joinable. There exists $\kappa_0 \le \left(D + \|P\| + \|P'\|\right)^{\mathcal{O}(1)}$ such that, if $\kappa \ge \kappa_0$, some $0 \le r, \ell \le \kappa_0$ and $\boldsymbol{h} \in H_r$, $\boldsymbol{h}' \in J_\ell$ satisfy

$$\|\boldsymbol{h}\|, \quad \|\boldsymbol{h}'\|, \quad \ell_V\big((\boldsymbol{p}, \boldsymbol{h}), (\boldsymbol{q}, \boldsymbol{h}')\big) \le \left(D + \mathrm{NORM}(H_0) + \mathrm{NORM}(J_0) + \|\boldsymbol{\delta}\| + \|\boldsymbol{\delta}'\|\right)^{\mathcal{O}(1)}.$$

Proof. Let $\boldsymbol{z} \in H_0$, $\boldsymbol{z}' \in J_0$ and $a, a' \le \min\{D, \kappa\}$ witness C-joinability. Write $\boldsymbol{b}_s := \boldsymbol{b} + s \cdot \boldsymbol{\delta}$ and $\boldsymbol{b}'_s := \boldsymbol{b}' + s \cdot \boldsymbol{\delta}'$ for $0 \le s \le \kappa$, and let P_0, P'_0 be the free periods. Let $M := D + \mathrm{NORM}(H_0) + \mathrm{NORM}(J_0) + \|\boldsymbol{\delta}\| + \|\boldsymbol{\delta}'\|$ and $n := |Q|$. Set $\Delta_* := \max\{1, \|T\|, \|P\|, \|P'\|\}$ and $L_{\mathrm{cov}} := \max\{2, n, \Delta_*, C\}^{c_{\mathrm{cov}}}$, where c_{cov} is the exponent for 3-VASS from Lemma 3.1. Choose

$$\kappa_0 := D + 6(n\Delta_*)^3 + C + L_{\mathrm{cov}}\Delta_* + 1 \le \left(D + \|P\| + \|P'\|\right)^{\mathcal{O}(1)}. \tag{13}$$

Fix representations of $\boldsymbol{z} \in H_0$ and $\boldsymbol{z}' \in J_0$. In the augmented VASS $W := U[P, P']$, we can first execute the added positive loops at $\boldsymbol{p}$ to reach $\boldsymbol{z}$, follow the original run ρ_0, and then execute the negative loops at $\boldsymbol{q}$ to reach $\boldsymbol{b}'$. This gives

$$\boldsymbol{p}(\boldsymbol{b}) \xrightarrow{\rho_P} \boldsymbol{p}(\boldsymbol{z}) \xrightarrow{\rho_0} \boldsymbol{q}(\boldsymbol{z}') \xrightarrow{\rho_{P'}} \boldsymbol{q}(\boldsymbol{b}'). \tag{14}$$

The last segment is enabled because each intermediate vector equals $\boldsymbol{b}'$ plus the periods not yet removed. Let $\widetilde{V}$ be obtained from V by adjoining the corresponding transitions $(0, 0, \boldsymbol{a})$ for $\boldsymbol{a} \in P$ and $(0, 0, -\boldsymbol{a}')$ for $\boldsymbol{a}' \in P'$. Both W and $\widetilde{V}$ have size polynomial in M.

If W is non-wide, choose $r = \ell = 0$ and apply Corollary 6.9 to obtain a run between the unshifted bases of length $M^{\mathcal{O}(1)}$. Suppose now that W is wide. Choose three linearly independent simple-cycle effects $\boldsymbol{d}_1, \boldsymbol{d}_2, \boldsymbol{d}_3$ in W, and let

$$\nu := \left|\det(\boldsymbol{d}_1, \boldsymbol{d}_2, \boldsymbol{d}_3)\right|, \qquad 1 \le \nu \le 6(n\Delta_*)^3.$$

The indices a, a' witness forward and backward C-coverability, respectively, but translating the bases by arbitrary indices may destroy their connection. We therefore round them up to multiples of ν, ensuring that the displacement between the new bases is an integer combination of cycle effects. Let

$$r := \nu \left\lceil \frac{a}{\nu} \right\rceil, \qquad \ell := \nu \left\lceil \frac{a'}{\nu} \right\rceil, \qquad 0 \le r, \ell \le D + \nu - 1 < \kappa_0 \le \kappa.$$

For every h with $\boldsymbol{\delta}(h) \ge 0$, we have $\boldsymbol{b}_r(h) \ge \boldsymbol{b}_a(h)$. If $\boldsymbol{\delta}(h) < 0$, the nonnegativity of $\boldsymbol{b}_\kappa$ gives

$$\boldsymbol{b}_r(h) \ge (\kappa - r)|\boldsymbol{\delta}(h)| \ge C + L_{\mathrm{cov}}\Delta_* + 1. \tag{15}$$

By Lemma 3.1, choose a covering run of length less than L_{cov} from $\boldsymbol{p}(\boldsymbol{b}_a)$ in $U[P_0]$, and replay it from $\boldsymbol{p}(\boldsymbol{b}_r)$. Coordinates with nonnegative translation follow values no smaller than before, and every other coordinate stays above C by (15). Thus $\boldsymbol{p}(\boldsymbol{b}_r)$ is forward C-coverable in W. The same argument for $\boldsymbol{\delta}'$ in the reversed system makes $\boldsymbol{q}(\boldsymbol{b}'_\ell)$ backward C-coverable in W.

Cramer's rule and the divisibility of r, ℓ by ν give integers t_1, t_2, t_3 such that

$$\ell \cdot \boldsymbol{\delta}' - r \cdot \boldsymbol{\delta} = \sum_{m=1}^{3} t_m \cdot \boldsymbol{d}_m. \tag{16}$$

Apply the cycle-insertion argument from Lemma 6.13 to the path in (14), using (16) to obtain a $\mathbb{Z}$-run from $\boldsymbol{p}(\boldsymbol{b}_r)$ to $\boldsymbol{q}(\boldsymbol{b}'_\ell)$ in W. By Corollary 4.7, C remains a transformation-area threshold for $[\![\widetilde{V}]\!]_B$. Hence Theorem 6.10 supplies a run between these bases in $\widetilde{V}$.

In both cases, write ρ for the resulting run and $|\rho| \le L \le M^{\mathcal{O}(1)}$ for its length bound. Removing the added period transitions and applying restricted-period absorption yields the following claim.

Claim 7.13. *There exist indices $0 \le r, \ell \le \kappa_0$ and members $\boldsymbol{h} \in H_r$, $\boldsymbol{h}' \in J_\ell$ whose norms and reachability distance $\ell_V((\boldsymbol{p}, \boldsymbol{h}), (\boldsymbol{q}, \boldsymbol{h}'))$ are at most $M^{\mathcal{O}(1)}$.*

The proof of the claim is in Appendix E.6. We are done by the claim. □

7.3 Completing the Induction on the Number of SCCs

We have established reachability-distance bounds for two terminal cases: hybrid sources with a nonzero free period and translated families satisfying the joining conditions. We now show that every R-large bounded-coordinate prefix yields one of these structures.

The R-Large Dichotomy. Let $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{\pi_0} \boldsymbol{r}(\boldsymbol{z})$ be a prefix with $i \in \mathrm{BD}(\pi_0, D)$. Recording coordinate i in states gives a 2-VASS of size at most $D\,\mathrm{SIZE}(U)$. The following lemma extracts either a hybrid set with a nonzero free period or a translated family with $\kappa+1$ members. Crucially, the period norms, costs, and translation direction are bounded polynomially in D, independently of κ. This allows us to choose κ before the growth threshold R.

Lemma 7.14. *For every integer $\kappa \geq 1$, there exists $D \leq R \leq (D+\kappa)^{\mathcal{O}(1)}$ such that the following holds. Let $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{\pi_0} \boldsymbol{r}(\boldsymbol{z})$ be an R-large prefix in U whose coordinate i stays below D, and let $\boldsymbol{s} = (\boldsymbol{p}, \boldsymbol{x})$. Choose $k \neq i$ with $\boldsymbol{z}(k) > \|\boldsymbol{x}\| + R$, and let j be the remaining coordinate. Then one of the following holds:*

1. *$\boldsymbol{z} \in H$ for a hybrid set H with a nonzero free period and $\mathrm{NORM}(H) \leq (D + \|\boldsymbol{x}\|)^{\mathcal{O}(1)}$.*
2. *There is a translated family $\{H_r\}_{r=0}^{\kappa}$ with*

$$H_r = \boldsymbol{b} + r \cdot \boldsymbol{\delta} + P^{\langle \boldsymbol{c}, \cdot \rangle \leq K}, \qquad \boldsymbol{\delta} = u \cdot \boldsymbol{e}_j - v \cdot \boldsymbol{e}_k,$$

such that $\boldsymbol{z} \in H_0$, $u, v \geq 1$, and

$$\|\boldsymbol{c}\|, \|P\|, \|\boldsymbol{\delta}\| \leq D^{\mathcal{O}(1)}, \quad \|\boldsymbol{b}\| \leq (D + \|\boldsymbol{x}\| + \kappa)^{\mathcal{O}(1)}.$$

Moreover, every member $\boldsymbol{h}$ of the resulting hybrid set or translated family satisfies

$$\ell_V(\boldsymbol{s}, (\boldsymbol{r}, \boldsymbol{h})) \leq \big(D + \|\boldsymbol{x}\| + \|\boldsymbol{h}\|\big)^{\mathcal{O}(1)}. \tag{17}$$

We will first choose an LPS Λ with polynomial parameters for the encoding induced by π_0. We recall the relevant 2-VASS results below. The proof is in Appendix E.7.

Lemma 7.15. *Let $U = (Q, T)$ be a 2-VASS with $S = \mathrm{SIZE}(U)$, and let p, q be states of U.*

- *There exists a finite family $\mathcal{L}$ of LPSs Λ satisfying $|\Lambda| \leq S^{\mathcal{O}(1)}$ and $\|\Lambda\| \leq \|T\|$ such that*

$$\mathrm{Reach}_U^q(p(\boldsymbol{x})) = \bigcup_{\Lambda \in \mathcal{L}} \{\boldsymbol{y} : (\boldsymbol{x}, \boldsymbol{y}) \in \mathrm{Reach}_\Lambda\} \qquad \text{for every } \boldsymbol{x} \in \mathbb{N}^2.$$

  Moreover, whenever $p(\boldsymbol{x}) \xrightarrow{*} q(\boldsymbol{y})$, $\ell_U(p(\boldsymbol{x}), q(\boldsymbol{y})) \leq \big(S + \|\boldsymbol{x}\| + \|\boldsymbol{y}\|\big)^{\mathcal{O}(1)}$.

- *Let $A = \boldsymbol{b}_A + P_A^{\langle \boldsymbol{c}_A, \cdot \rangle \leq K_A}$ be a hybrid set at p. There is a finite family $\mathcal{G}$ of hybrid sets at q such that $\mathrm{Reach}_U^q(p(A)) = \bigcup_{G \in \mathcal{G}} G$, and every $G = \boldsymbol{b}_G + P_G^{\langle \boldsymbol{c}_G, \cdot \rangle \leq K_G} \in \mathcal{G}$ satisfies*

$$\mathrm{NORM}(G) \leq \big(S + \mathrm{NORM}(A)\big)^{\mathcal{O}(1)}, \qquad \|P_G\|, \|\boldsymbol{c}_G\|, |P_G| \leq \big(S + \|P_A\| + \|\boldsymbol{c}_A\|\big)^{\mathcal{O}(1)}.$$

We now apply Lemma 7.15 to prove the R-large dichotomy.

Proof of Lemma 7.14. Encode coordinate i and apply Lemma 7.15. Let $L \leq D^{\mathcal{O}(1)}$ bound the lengths of all resulting LPS, and let $E := L\,\|T\|$. For any $\kappa \geq 1$, choose

$$R := \max\{D, \kappa E\} \leq (D+\kappa)^{\mathcal{O}(1)}.$$

Choose such an LPS $\Lambda = \alpha_0 \beta_1^* \alpha_1 \cdots \beta_m^* \alpha_m$ containing a witnessing run from $\boldsymbol{x}\restriction_{\{j,k\}}$ to $\boldsymbol{z}\restriction_{\{j,k\}}$, where $[3] = \{i, j, k\}$. Delete unused cycle blocks and zero-effect cycles, preserving the witnessing run.

Throughout the following claims, we identify each hybrid set at a fixed state of the 2-VASS with its lift to $\mathbb{N}^3$, obtained by inserting the recorded value of coordinate i into its base and zero into every period. Similarly, cycle effects are lifted with zero in coordinate i. For the chosen LPS Λ, let

$$\mathcal{S}_\Lambda := \left\{ \boldsymbol{h} \in \mathbb{N}^3 : \boldsymbol{h}(i) = \boldsymbol{z}(i),\ \boldsymbol{x} \restriction_{\{j,k\}} \xrightarrow{\Lambda}_{\mathbb{N}} \boldsymbol{h} \restriction_{\{j,k\}} \right\}.$$

The first claim handles the free-period case.

Claim 7.16. *If a cycle used by the chosen run has nonzero nonnegative effect, there exists a hybrid set H with a nonzero free period such that $\boldsymbol{z} \in H \subseteq \mathcal{S}_\Lambda$ and $\mathrm{NORM}(H) \le (D + \|\boldsymbol{x}\|)^{\mathcal{O}(1)}$.*

For the remaining case, assume that no used cycle has nonzero nonnegative effect. Every used cycle increasing coordinate k must then decrease coordinate j. The large growth in coordinate k forces such a cycle to occur at least κ times. Choosing the last cycle block with this property ensures that the remaining suffix has enough reserve in coordinate k to delete up to κ repetitions, even at intermediate positions within cycles. Hence, we have the following claim.

Claim 7.17. *There exists a run $\eta_- \beta^\kappa \eta_+$ from the encoded source to $\boldsymbol{z}$, with vector $\boldsymbol{a}$ immediately before β^κ, satisfying the following properties:*

1. *β is a cycle with $\mathrm{EFF}(\beta) = -u \cdot \boldsymbol{e}_j + v \cdot \boldsymbol{e}_k$, where $1 \le u, v \le E$.*
2. *The transition-effect sequences of η_- and η_+ are captured by LPSs Λ_- and Λ_+, respectively, obtained by cutting Λ after the selected cycle β, with that cycle included in Λ_-.*
3. *Every vector $\boldsymbol{w}$ along η_+ satisfies $\boldsymbol{w}(k) \ge \kappa v$.*

Fix the run $\eta_- \beta^\kappa \eta_+$ and the LPSs Λ_-, Λ_+ supplied by Claim 7.17. Deleting copies of β shifts the final vector along $\boldsymbol{\delta} = -\mathrm{EFF}(\beta) = u \cdot \boldsymbol{e}_j - v \cdot \boldsymbol{e}_k$. The lower bound along η_+ ensures that these deletions preserve nonnegativity. Using hybrid descriptions of Λ_- and Λ_+, we obtain the following translated family. The full proofs are given in Appendix E.8.

Claim 7.18. *There is a translated family $\{H_s\}_{s=0}^{\kappa}$ with $\boldsymbol{z} \in H_0$ and direction $\boldsymbol{\delta} = u \cdot \boldsymbol{e}_j - v \cdot \boldsymbol{e}_k$. Its periods and costs have norm at most $D^{\mathcal{O}(1)}$, and its base has norm at most $(D + \|\boldsymbol{x}\| + \kappa)^{\mathcal{O}(1)}$. Every member is reachable from the encoded source at the fixed encoded target state.*

These claims establish the two outcomes of the dichotomy. In both cases, every member of the resulting hybrid set or translated family is reachable from the fixed source in the bounded-coordinate 2-VASS. The short-witness bound in Lemma 7.15, followed by lifting to V, gives (17) and completes the proof. □

Completion of the Induction. The two outcomes of Lemma 7.14 yield a polynomial distance bound for every R-large prefix satisfying the bounded-prefix stopping condition, leaving only R-small prefixes for the inductive step.

Lemma 7.19 (R-large: terminal case)**.** *There exists $D \le R \le D^{\mathcal{O}(1)}$ such that the following holds. Let $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{\pi_0} \boldsymbol{r}(\boldsymbol{z}) \xrightarrow{\pi_1} \boldsymbol{q}(\boldsymbol{y})$ be a run in U whose target $\boldsymbol{q}(\boldsymbol{y})$ is backward C-coverable, where $\boldsymbol{s} = (\boldsymbol{p}, \boldsymbol{x})$ and $\boldsymbol{t} = (\boldsymbol{q}, \boldsymbol{y})$. Assume that π_0 is R-large and $i \in \mathrm{BD}(\pi_0, D)$. If the configuration $\boldsymbol{r}(\boldsymbol{w})$ is forward C-coverable whenever $\boldsymbol{w}(i) \ge \boldsymbol{z}(i)$ and $\boldsymbol{w}(j) \ge D$ for all $j \ne i$, then*

$$\ell_V(\boldsymbol{s}, \boldsymbol{t}) \le \left(D + \|\boldsymbol{x}\| + \|\boldsymbol{t}\|\right)^{\mathcal{O}(1)}.$$

Proof sketch. The extracted period norms are polynomially bounded in D, independently of κ. We can therefore choose $\kappa \geq \max\{2D, \kappa_0\}$ given in Lemma 7.12, and then choose R by Lemma 7.14, both polynomially bounded in D. Apply the dichotomy. In the free-period case, invoke Lemma 7.10. In the translated-family case, the covering hypothesis and the choice of κ give C-joinability with the constant target family $J_\ell = \{\boldsymbol{y}\}$, so invoke Lemma 7.12. In either case, prepend the short prefix supplied by (17). The full proof is in Appendix E.9. □

Proof of Theorem 7.1. Fix V, B, C and $\boldsymbol{t} = (\boldsymbol{q}, \boldsymbol{y})$, and assume that $\boldsymbol{q}(\boldsymbol{y})$ is backward C-coverable. Choose D uniformly for all sequential subinstances by Lemma 7.2, and then choose κ, R as in the proof of Lemma 7.19, so that

$$B + C + \textsc{size}(V) \leq D \leq (B + C + \textsc{size}(V))^{\mathcal{O}(1)}, \qquad D \leq R \leq D^{\mathcal{O}(1)}. \tag{18}$$

For $k \geq 1$, we construct a nondecreasing bound $L_k(N)$ on $\ell_V(\boldsymbol{s}, \boldsymbol{t})$, uniform over reachable sequential subinstances with at most k components, source norm at most N, and backward C-coverable target $\boldsymbol{q}(\boldsymbol{y})$. Let $M := D + \|\boldsymbol{t}\|$, and choose a fixed nondecreasing polynomial P large enough for the bounds below. Its degree and the choices of D, R do not depend on k.

Base Case: $k = 1$. If $\boldsymbol{p}(\boldsymbol{x})$ is forward C-coverable, apply Corollary 6.14. Otherwise, Lemma 7.2 captures the entire run in a 2-VASS of size at most $D\,\textsc{size}(U) \leq D^{\mathcal{O}(1)}$, since there is no later component. Apply Lemma 7.15 and lift the resulting run to V. In both cases,

$$L_1(N) \leq P(M + N).$$

For $k \geq 1$, assume that L_k has been constructed. Consider a run from $\boldsymbol{p}(\boldsymbol{x})$ to $\boldsymbol{q}(\boldsymbol{y})$ in a sequential subinstance U with at most $k+1$ components and $\|\boldsymbol{s}\| \leq N$. We distinguish the terminal and inductive cases.

Terminal Cases. If $\boldsymbol{p}(\boldsymbol{x})$ is forward C-coverable, apply Corollary 6.14. Otherwise, write the run as $\pi_0\pi_1$ by Lemma 7.2. If $\pi_1 = \varepsilon$, apply Lemma 7.15 to the 2-VASS encoding of π_0 and lift the run to V. If $\pi_1 \neq \varepsilon$ and π_0 is R-large, its target satisfies the covering condition supplied by Lemma 7.2. By applying Lemma 7.19, the reachability-distance bound holds. Each case gives a run in V with

$$\ell_V(\boldsymbol{s}, \boldsymbol{t}) \leq P(M + N).$$

Inductive Case. The remaining case has an R-small prefix π_0 ending at a later component entry $\boldsymbol{r}(\boldsymbol{z})$. Let $\boldsymbol{u} := (\boldsymbol{r}, \boldsymbol{z})$. Since $R \geq D \geq B$,

$$\|\boldsymbol{u}\| \leq \max\{B, \|\boldsymbol{x}\| + R\} \leq \|\boldsymbol{s}\| + R \leq N + R. \tag{19}$$

By Lemma 7.15, replace π_0 by a run in V from $\boldsymbol{s}$ to $\boldsymbol{u}$ of length at most $P(M + N + R)$. The unchanged suffix has at most k components and preserves the target's backward covering run, which stays in its component. Apply the induction hypothesis to this suffix and prepend the shortened prefix. Together with the terminal-case bounds, this allows us to choose

$$L_{k+1}(N) \leq P(M + N + R) + L_k(N + R). \tag{20}$$

Note that $k \leq (B+1)^2$. By (18), solving this recurrence gives

$$\ell_V(\boldsymbol{s}, \boldsymbol{t}) \leq L_k(\|\boldsymbol{s}\|) \leq (B + C + \textsc{size}(V, \boldsymbol{s}, \boldsymbol{t}))^{\mathcal{O}(1)}.$$

The degree is independent of k, proving (6). □

8 Removing the Coverability Assumption

We now remove the coverability assumption from Theorem 7.1, using the same additive-growth induction. Pumpable sources and hybrid sets with a nonzero free period can be handled using the semi-coverable result. The remaining difficulty arises when an R-large prefix yields a translated family but the target is not backward C-coverable.

To handle this case, we reverse the remaining suffix and regard the extracted translated family as a fixed target family. We show that an arbitrary source can reach a suitable member of such a family by a polynomial-length run, using Lemma 7.12 when another translated family arises. This provides the missing terminal case for the final induction.

Throughout this section, fix a 5-VAS V, an encoding bound $B \in \mathbb{N}$, and a sequential sub-VASS U of $[\![V]\!]_B$. Let C be a transformation-area threshold for $[\![V]\!]_B$, and choose D uniformly as in (7).

8.1 Pumpable Sources and Hybrid Sets with a Nonzero Free Period

We first extend the pumpable-source bound to arbitrary targets. The extraction in Lemma 7.9 does not require backward coverability of the target. It produces a forward C-coverable hybrid set, unless the extraction process reaches the target, which already gives a short witness. We can therefore apply Theorem 7.1 after adjoining its periods as self-loops. The proof is given in Appendix F.1.

Lemma 8.1. *Suppose $\boldsymbol{p}(\boldsymbol{x})$ is pumpable in U and $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{*} \boldsymbol{q}(\boldsymbol{y})$ in U. For $\boldsymbol{s} = (\boldsymbol{p}, \boldsymbol{x})$ and $\boldsymbol{t} = (\boldsymbol{q}, \boldsymbol{y})$,*

$$\ell_V(\boldsymbol{s}, \boldsymbol{t}) \leq \big(D + \|\boldsymbol{s}\| + \|\boldsymbol{t}\|\big)^{\mathcal{O}(1)}.$$

The preceding lemma handles fixed pumpable sources without requiring target coverability. We now extend this result to hybrid sources and targets, allowing the source base to be either forward C-coverable or pumpable after adjoining its free periods.

Lemma 8.2. *Let $H = \boldsymbol{a} + P^{\langle \boldsymbol{c}, \cdot \rangle \leq K}$ and J be hybrid sets at the source state $\boldsymbol{p}$ and target state $\boldsymbol{q}$ of U, respectively. Write P_0 for the free periods of H. Suppose that $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{*} \boldsymbol{q}(\boldsymbol{y})$ in U for some $\boldsymbol{x} \in H$ and $\boldsymbol{y} \in J$. If $\boldsymbol{p}(\boldsymbol{a})$ is forward C-coverable or pumpable in $U[P_0]$, then there exist $\boldsymbol{h} \in H$ and $\boldsymbol{h}' \in J$ such that*

$$\|\boldsymbol{h}\|, \ \|\boldsymbol{h}'\|, \quad \ell_V\big((\boldsymbol{p}, \boldsymbol{h}), (\boldsymbol{q}, \boldsymbol{h}')\big) \leq \big(D + \mathrm{NORM}(H) + \mathrm{NORM}(J)\big)^{\mathcal{O}(1)}. \tag{21}$$

The proof adds the source periods and the negatives of the target periods as self-loops. The augmented source is forward C-coverable or pumpable, so Theorem 7.1 or the preceding pumpable-source bound gives a short run between the bases. Absorbing restricted periods yields the required members of the original hybrid sets. The details are given in Appendix F.2.

In particular, taking the target hybrid set to be a singleton handles the free-period outcome of the R-large dichotomy without any target coverability assumption. It remains to handle the translated-family outcome.

8.2 Reaching a Fixed Translated Family

We now extend the induction of Section 7 to a fixed translated-family target with a backward C-coverable member. The first ingredient is a small-intersection property for hybrid sets, with a bound independent of their cost bounds. Its proof is given in Appendix F.3.

Lemma 8.3. *Let $H, J \subseteq \mathbb{N}^3$ be hybrid sets with $H \cap J \neq \emptyset$. There is $\boldsymbol{h} \in H \cap J$ such that*

$$\|\boldsymbol{h}\| \leq \big(D + \textsc{norm}(H) + \textsc{norm}(J)\big)^{\mathcal{O}(1)}.$$

We now establish the distance bound for an arbitrary source and a fixed translated-family target. The induction follows Theorem 7.1, keeping the target family fixed and replacing its terminal cases by the preceding hybrid-set bounds.

Lemma 8.4. *Let $\mathcal{J} = \{J_\ell\}_{0 \leq \ell \leq \kappa}$ be a translated family at the target state $\boldsymbol{r}'$ of U with*

$$J_\ell = \boldsymbol{b}' + \ell \cdot \boldsymbol{\delta}' + P'^{\langle \boldsymbol{c}', \cdot \rangle \leq K'}.$$

Suppose that J_a is backward C-coverable in U for some $0 \leq a \leq \min\{D, \kappa\}$. Suppose that $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{*} \boldsymbol{r}'(\boldsymbol{z}')$ in U for some $\boldsymbol{z}' \in J_0$. There exists $\kappa_0 \leq \big(D + \|P'\|\big)^{\mathcal{O}(1)}$ such that, if $\kappa \geq \kappa_0$, then some $0 \leq \ell \leq \kappa_0$ and $\boldsymbol{h}' \in J_\ell$ satisfy

$$\|\boldsymbol{h}'\|, \quad \ell_V(\boldsymbol{s}, (\boldsymbol{r}', \boldsymbol{h}')) \leq \big(D + \|\boldsymbol{s}\| + \textsc{norm}(J_0) + \|\boldsymbol{\delta}'\|\big)^{\mathcal{O}(1)}, \qquad \boldsymbol{s} = (\boldsymbol{p}, \boldsymbol{x}). \tag{22}$$

Proof sketch. Follow the induction on the number of SCCs in Theorem 7.1, keeping $\mathcal{J}$ fixed. There are four terminal cases. A forward C-coverable source is handled by Lemma 8.2. A bounded-coordinate prefix reaching J_0 is handled by Lemmas 7.15 and 8.3. An R-large prefix yielding a nonzero free period is handled by Lemma 8.2. An R-large prefix yielding a translated family is handled by Lemma 7.12. Only an R-small prefix ending at a later component requires induction, with an additive increase of at most R in the source norm. The full case analysis is in Appendix F.4. □

8.3 The Final Induction

We are now ready to complete the proof of Proposition 6.5.

Proof sketch of Proposition 6.5. Use additive-growth induction as in Theorem 7.1. A forward C-coverable source is terminal by Theorem 7.1. Otherwise, extract a bounded-coordinate prefix. If it reaches the target, apply Lemma 7.15. If it is R-large, use Lemma 8.2 for a nonzero free period, or reverse the suffix and use Lemma 8.4 for a translated family. Only an R-small prefix ending at a later component requires induction, increasing the source norm by at most R. The full proof is in Appendix F.5. □

AI Methodology. The main mathematical ideas and proof strategies were developed by the authors. OpenAI's ChatGPT was used to assist in checking and refining detailed technical arguments, particularly in Section 8 and Appendix F. It was also used to prepare, audit, and revise the manuscript. The authors have verified all content, including references, and take full responsibility for the manuscript.

A Proofs for Main Results

A.1 Proof of Theorem 1.1

Theorem 1.1. *The reachability problem for* 5*-VAS is* PSPACE*-complete under binary encoding and* NL*-complete under unary encoding.*

Recall that the size $\mathrm{SIZE}(V, \boldsymbol{s}, \boldsymbol{t})$ is defined under unary encoding as in Section 3. For d-VAS, we use the following measure, which is polynomially equivalent to the binary encoding length:

$$\mathrm{SIZE}_{\mathrm{BIN}}(V, \boldsymbol{s}, \boldsymbol{t}) := 1 + d|T|\lceil \log_2(\|T\| + 2)\rceil + d\lceil \log_2(\|\boldsymbol{s}\| + 2)\rceil + d\lceil \log_2(\|\boldsymbol{t}\| + 2)\rceil.$$

The main theorem follows by a nondeterministic brute-force algorithm as follows.

Proof of Theorem 1.1. The NL-hardness and the PSPACE-hardness are given in [CFZ26, Theorem 1.2] and [KL27, Theorem 1], respectively. We only focus on the upper bound.

Let $c \geq 1$ be an integer constant satisfying Theorem 1.2. Construct a nondeterministic Turing machine that, given a 5-VAS $(V, \boldsymbol{s}, \boldsymbol{t})$ of size $M := \mathrm{SIZE}(V, \boldsymbol{s}, \boldsymbol{t})$, guesses at most $\ell := M^c$ transitions and verifies that they induce a run from $\boldsymbol{s}$ to $\boldsymbol{t}$. By this theorem, if $\boldsymbol{s} \xrightarrow{*} \boldsymbol{t}$, then there exists a witness of length at most ℓ. Hence the machine accepts if and only if $\boldsymbol{s} \xrightarrow{*} \boldsymbol{t}$ in V.

Then we analyze the space complexity of this algorithm. The machine needs to store the current configuration, which has a norm at most $\|\boldsymbol{s}\| + \ell \cdot M \leq (\ell + 1) \cdot M \leq 2M^{c+1}$. The five counters, a step counter, and an index of the current transition therefore require $O(\log M)$ space. The transition can be read from the input using that index. Also ℓ can be computed within the same space. For unary encoding, this uses logarithmic space, and hence reachability in unary 5-VAS is in NL. For binary encoding, note that $\log M = O(\mathrm{SIZE}_{\mathrm{BIN}}(V, \boldsymbol{s}, \boldsymbol{t}))$. So the same algorithm uses polynomial space, and hence reachability in binary 5-VAS is in NPSPACE = PSPACE. □

A.2 Proof of Corollary 1.4

Corollary 1.4. *For every* 5*-VAS V, its reachability relation can be represented by a finite set of mixed LPS, with all parameters polynomially bounded in* $\mathrm{SIZE}(V)$.

Proof of Corollary 1.4. Fix a polynomial factorization of Reach_V given by Theorem 1.3. Let R be a factor captured by $[\![V]\!]_{I,B}$ where $|I| = 3$ and B is polynomial in $\mathrm{SIZE}(V)$. Applying [BEF+21, Theorem 3.1] to this 2-VASS and restoring the original transition effects of V gives a finite family $\mathcal{S}_R$ of $\mathbb{N}$-LPSs satisfying

$$R \subseteq \bigcup_{\Lambda \in \mathcal{S}_R} \mathrm{Reach}_\Lambda \subseteq \mathrm{Reach}_V .$$

The second holds because the resulting $\mathbb{N}$-runs use only transitions of V. We still have $|\Lambda| \leq \mathrm{poly}(\mathrm{SIZE}(V))$ and $\|\Lambda\| \leq \mathrm{SIZE}(V)$.

Let $\mathcal{S}$ be the finite family of concatenations obtained by choosing an LPS from $\mathcal{S}_R$ for each bounded-counter factor and keeping the mixed-LPS factors. These concatenations cover Reach_V by the original factorization. Conversely, each chosen mixed LPS captures a subset of Reach_V, so transitivity gives $\mathrm{Reach}_V = \bigcup_{\Lambda \in \mathcal{S}} \mathrm{Reach}_\Lambda$. Since the number of factors is polynomially bounded, so are the parameters of the resulting mixed LPSs. □

A.3 Proof of Theorem 1.5

Theorem 1.5. *Reachability in $(d+3)$-VAS is in F_d for every fixed $d \geq 3$, even under binary encoding.*

We use F_d for the fast-growing functions and F_d for the associated complexity classes. See [Sch16, Section 2] for more details.

Proof of Theorem 1.5. Let n be the bit-length of the input. The unary size is at most $2^{\mathsf{poly}(n)}$, so Lemma A.1 gives a run of length at most $F_d(E(n))$ for some elementary function E whenever the target is reachable. Exhaustive search in V up to this length decides reachability. For fixed $d \geq 3$, the preprocessing and search overhead are absorbed in F_d [Sch16, Lemma 4.6 and Section 4.2.1]. □

Lemma A.1. *For every fixed $d \geq 3$, there is an elementary function g_d such that every reachable $(d+3)$-VAS instance $(V, \boldsymbol{s}, \boldsymbol{t})$ satisfies*

$$\ell_V(\boldsymbol{s}, \boldsymbol{t}) \leq F_d\big(g_d(\textsc{size}(V, \boldsymbol{s}, \boldsymbol{t}))\big).$$

Proof. Put $D := d+3$ and $M := \textsc{size}(V, \boldsymbol{s}, \boldsymbol{t})$. Apply Theorem 1.3 and choose a witnessing composition for $\boldsymbol{s} \xrightarrow{*} \boldsymbol{t}$ with polynomially bounded factor count and parameters. We build a VASS W using the original D counters and dedicated auxiliary counters that start and finish at zero.

Claim A.2. *Let Λ be a $\mathbb{Z}$-LPS over m counters. There exists an $(m+1)$-dimensional $\mathbb{N}$-LPS $\widehat{\Lambda}$ such that*

$$\mathrm{Reach}_{(\Lambda,\mathbb{Z})} = \left\{(\boldsymbol{s}, \boldsymbol{t}) : ((\boldsymbol{s}, 0), (\boldsymbol{t}, 0)) \in \mathrm{Reach}_{\widehat{\Lambda}}\right\}.$$

Moreover, $|\widehat{\Lambda}| \leq |\Lambda| + 2$, $\left\|\widehat{\Lambda}\right\| \leq \max\{\|\Lambda\|, 1\}$.

Proof of Claim. We append a zero auxiliary coordinate to every vector of Λ to obtain an $(m+1)$-dimensional $\mathbb{Z}$-LPS $\Lambda^{\uparrow}$, and define

$$\widehat{\Lambda} := (\mathbf{1}, 1)^* \Lambda^{\uparrow} (-\mathbf{1}, -1)^*.$$

Every $\mathbb{Z}$-witness of Λ with nonnegative endpoints can be executed under $\mathbb{N}$-semantics by repeating the positive vector sufficiently many times before it and the negative vector equally many times afterward. Conversely, a witness of $\widehat{\Lambda}$ whose auxiliary coordinate starts and ends at zero uses these two vectors equally often. Removing their occurrences gives a $\mathbb{Z}$-witness of Λ. The parameter bounds follow directly from the construction. ■

For a mixed LPS Λ, discard empty components and apply Claim A.2 to each $\mathbb{Z}$-component, using distinct auxiliary counters that are unchanged elsewhere and start and finish at zero. Write the resulting $\mathbb{N}$-LPS as $\widehat{\Lambda} = \alpha_0 \beta_1^* \alpha_1 \cdots \beta_k^* \alpha_k$. We construct a VASS $V_{\widehat{\Lambda}}$ by realizing each α_i as a directed path and each β_i as a simple cycle, using fresh states for distinct occurrences. Connect these paths and cycles in their order in $\widehat{\Lambda}$ by forward zero-effect transitions, and let $p_{\widehat{\Lambda}}$ and $q_{\widehat{\Lambda}}$ be the entry and exit states. Every run from $p_{\widehat{\Lambda}}$ to $q_{\widehat{\Lambda}}$ follows the fixed paths and repeats each cycle an arbitrary number of times. Consequently, for all $\boldsymbol{s}, \boldsymbol{t} \in \mathbb{N}^D$,

$$(\boldsymbol{s}, \boldsymbol{t}) \in \mathrm{Reach}_{\Lambda} \iff p_{\widehat{\Lambda}}(\boldsymbol{s}, \mathbf{0}) \xrightarrow{*} q_{\widehat{\Lambda}}(\boldsymbol{t}, \mathbf{0}) \text{ in } V_{\widehat{\Lambda}}.$$

Deleting the connecting transitions recovers an LPS witness of no greater length. Moreover, $V_{\widehat{\Lambda}}$ has $O(|\Lambda| + 1)$ states and transitions, with transition norm at most $\max\{\|\Lambda\|, 1\}$. Each SCC contains at most one simple cycle, so its cycle-dimension is at most one.

For a bounded-counter factor captured by $[\![V]\!]_{I,B}$ with $|I| = 3$, we restore all D original transition coordinates, and add a private entry and exit connected to the encoded states by zero-effect edges. This gadget has polynomial size and cycle-dimension at most d, since every cycle has zero effect on the three encoded coordinates. Its runs become runs of V after deleting the entry and exit edges, and it admits every witness of the factor.

Give the factor gadgets disjoint states and auxiliary counters, and connect them in sequence by forward zero-effect edges to obtain W. Let $s' := p_{\mathrm{in}}(\boldsymbol{s}, \mathbf{0})$ and $t' := p_{\mathrm{out}}(\boldsymbol{t}, \mathbf{0})$ be the global entry and exit configurations. The chosen witnesses give $s' \xrightarrow{*} t'$ in W, with $\textsc{size}(W, s', t') \leq \mathsf{poly}(M)$ and $\dim_{\mathrm{com}}(W) \leq d$, since every SCC stays inside one gadget. Every run from s' to t' restores each auxiliary counter to zero. Deleting the padding loops and connecting edges therefore recovers mixed-LPS witnesses and runs of V for the bounded-counter gadgets. By Theorem 1.3, these yield a run of V of no greater length. Thus [CFZ26, Theorem 2.1] gives

$$\ell_V(\boldsymbol{s}, \boldsymbol{t}) \leq \ell_W(s', t') \leq F_d\big(g(\textsc{size}(W, s', t'))\big),$$

where g is a fixed nondecreasing elementary function, independent of the number of counters. Composing g with the polynomial size bound gives g_d. □

B Quantifying van Leeuwen's Transformation Area

Fix a d-VAS V with transition set T. A set $A \subseteq \mathbb{N}^d$ is a transformation area if every $\mathbb{Z}$-run between configurations in A can be rearranged into a run in V with the same Parikh image. For $i \in [d]$ and $C \in \mathbb{N}$, define

$$\mathcal{T}_C^i := \{\boldsymbol{x} \in \mathbb{N}^d : \boldsymbol{x}(j) \geq C \text{ for every } j \in [d] \setminus \{i\}\}, \qquad \mathcal{T}_C := \bigcup_{i \in [d]} \mathcal{T}_C^i.$$

The original result of van Leeuwen [vL74, Theorem 5.3] states that for every $i \in [d]$, there exists an effectively computable threshold $C \in \mathbb{N}$ such that $\mathcal{T}_C^i$ is a transformation area. We give an explicit threshold linear in d and prove that the union $\mathcal{T}_C$ is also a transformation area.

Theorem B.1. *Let $C := 2d\,\|T\|\,(\|T\| + 1)$. Then $\mathcal{T}_C$ is a transformation area of V.*

We first show how to rearrange a $\mathbb{Z}$-run whose endpoints lie in the same $\mathcal{T}_C^i$ into a valid run.

Lemma B.2. *For every $i \in [d]$, the set $\mathcal{T}_{C_0}^i$ is a transformation area of V, where $C_0 := 2d\,\|T\|$.*

Proof. Fix $i \in [d]$ and a $\mathbb{Z}$-run $\boldsymbol{s} \xrightarrow{\pi}_{\mathbb{Z}} \boldsymbol{t}$ with $\boldsymbol{s}, \boldsymbol{t} \in \mathcal{T}_{C_0}^i$. Write $\pi = \boldsymbol{a}_1 \cdots \boldsymbol{a}_n$ and $\boldsymbol{z} := \boldsymbol{t} - \boldsymbol{s} = \sum_{\ell=1}^n \boldsymbol{a}_\ell$. The case $n = 0$ is immediate. As in the proof of Corollary 4.4, append n copies of $-\boldsymbol{z}/n$ and apply Lemma 4.3. Let $\boldsymbol{P}_0, \ldots, \boldsymbol{P}_{2n}$ be the partial sums in the resulting order, so that

$$\boldsymbol{P}_0 = \boldsymbol{P}_{2n} = \mathbf{0}, \qquad \|\boldsymbol{P}_k\| \leq d\,\|T\|\,.$$

Choose $r \in [0, 2n-1]$ minimizing $\boldsymbol{P}_r(i)$, and rotate the sequence to begin just after $\boldsymbol{P}_r$. Since the total sum is zero, every partial sum $\boldsymbol{R}$ of the rotated sequence has the form $\boldsymbol{P}_k - \boldsymbol{P}_r$, including those after the sequence wraps around. Thus

$$\boldsymbol{R}(i) \geq 0, \qquad \boldsymbol{R}(j) \geq -2d\,\|T\| \text{ for } j \neq i. \tag{23}$$

Delete the auxiliary vectors from the rotated sequence, and let ρ be the resulting path. For a nonempty prefix of ρ, let $\boldsymbol{p}$ be its effect and let $\boldsymbol{R}$ be the partial sum of the rotated augmented sequence ending at the same original occurrence. If that augmented prefix contains m auxiliary vectors, then

$$\boldsymbol{s} + \boldsymbol{p} = \boldsymbol{s} + \boldsymbol{R} + (m/n) \cdot \boldsymbol{z} \geq \boldsymbol{R} + \boldsymbol{s} + \min(\boldsymbol{0}, \boldsymbol{t} - \boldsymbol{s}) = \boldsymbol{R} + \min(\boldsymbol{s}, \boldsymbol{t}).$$

Using (23), we obtain $\boldsymbol{s}(i) + \boldsymbol{p}(i) \geq \boldsymbol{R}(i) + \min(\boldsymbol{s}(i), \boldsymbol{t}(i)) \geq 0$. For every $j \neq i$, the endpoints instead satisfy $\min(\boldsymbol{s}(j), \boldsymbol{t}(j)) \geq C_0$, so $\boldsymbol{s}(j) + \boldsymbol{p}(j) \geq -2d \|T\| + C_0 = 0$. Hence ρ is a run from $\boldsymbol{s}$ to $\boldsymbol{t}$. Every original occurrence is used exactly once, so $\boldsymbol{\psi}^{\rho} = \boldsymbol{\psi}^{\pi}$. □

For the union, a short prefix brings every coordinate to at least C_0, after which Lemma B.2 applies. The larger threshold $C = (\|T\| + 1)C_0$ pays for the decrease along that prefix.

Proof of Theorem B.1. Let $C_0 := 2d \|T\|$, so $C = (\|T\| + 1)C_0$. Consider a $\mathbb{Z}$-run $\boldsymbol{s} \xrightarrow{\pi}_{\mathbb{Z}} \boldsymbol{t}$ with $\boldsymbol{s} \in \mathcal{T}_C^i$ and $\boldsymbol{t} \in \mathcal{T}_C^j$. If $\boldsymbol{s}, \boldsymbol{t} \in \mathcal{T}_{C_0}^j$, then Lemma B.2 rearranges π into the required run.

Otherwise $\boldsymbol{s} \in \mathcal{T}_C^i \setminus \mathcal{T}_{C_0}^j$, which implies that $i \neq j$ and $\boldsymbol{s}(i) < C_0$. Since $\boldsymbol{t} \in \mathcal{T}_C^j$ and $i \neq j$, we have $\boldsymbol{t}(i) \geq C$. For the occurrences of π with strictly positive i-th component, the sum of these components is therefore at least $\boldsymbol{t}(i) - \boldsymbol{s}(i) \geq C_0 - \boldsymbol{s}(i)$. Choose α as a subsequence consisting only of such occurrences. Take such occurrences until that coordinate first reaches C_0. Each increases it by at least one, so $|\alpha| \leq C_0 - \boldsymbol{s}(i) \leq C_0$. Coordinate i is nondecreasing along α, and every other coordinate stays at least

$$C - \|T\| \, |\alpha| \geq C - \|T\| \cdot C_0 = C_0.$$

Thus α induces a run to some $\boldsymbol{w} \geq C_0 \cdot \boldsymbol{1}$.

The Parikh image $\boldsymbol{\psi}^{\pi} - \boldsymbol{\psi}^{\alpha}$ gives a $\mathbb{Z}$-run from $\boldsymbol{w}$ to $\boldsymbol{t}$ in V. Since $\boldsymbol{w}, \boldsymbol{t} \in \mathcal{T}_{C_0}^j$, Lemma B.2 rearranges this $\mathbb{Z}$-run into a run β. Then $\alpha\beta$ is a run from $\boldsymbol{s}$ to $\boldsymbol{t}$ with Parikh image $\boldsymbol{\psi}^{\pi}$. This proves that $\mathcal{T}_C$ is a transformation area. Each $\mathcal{T}_C^i$ is a subset of $\mathcal{T}_C$, so it is a transformation area as well. □

C Proofs for Section 5

The following is the formal version of Theorem 1.3.

Theorem C.1. *For every fixed dimension $d \geq 3$, there exists a polynomial P_d such that every d-VAS V admits a finite family $\mathcal{F}$ of sequences of relations $(R_1, \ldots, R_m)$ satisfying the following properties:*

$$\mathrm{Reach}_V = \bigcup_{(R_1, \ldots, R_m) \in \mathcal{F}} R_1 \circ \cdots \circ R_m.$$

Every sequence in $\mathcal{F}$ has $1 \leq m \leq P_d(\mathrm{SIZE}(V))$. Each factor is $R_i = \mathrm{Reach}_{V_i} \subseteq \mathrm{Reach}_V$, where V_i has one of the following forms:

1. *$V_i = [\![V]\!]_{I,B}$ for some $I \subseteq [d]$ with $|I| = 3$ and $B \leq P_d(\mathrm{SIZE}(V))$;*
2. *$V_i = \Lambda$ for some mixed LPS Λ, with $\|\Lambda\|, |\Lambda| \leq P_d(\mathrm{SIZE}(V))$.*

Moreover, for every factor R_i and every $(\boldsymbol{x}, \boldsymbol{y}) \in R_i$,

$$\ell_V(\boldsymbol{x}, \boldsymbol{y}) \leq \ell_{V_i}(\boldsymbol{x}, \boldsymbol{y}).$$

Proof. Let $N := \mathrm{SIZE}(V)$ and fix a nondecreasing polynomial P_0 sufficiently large for all applications of Lemma 5.1, including those with at most two zero-effect dummy coordinates. Set $B_0 := 0$ and $B_{k+1} := P_0(B_k + N)$ for $k \in \{0, 1, 2\}$.

We construct finite families $\mathcal{V}_0, \ldots, \mathcal{V}_3$ of sequences of factors, starting with $\mathcal{V}_0 := \{(V)\}$. At level $k \in \{0, 1, 2\}$, consider a sequence $(V_1, \ldots, V_m) \in \mathcal{V}_k$. We leave its mixed-LPS factors unchanged and replace each factor of the form $V_i = [\![V]\!]_{I,B_k}$ with $|I| = k$ as follows.

If $k < 2$, add $2 - k$ zero-effect dummy coordinates and encode them together with the coordinates in I. We can then apply Lemma 5.1, since exactly two coordinates are encoded. Projecting away the dummy coordinates yields a finite family of compositions whose factors are either mixed LPS or encodings $[\![V]\!]_{I\cup\{j\},D}$ for some original coordinate $j \notin I$ and $D \le B_{k+1}$. Every resulting mixed LPS has zero effect on the dummy coordinates, so this projection preserves its witnesses and their lengths. We enlarge D to B_{k+1} and replace V_i by each of these compositions.

Performing these replacements for every nonterminal factor and taking all possible choices gives $\mathcal{V}_{k+1}$. Thus every nonterminal factor in $\mathcal{V}_{k+1}$ has $k + 1$ bounded original coordinates. After three levels, every factor is either a mixed LPS or an encoding with three bounded coordinates, and the recursion terminates. Each level-k replacement introduces at most $2B_k + 2$ factors. Hence every resulting sequence contains at most

$$\prod_{k=0}^{2}(2B_k + 2)$$

factors, and all counter bounds and mixed-LPS parameters are at most B_3. Since the recursion has constant depth, these bounds are polynomial in N.

Every reachable pair remains represented after each replacement by Lemma 5.1. Conversely, every resulting factor defines a subset of Reach_V, and the distance inequality is preserved by the lemma and by projection of the dummy coordinates. Thus $\mathcal{V}_3$ gives a polynomial factorization. □

Effective Construction. We consider candidate mixed LPSs with polynomially bounded parameters. Each $\mathbb{N}$-component may be an arbitrary LPS using transitions of V. For each $\mathbb{Z}$-component, guess a B-encoding of V with two encoded coordinates, possibly after adding at most two zero-effect dummy coordinates. Choose a state $\boldsymbol{q} \in [0, B]^2$ and a concatenation of paths in the corresponding $2B$-encoding from $\boldsymbol{q}$ back to $\boldsymbol{q}$, with repeatable cycles inserted at matching states. Let Γ be the LPS obtained by recording the transition-effect sequences of these paths and cycles. Let C be the transformation-area threshold supplied by Theorem 4.2 for this $2B$-encoding. Restore the original transition effects and form the block

$$(\chi^C_{\boldsymbol{q}}, \mathbb{N})(\Gamma, \mathbb{Z})(\chi^C_{\boldsymbol{q}}, \mathbb{N}),$$

projecting away any dummy coordinates afterward.

By Lemma 5.2, every such block defines a subrelation of Reach_V and satisfies the distance inequality. Thus every candidate mixed LPS is sound, while the preceding construction guarantees that these candidates include all factors needed for the factorization. The encoded paths and cycles, checking gadgets, and polynomial parameter bounds can all be verified in polynomial time.

D Proofs for Section 6

D.1 Proof of Theorem 6.1

Proof of Theorem 6.1. We follow the construction in [CFZ26, proof of Proposition 5.1]. A witnessing run with two coordinates bounded throughout by a polynomial threshold gives Case II directly. Branches admitting polynomial-length witnesses can be represented by acyclic unfoldings and satisfy Case I. In the remaining branch, the construction yields a sequential 5-VASS V'' of polynomial size with $E := \mathrm{CYC}(V'')$ and $\dim E \le 4$. As shown in [CFZ26, Appendix D.5], a witnessing run splits into two segments, each bounding a pair I_i of coordinates by a polynomial threshold B, with the two coordinate functionals linearly independent on E. Encode I_i into states while retaining all five counters, and connect the two copies at the shared intermediate configuration by one zero-effect transition. Every cycle in the i-th copy has effect in

$$E \cap \{\boldsymbol{z} : \boldsymbol{z}\restriction_{I_i} = \boldsymbol{0}\},$$

whose dimension is $\dim E - 2 \le 2$. The chosen run lifts to the resulting V', while deleting the bridge and forgetting the encoded states projects every run back to V. Thus $\ell_V(\boldsymbol{s},\boldsymbol{t}) \le \ell_{V'}(s',t')$. All sizes and thresholds remain polynomial in $\mathrm{SIZE}(V,\boldsymbol{s},\boldsymbol{t})$. □

D.2 Proof of Theorem 1.2

Theorem 1.2 (Polynomial Reachability Distance)**.** *There exists a constant c such that for every reachable 5-VAS instance $(V,\boldsymbol{s},\boldsymbol{t})$, one has $\ell_V(\boldsymbol{s},\boldsymbol{t}) \le \mathrm{SIZE}(V,\boldsymbol{s},\boldsymbol{t})^c$.*

Proof. Let $(V,\boldsymbol{s},\boldsymbol{t})$ be a reachable 5-VAS and $M := \mathrm{SIZE}(V,\boldsymbol{s},\boldsymbol{t})$. We obtain $B \le M^{\mathcal{O}(1)}$ and the two case distinctions by applying Theorem 6.1 to V.

For Case I, assume that $V'' = (Q,T)$. Then V' has $2|Q|\cdot(B+1)^2$ states, at most $2|T|\cdot(B+1)^2+1$ transitions, and norm at most $\|T\|$, which implies that

$$\mathrm{SIZE}(V') \le (2(B+1)^2+1)\cdot\mathrm{SIZE}(V'').$$

Theorem 6.1 gives that $\ell_V(\boldsymbol{s},\boldsymbol{t}) \le \ell_{V'}(s',t')$ and $B, \mathrm{SIZE}(V'') \le \mathrm{SIZE}(V,\boldsymbol{s},\boldsymbol{t})^{\mathcal{O}(1)}$. Since

$$B + \mathrm{SIZE}(V',s',t') \le \mathcal{O}((B+1)^2)\cdot\mathrm{SIZE}(V,\boldsymbol{s},\boldsymbol{t})^{\mathcal{O}(1)} \le \mathrm{SIZE}(V,\boldsymbol{s},\boldsymbol{t})^{\mathcal{O}(1)},$$

by Proposition 6.2, $\ell_V(\boldsymbol{s},\boldsymbol{t}) \le \ell_{V'}(s',t') \le \mathrm{SIZE}(V,\boldsymbol{s},\boldsymbol{t})^{\mathcal{O}(1)}$.

In Case II, we fix the SCCs and bridges visited by a witnessing run π and obtain U. By Proposition 6.5,

$$\ell_V(\boldsymbol{s},\boldsymbol{t}) \le (B+M)^{\mathcal{O}(1)} \le M^{\mathcal{O}(1)}.$$

The exponents are uniform over all instances, so one constant suffices for both cases. □

D.3 Proof of Lemma 6.8

Geometrically, a proper cycle cone is contained in a half-space with a nonzero normal. The existence and the description size of the separating plane are guaranteed by the following lemma, which is a part of [CFZ26, Lemma 4.5].

Lemma D.1. *Let $X \subseteq \mathbb{Z}^d$ be a finite set of vectors such that $\mathrm{cone}(X) \neq \mathrm{span}(X)$. There exists a nonzero vector $\boldsymbol{n} \in \mathbb{Z}^d$ such that*

1. *$\langle \boldsymbol{n}, \boldsymbol{x} \rangle \geq 0$ for all $\boldsymbol{x} \in X$;*
2. *$\|\boldsymbol{n}\| \leq (d+1) \cdot (r \|X\|)^r$, where $r := \dim(\mathrm{span}(X))$.*

Proof of Lemma 6.8. Assume that $U = (Q, T)$, $s = p(\boldsymbol{x})$, and $t = q(\boldsymbol{y})$. If $\dim_{\mathrm{cyc}}(U) \leq 2$, take $\widetilde{U} = U$ with unchanged endpoints. Otherwise $\mathrm{CYC}(U) = \mathbb{Q}^3$. Let $X := \{\mathrm{EFF}(\theta) : \theta \text{ is a simple cycle in } U\}$. Thus X is finite and $\mathrm{cone}(X) \neq \mathrm{span}(X)$. By Lemma D.1, there exists a nonzero integer normal $\boldsymbol{h}$ satisfying

$$\langle \boldsymbol{h}, \mathrm{EFF}(\theta) \rangle \geq 0 \text{ for every simple cycle } \theta, \qquad \|\boldsymbol{h}\| \leq 4(3|Q| \cdot \|T\|)^3.$$

Let $r(\boldsymbol{z})$ occur on a run from $p(\boldsymbol{x})$ to $q(\boldsymbol{y})$. We delete simple cycles separately from its prefix and suffix. Then each remaining path has length at most $|Q| - 1$, while every deleted cycle has nonnegative inner product with $\boldsymbol{h}$. Consequently, with $D := (|Q| - 1) \|T\| \cdot \|\boldsymbol{h}\|_1$,

$$\langle \boldsymbol{h}, \boldsymbol{x} \rangle - D \leq \langle \boldsymbol{h}, \boldsymbol{z} \rangle \leq \langle \boldsymbol{h}, \boldsymbol{y} \rangle + D.$$

In particular, all these inner products lie in $[-H, H]$, where

$$H := (\|\boldsymbol{x}\| + \|\boldsymbol{y}\| + (|Q| - 1) \|T\|) \cdot \|\boldsymbol{h}\|_1 \leq \mathrm{SIZE}(U, s, t)^{\mathcal{O}(1)}.$$

Define $\widetilde{U}$ with states $(r, b) \in Q \times [-H, H]$ and transitions

$$\big((r, b), \boldsymbol{a}, (r', b + \langle \boldsymbol{h}, \boldsymbol{a} \rangle)\big)$$

whenever $(r, \boldsymbol{a}, r') \in T$ and both states belong to $Q \times [-H, H]$. Take $\widetilde{s} = (p, \langle \boldsymbol{h}, \boldsymbol{x} \rangle)(\boldsymbol{x})$ and $\widetilde{t} = (q, \langle \boldsymbol{h}, \boldsymbol{y} \rangle)(\boldsymbol{y})$. The bound above lifts every original run from s to t to $\widetilde{U}$, and projection gives the converse, preserving lengths. Every cycle of $\widetilde{U}$ returns its recorded inner product to the same value, so its effect lies in $\boldsymbol{h}^\perp := \{\boldsymbol{x} \in \mathbb{Q}^3 : \langle \boldsymbol{h}, \boldsymbol{x} \rangle = 0\}$. Thus $\dim_{\mathrm{cyc}}(\widetilde{U}) \leq 2$. There are at most $|Q| \cdot (2H + 1)$ states and $|T| \cdot (2H + 1)$ transitions, with unchanged transition vectors and endpoint norms, proving the polynomial size bound. □

D.4 Proof of Lemma 6.11

Lemma 6.11. *Let $U = (Q, T)$ be a d-VASS. If $p(\boldsymbol{x}) \xrightarrow{*}_{\mathbb{Z}} q(\boldsymbol{y})$ in U, there is such a $\mathbb{Z}$-run of length at most $(\mathrm{SIZE}(U) + \|\boldsymbol{x}\| + \|\boldsymbol{y}\|)^{\mathsf{poly}(d)}$.*

Proof. We fix a witnessing path π and mark the first occurrence of every visited state. The at most $|Q|$ marks divide π into at most $|Q|$ segments, including the suffix after the last mark. Delete simple cycles from each segment until a simple path remains. The concatenation η of the remaining paths has length at most $|Q|^2$ and still visits every state visited by π, since every marked occurrence remains as an endpoint of one segment. Each deleted simple cycle has length at most $|Q|$ and effect of norm at most $K := |Q| \cdot \|T\|$.

There are at most $(2K + 1)^d$ distinct effects among the deleted cycles. Choose one representative cycle for each such effect, and let $\boldsymbol{A}$ be the matrix of these effects. The original multiplicities give a

nonnegative integer solution of $\boldsymbol{A}\boldsymbol{z} = \boldsymbol{b}$, where $\boldsymbol{b} := \boldsymbol{y} - \boldsymbol{x} - \mathrm{EFF}(\eta)$. By Theorem 3.3, a minimal solution of this system satisfies

$$\|\boldsymbol{z}\|_1 \leq \left(1 + (2K+1)^d K + \|\boldsymbol{x}\| + \|\boldsymbol{y}\| + |Q|^2 \|T\|\right)^d \leq \left(\mathrm{SIZE}(U) + \|\boldsymbol{x}\| + \|\boldsymbol{y}\|\right)^{\mathcal{O}(d^2)}.$$

Here the augmented matrix $\begin{pmatrix} \boldsymbol{A} & \boldsymbol{b} \end{pmatrix}$ has rank at most d. If $\boldsymbol{b} = \boldsymbol{0}$, the zero solution suffices.

Since every state visited by π still appears in η, each representative cycle can be placed at a state visited by η. Using the multiplicities given by $\boldsymbol{z}$, we obtain a $\mathbb{Z}$-run of total effect $\boldsymbol{y} - \boldsymbol{x}$. Its length is at most $|\eta| + |Q| \cdot \|\boldsymbol{z}\|_1 \leq \left(\mathrm{SIZE}(U) + \|\boldsymbol{x}\| + \|\boldsymbol{y}\|\right)^{\mathsf{poly}(d)}$. □

D.5 Proof of Lemma 6.17

Lemma 6.17. *Let V be a 5-VAS, $B \in \mathbb{N}$, and C a transformation-area threshold for $[\![V]\!]_B$. Let H_1, H_2 be hybrid sets at states $\boldsymbol{p}, \boldsymbol{q}$ of a sequential sub-VASS U of $[\![V]\!]_B$. Suppose that there exist $\boldsymbol{x} \in H_1$ and $\boldsymbol{y} \in H_2$ such that $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{*} \boldsymbol{q}(\boldsymbol{y})$ in U. Assume that H_1 is forward C-coverable in U and H_2 is backward C-coverable in U. Then there exist $\boldsymbol{x}' \in H_1$ and $\boldsymbol{y}' \in H_2$ such that*

$$\left\|\boldsymbol{x}'\right\|, \left\|\boldsymbol{y}'\right\|, \ell_V\left((\boldsymbol{p}, \boldsymbol{x}'), (\boldsymbol{q}, \boldsymbol{y}')\right) \leq \left(B + C + \mathrm{SIZE}(V) + \mathrm{NORM}(H_1) + \mathrm{NORM}(H_2)\right)^{\mathcal{O}(1)}.$$

Proof. Write $H_i = \boldsymbol{b}_i + P_i^{\langle \boldsymbol{c}_i, \cdot \rangle \leq K_i}$ for $i = 1, 2$, write $\boldsymbol{c}_i(\boldsymbol{a})$ for the cost of $\boldsymbol{a} \in P_i$, and fix representations $\boldsymbol{x} = \boldsymbol{b}_1 + \sum_{\boldsymbol{a} \in P_1} \lambda^*_{\boldsymbol{a}} \cdot \boldsymbol{a}$ and $\boldsymbol{y} = \boldsymbol{b}_2 + \sum_{\boldsymbol{a} \in P_2} \mu^*_{\boldsymbol{a}} \cdot \boldsymbol{a}$ satisfying K_1 and K_2, respectively. Consider the augmented VASS $U[P_1, P_2]$. The fixed representations and the hypothesis give $\boldsymbol{p}(\boldsymbol{b}_1) \xrightarrow{*} \boldsymbol{q}(\boldsymbol{b}_2)$ in $U[P_1, P_2]$. Moreover, the two bases are respectively forward and backward C-coverable in $U[P_1, P_2]$. By Corollary 4.7, C remains a transformation-area threshold after adjoining these periods. Hence Corollary 6.14 gives a run ρ in the augmented VAS from $(\boldsymbol{p}, \boldsymbol{b}_1)$ to $(\boldsymbol{q}, \boldsymbol{b}_2)$ of length $L \leq M^{\mathcal{O}(1)}$, where $M := B + C + \mathrm{SIZE}(V) + \mathrm{NORM}(H_1) + \mathrm{NORM}(H_2)$.

For $\boldsymbol{a} \in P_1$, let $\lambda_{\boldsymbol{a}}$ be the number of occurrences of $+\boldsymbol{a}$ in ρ, and for $\boldsymbol{a} \in P_2$, let $\mu_{\boldsymbol{a}}$ be the number of occurrences of $-\boldsymbol{a}$. Let

$$\boldsymbol{x}' := \boldsymbol{b}_1 + \sum_{\boldsymbol{a} \in P_1} \lambda_{\boldsymbol{a}} \cdot \boldsymbol{a}, \qquad \boldsymbol{y}' := \boldsymbol{b}_2 + \sum_{\boldsymbol{a} \in P_2} \mu_{\boldsymbol{a}} \cdot \boldsymbol{a}.$$

Delete all added transitions from ρ and start the remaining path at $\boldsymbol{x}'$. At each position on the resulting path, its vector equals the corresponding vector of ρ plus the sum of the unused positive periods and the absolute effects of the negative periods already removed. Hence the resulting path is a run in V from $(\boldsymbol{p}, \boldsymbol{x}')$ to $(\boldsymbol{q}, \boldsymbol{y}')$, of length at most L.

We still need to verify the cost bounds. Since ρ has length at most L, the total cost of the added period occurrences on side i is at most $\|\boldsymbol{c}_i\| L$. Hence, if $K_i \geq \|\boldsymbol{c}_i\| L$, the corresponding vector already belongs to H_i.

If $K_i < \|\boldsymbol{c}_i\| L$, apply Lemma 6.16 to the original member $\boldsymbol{x}$ or $\boldsymbol{y}$ on side i. This gives a subset $\widehat{H}_i$ of H_i containing that member, with only free periods and base norm at most $\|\boldsymbol{b}_i\| + K_i \|P_i\| \leq M^{\mathcal{O}(1)}$. Thus the original connected pair $(\boldsymbol{p}(\boldsymbol{x}), \boldsymbol{q}(\boldsymbol{y}))$ remains available. Since the base increases and the free periods are unchanged, forward or backward C-coverability is preserved by monotonicity. We repeat the construction with the new hybrid sets. Each side can be absorbed at most once, so at most two repetitions are needed. The final vectors belong to the original hybrid sets H_1, H_2, have norm at most $M^{\mathcal{O}(1)}$, and are connected in V by a run of length at most $M^{\mathcal{O}(1)}$. □

E Proofs for Section 7

E.1 Proof of Lemma 7.2

Lemma 7.2. *Let $U = (U_1)u_1(U_2)\cdots u_{k-1}(U_k)$ be a sequential 3-VASS, and fix a run $p(\boldsymbol{x}) \xrightarrow{\pi} q(\boldsymbol{y})$ whose source is not forward C-coverable. Then there exists an integer $D \le (C + \textsc{size}(U))^{\mathcal{O}(1)}$ such that one can find a coordinate $i \in [3]$ and a prefix $p(\boldsymbol{x}) \xrightarrow{\pi_0} r(\boldsymbol{z})$ with $i \in \textsc{bd}(\pi_0, D)$ satisfying one of the following:*

1. *$r(\boldsymbol{z}) = q(\boldsymbol{y})$;*
2. *$r(\boldsymbol{z})$ is the entry configuration of a later component U_h, and every $\boldsymbol{w} \in \mathbb{N}^3$ satisfying*

$$\boldsymbol{w}(i) \ge \boldsymbol{z}(i), \qquad \boldsymbol{w}(j) \ge D \quad \text{for all } j \ne i$$

has $r(\boldsymbol{w})$ forward C-coverable in U_h.

Moreover, π_0 is captured by a 2-VASS of size at most $D\,\textsc{size}(U)$.

Proof. Assume that $U = (Q, T)$ with $n := |Q|$. Let $D := D_0 + \|T\|$, where

$$L_c := \max\{2, n, \|T\|, C\}^{c_{\mathrm{cov}}}, \qquad D_0 := \max\{(C + n\|T\| + \textsc{size}(U))^{c_{\mathrm{ext}}}, C + L_c\|T\| + 1\}.$$

Here $c_{\mathrm{cov}}, c_{\mathrm{ext}} \ge 1$ are fixed integer constants from Lemmas 3.1 and 3.2 for $d = 3$. Write the witness as $\pi = \xi_1 u_1 \xi_2 \cdots u_{k-1}\xi_k$, where ξ_j stays in U_j. Some coordinate remains below D_0 along ξ_1. Otherwise, Lemma 3.2 gives a run in U_1 from $p(\boldsymbol{x})$ to a configuration $q'(\boldsymbol{w})$ with $\boldsymbol{w} \ge (C + n\|T\|)\cdot\mathbf{1}$. A simple path from q' back to p has length at most $n-1$ and decreases each coordinate by at most $(n-1)\|T\|$. It is fireable at $\boldsymbol{w}$ and returns with at least $C\cdot\mathbf{1}$, contradicting the hypothesis.

Choose $i \in \textsc{bd}(\xi_1, D_0)$ and put

$$h := \min\{j \in \{2, \dots, k\} : i \notin \textsc{bd}(\xi_j, D_0)\},$$

with $h := k+1$ if the set is empty. If $h \le k$, set $\pi_0 := \xi_1 u_1 \cdots \xi_{h-1}u_{h-1}$. Otherwise, set $\pi_0 := \pi$. Every included component segment keeps coordinate i below D_0, including ξ_1. Across each included bridge, its value satisfies

$$\boldsymbol{w}'(i) = \boldsymbol{w}(i) + \textsc{eff}(u_j)(i) < D_0 + \|T\| = D.$$

This proves that $i \in \textsc{bd}(\pi_0, D)$.

Suppose $h \le k$, and let $r(\boldsymbol{z})$ be the entry of U_h. There is a prefix η of ξ_h reaching $q'(\boldsymbol{w})$ with $\boldsymbol{w}(i) \ge D_0$. Choose a simple path ζ from q' to r in U_h. Coordinate i is nonnegative along η, and every prefix ζ' of ζ satisfies

$$\boldsymbol{w}(i) + \textsc{eff}(\zeta')(i) \ge D_0 - (n-1)\|T\| \ge C.$$

Projecting $\rho := \eta\zeta$ onto coordinate i gives a run from $r(\boldsymbol{z}(i))$ to $r(v)$ with $v \ge C$. By Lemma 3.1, there is such a covering run of length less than L_c. Now fix any $\boldsymbol{w}$ satisfying the conditions in the second outcome. Lift the shortened run to a $\mathbb{Z}$-run ρ' in U_h from $r(\boldsymbol{w})$ to $r(v)$. Since $\boldsymbol{w}(i) \ge \boldsymbol{z}(i)$, coordinate i is nonnegative along ρ' and $v(i) \ge C$. For every $j \ne i$ and every prefix η' of ρ', we have $\boldsymbol{w}(j) + \textsc{eff}(\eta')(j) \ge D - L_c\|T\| > C$. Therefore, ρ' is a run and $r(\boldsymbol{w})$ is forward C-coverable in U_h.

Finally, encoding coordinate i in $\{0, \dots, D-1\}$ gives the 2-VASS $[\![U]\!]_{\{i\},D-1}$ with

$$\textsc{size}([\![U]\!]_{\{i\},D-1}) \le D\,\textsc{size}(U).$$

Projection and lifting preserve run lengths and represent exactly the runs keeping coordinate i below D. The construction remains valid with any larger D if D_0 is replaced by $D - \|T\|$. □

E.2 Proof of Lemma 7.5

Lemma 7.5. *Let $U = (U_1)u_1 \cdots u_{k-1}(U_k)$ be a sequential 3-VASS, let $p(\boldsymbol{x}) \xrightarrow{\pi} t$ be a run in U, and let $\boldsymbol{a}_0 \in \mathbb{N}^3 \setminus \{\mathbf{0}\}$ be the effect of an enabled cycle at $p(\boldsymbol{x})$. There exist*

$$D \leq (C + \textsc{size}(U) + \|\boldsymbol{a}_0\|)^{\mathcal{O}(1)} \tag{10}$$

and a possibly empty prefix $p(\boldsymbol{x}) \xrightarrow{\pi_*} r_*(\boldsymbol{z}_*)$ that is D-alternating on $[3] \setminus \operatorname{supp}(\boldsymbol{a}_0)$. Either $r_*(\boldsymbol{z}_*) = t$, or it is the entry of some component U_* such that $r_*(\boldsymbol{z}_*)$ is forward C-coverable in $U_*[\{\boldsymbol{a}_0\}]$.

Proof. Write $U = (Q, T)$ and put $I := [3] \setminus \operatorname{supp}(\boldsymbol{a}_0)$ and $n := |Q|$. Choose D by the formula in the proof of Lemma 7.2, using the parameters of $U[\{\boldsymbol{a}_0\}]$ and the threshold C. These parameters bound those of every suffix augmented with the same loop at its source, so the same D works throughout and satisfies (10). Let $r_*(\boldsymbol{z}_*)$ be the first component entry on π that is forward C-coverable after adjoining this loop there. If no such entry exists, take $r_*(\boldsymbol{z}_*) := t$. Let π_* be the original prefix ending at this position. If $\pi_* = \varepsilon$, its empty partition proves the claim. In particular, this handles $I = \varnothing$, since repetitions of the added loop already give forward C-coverability at the source.

We divide π_* into bounded prefixes by repeated applications of Lemma 7.2. Suppose the unprocessed suffix starts at an original component entry $p_j(\boldsymbol{x}_j)$ before $r_*(\boldsymbol{z}_*)$. By its choice, this entry is not forward C-coverable in its component augmented with a self-loop of effect $\boldsymbol{a}_0$. A run returning to this entry cannot cross a bridge, so the entry is also not forward C-coverable in the augmented suffix. Execute the loop D times and replay the unprocessed part of π_* translated by $D \cdot \boldsymbol{a}_0$. By monotonicity this is a run from $p_j(\boldsymbol{x}_j)$, including the initial loops, and every coordinate in $\operatorname{supp}(\boldsymbol{a}_0)$ reaches at least D within its first component. By the construction in the proof of Lemma 7.2, the extracted prefix contains all the initial loop occurrences. Its bounded coordinate therefore lies in I. Removing these loops and undoing the translation by $D \cdot \boldsymbol{a}_0$ gives a prefix ξ of the original run with

$$\textsc{bd}(\xi, D) \cap I \neq \varnothing.$$

If ξ stops at a later entry before $r_*(\boldsymbol{z}_*)$, apply the same construction to the remaining suffix. Continuing in this way gives a partition

$$\pi_* = \xi_1 \xi_2 \cdots \xi_m$$

such that

$$\textsc{bd}(\xi_j, D) \cap I \neq \varnothing \qquad \text{for every } j \in [m]. \tag{24}$$

The same threshold D is used throughout, and each nonterminal application passes at least one bridge, so $m \leq k$. Repeatedly merge adjacent segments that share a coordinate in I remaining below D. Since

$$\textsc{bd}(\alpha\beta, D) \cap I = \big(\textsc{bd}(\alpha, D) \cap I\big) \cap \big(\textsc{bd}(\beta, D) \cap I\big),$$

each merged segment still satisfies (24). When no further merge is possible, the resulting partition satisfies (9). □

E.3 Proof of Lemma 7.7

Lemma 7.7. *Let $U = (Q, T)$ be a 3-VASS, and let $i, j \in [3]$ be distinct coordinates. Suppose $p(\boldsymbol{x}) \xrightarrow{\pi} q(\boldsymbol{y})$ is a run with $D \geq 1$ satisfying $j \in \textsc{bd}(\pi, D)$ and*

$$\boldsymbol{y}(i) \geq D' := \boldsymbol{x}(i) + D|Q|\,\|T\| + 1.$$

Then there exist a configuration $r(\boldsymbol{z})$ along π and a cycle θ fireable at $r(\boldsymbol{z})$ such that

$$\textsc{eff}(\theta)(i) > 0, \qquad \textsc{eff}(\theta)(j) = 0, \qquad |\theta| \leq (D + \textsc{size}(U) + \boldsymbol{x}(i))^{\mathcal{O}(1)}.$$

Proof. By taking a prefix, we assume w.l.o.g. that $q(\boldsymbol{y})$ is the first configuration where coordinate i has value at least D'. Let $r_0(\boldsymbol{u}_0), \ldots, r_s(\boldsymbol{u}_s)$ be the successive strict records of coordinate i, including $r_0(\boldsymbol{u}_0) := p(\boldsymbol{x})$ and $r_s(\boldsymbol{u}_s) := q(\boldsymbol{y})$. Each record exceeds the preceding one by at most $\|T\|$, so

$$s\,\|T\| \geq \boldsymbol{u}_s(i) - \boldsymbol{x}(i) \geq D|Q|\,\|T\| + 1, \qquad s \geq D|Q| + 1.$$

Here $\|T\| \geq 1$, since coordinate i increases along π. The first s records precede the crossing, and their state and j-coordinate pairs take at most $D|Q|$ values. Hence there are $0 \leq a < b < s$ with

$$r_a = r_b, \qquad \boldsymbol{u}_a(j) = \boldsymbol{u}_b(j), \qquad \boldsymbol{u}_a(i) < \boldsymbol{u}_b(i).$$

The intervening sub-run keeps coordinate i below D' and coordinate j below D. Encode these two coordinates into states, obtaining a 1-VASS with at most $|Q|DD'$ states and transition norm at most $\|T\|$. The sub-run witnesses coverability of $\widehat{r}_b(0)$ from $\widehat{r}_a(\boldsymbol{u}_a(k))$, where $\{k\} = [3] \setminus \{i, j\}$ and $\widehat{r}_c := (r_c, \boldsymbol{u}_c(i), \boldsymbol{u}_c(j))$ for $c \in \{a, b\}$. By Lemma 3.1, there is such a covering run of length

$$\max\{2, |Q|DD', \|T\|\}^{\mathcal{O}(1)} \leq (D + \textsc{size}(U) + \boldsymbol{x}(i))^{\mathcal{O}(1)}.$$

Lifting this run to U gives an enabled cycle θ at $r_a(\boldsymbol{u}_a)$, since $r_a = r_b$. Its encoded endpoints guarantee $\textsc{eff}(\theta)(i) = \boldsymbol{u}_b(i) - \boldsymbol{u}_a(i) > 0$ and $\textsc{eff}(\theta)(j) = 0$, as required. □

E.4 Proof of Lemma 7.9

Lemma 7.9. *Let $p(\boldsymbol{x}) \xrightarrow{\pi} t = q(\boldsymbol{y})$ be a run in a sequential 3-VASS $U = (Q, T)$, with $p(\boldsymbol{x})$ pumpable, and let $C \in \mathbb{N}$. There exist a partition of π*

$$p(\boldsymbol{x}) \xrightarrow{\tau} r(\boldsymbol{z}) \xrightarrow{\sigma} t,$$

and a hybrid set H containing $\boldsymbol{z}$, with $\textsc{norm}(H) \leq (C + \textsc{size}(U, p(\boldsymbol{x})))^{\mathcal{O}(1)}$, such that every $r(\boldsymbol{h})$ with $\boldsymbol{h} \in H$ has a prefix witness $p(\boldsymbol{x}) \xrightarrow{\rho} r(\boldsymbol{h})$ of length

$$|\rho| \leq (C + \textsc{size}(U, p(\boldsymbol{x}), r(\boldsymbol{h})))^{\mathcal{O}(1)}.$$

Either σ is empty, or H is forward C-coverable in the component containing r.

Proof. Let $M := C + \textsc{size}(U, p(\boldsymbol{x}))$. Pumpability gives a coordinate i for which $p(\boldsymbol{x} + \boldsymbol{e}_i)$ is coverable from $p(\boldsymbol{x})$. By Lemma 3.1, there is a cycle θ_0 fireable from $p(\boldsymbol{x})$ such that

$$\boldsymbol{a}_0 := \textsc{eff}(\theta_0) \in \mathbb{N}^3 \setminus \{\boldsymbol{0}\}, \qquad |\theta_0|, \|\boldsymbol{a}_0\| \leq M^{O(1)}.$$

Permute coordinates so that $\boldsymbol{a}_0(3) > 0$, and put $I := [3] \setminus \operatorname{supp}(\boldsymbol{a}_0) \subseteq \{1, 2\}$.

Apply Lemma 7.5 to obtain a partition $\pi = \pi_*\sigma_*$, where π_* ends at $r_*(\boldsymbol{z}_*)$ and is D-alternating on I for some $D \leq M^{O(1)}$. If σ_* is nonempty, its starting configuration is forward C-coverable in its component after adjoining $\boldsymbol{a}_0$ as a self-loop. Since $|I| \leq 2$, the prefix π_* is also D-alternating

on $\{1,2\}$. Apply Lemma 7.6 to π_* and enlarge D_1 if necessary so that $D_1 > \max\{D, \|\boldsymbol{x}\|\}$. Then $D_1, L \le M^{\mathcal{O}(1)}$. We distinguish the two outcomes of Lemma 7.6.

Suppose first that $j \in \mathrm{BD}(\pi_*, D_1)$ for some $j \in [3]$. Encode this coordinate into states, obtaining the 2-VASS $\widetilde{U} := [\![U]\!]_{\{j\},D_1-1}$ of size at most $D_1 \,\mathrm{SIZE}(U) \le M^{\mathcal{O}(1)}$. Apply Theorem 6.3 with the encoded singleton source, and choose a hybrid set at state $(r_*, \boldsymbol{z}_*(j))$ containing the projection of $\boldsymbol{z}_*$. Inserting $\boldsymbol{z}_*(j)$ into its base and zeros into its periods gives a hybrid set G containing $\boldsymbol{z}_*$, with $\mathrm{NORM}(G) \le M^{\mathcal{O}(1)}$. For each $\boldsymbol{h} \in G$, apply Lemma 6.4 in $\widetilde{U}$ with the same singleton source and the encoded $r_*(\boldsymbol{h})$ as target. The resulting short run lifts to U. Thus every $\boldsymbol{h} \in G$ has a run from $p(\boldsymbol{x})$ to $r_*(\boldsymbol{h})$ in U of length at most $(M + \|\boldsymbol{h}\|)^{\mathcal{O}(1)}$. If $r_*(\boldsymbol{z}_*) = t$, take $\tau := \pi$, $\sigma := \varepsilon$, and $H := G$, which proves the lemma in this case. Otherwise, $r_*(\boldsymbol{z}_*)$ is a component entry that is forward C-coverable after adjoining the $\boldsymbol{a}_0$-loop there. Adjoin $\boldsymbol{a}_0$ as a free period and write

$$G^+ := G + \boldsymbol{a}_0^* = \boldsymbol{b} + P^{\langle \boldsymbol{c},\cdot\rangle \le K}, \qquad \boldsymbol{z}_* \in G^+.$$

For every $\boldsymbol{h} \in G$ and $k \in \mathbb{N}$, reach $\boldsymbol{h} + k \cdot \boldsymbol{a}_0 \in G^+$ by pumping θ_0^k at the source and translating a witness reaching $\boldsymbol{h}$. Since $k \le \|\boldsymbol{h} + k \cdot \boldsymbol{a}_0\|$, this again yields polynomial-length witnesses.

Claim E.1. *There is a hybrid set H with $\boldsymbol{z}_* \in H \subseteq G^+$ and $\mathrm{NORM}(H) \le M^{\mathcal{O}(1)}$ that is forward C-coverable in the component containing r_*.*

Proof of Claim. By Lemma 3.1, choose a run ρ_{cov} of length $M^{\mathcal{O}(1)}$ from $r_*(\boldsymbol{z}_*)$ covering $r_*(C \cdot \mathbf{1})$ in this component augmented with the $\boldsymbol{a}_0$-loop. Let $\boldsymbol{v} := \max\{\mathrm{DROP}(\rho_{\mathrm{cov}}), C \cdot \mathbf{1} - \mathrm{EFF}(\rho_{\mathrm{cov}}), \mathbf{0}\}$, so $\boldsymbol{z}_* \ge \boldsymbol{v}$ and $\|\boldsymbol{v}\| \le M^{\mathcal{O}(1)}$. Choose a representation $\boldsymbol{z}_* = \boldsymbol{b} + P\boldsymbol{\eta}$ in G^+ with $\langle \boldsymbol{c}, \boldsymbol{\eta}\rangle \le K$. Starting from $\boldsymbol{b}$, add unused period occurrences from this representation that increase a deficient coordinate until the resulting base dominates $\boldsymbol{v}$. Each addition decreases the total integer deficit, giving $\mathbf{0} \le \boldsymbol{\mu} \le \boldsymbol{\eta}$ with $\|\boldsymbol{\mu}\|_1 \le 3\|\boldsymbol{v}\|$ and $\boldsymbol{b}' := \boldsymbol{b} + P\boldsymbol{\mu} \ge \boldsymbol{v}$. Set

$$H := \boldsymbol{b}' + P^{\langle \boldsymbol{c},\cdot\rangle \le K - \langle \boldsymbol{c},\boldsymbol{\mu}\rangle}.$$

Then $\boldsymbol{z}_* \in H \subseteq G^+$ and $\mathrm{NORM}(H) \le M^{\mathcal{O}(1)}$. Since $\boldsymbol{a}_0$ is a free period of H and $\boldsymbol{b}' \ge \boldsymbol{v}$, the same run ρ_{cov} witnesses forward C-coverability of H. ∎

Take H from Claim E.1, $\tau := \pi_*$, and let σ be the original remaining suffix. The witnesses for G^+ give the required length bound for every member of H.

Otherwise, $\mathrm{BD}(\pi_*, D_1) = \emptyset$. Since D-alternation on I with $|I| \le 1$ implies that some coordinate stays below D_1 throughout π_*, we must have $I = \{1,2\}$ and $\boldsymbol{a}_0 = (0,0,a)$ for some $a \ge 1$. Lemma 7.6 supplies two ordered cycles θ_1, θ_2 and a partition $\pi_* = \rho_1\rho_2\rho$. Apply Lemma 7.8 with common cycle-length bound $\max\{|\theta_0|, L\} \le M^{\mathcal{O}(1)}$. The resulting $U^{(q_1,q_2)}$ and period $\boldsymbol{a}_+ \ge \mathbf{1}$ satisfy $\mathrm{SIZE}(U^{(q_1,q_2)}), \|\boldsymbol{a}_+\| \le M^{\mathcal{O}(1)}$. Its first property gives a run to $f(\boldsymbol{z}_2)$ on which coordinate 2 remains below D_1. Encode this coordinate into states, obtaining $\widetilde{U}^{(q_1,q_2)} := [\![U^{(q_1,q_2)}]\!]_{\{2\},D_1-1}$. Its size satisfies

$$\mathrm{SIZE}(\widetilde{U}^{(q_1,q_2)}) \le D_1 \,\mathrm{SIZE}(U^{(q_1,q_2)}) \le M^{\mathcal{O}(1)}.$$

As above, apply Theorem 6.3 with the encoded singleton source and choose a hybrid set at state $(f, \boldsymbol{z}_2(2))$ containing the projection of $\boldsymbol{z}_2$. Recovering the fixed second coordinate gives a hybrid set G with

$$\boldsymbol{z}_2 \in G \subseteq \mathrm{Reach}^f_{U^{(q_1,q_2)}}(p^{(0)}(\boldsymbol{x})), \qquad \mathrm{NORM}(G) \le M^{\mathcal{O}(1)}.$$

For each $\boldsymbol{h} \in G$, Lemma 6.4 in $\widetilde{U}^{(q_1,q_2)}$, with the same singleton source and the encoded $f(\boldsymbol{h})$ as target, gives a run to $f(\boldsymbol{h})$ in $U^{(q_1,q_2)}$ of length at most $(M + \|\boldsymbol{h}\|)^{\mathcal{O}(1)}$.

Set $H := G + \boldsymbol{a}_+^*$, assigning the added period cost zero. Then $\boldsymbol{z}_2 \in H$ and $\textsc{norm}(H) \le M^{\mathcal{O}(1)}$. For $\boldsymbol{v} = \boldsymbol{h} + k \cdot \boldsymbol{a}_+ \in H$, we have $\boldsymbol{h} \le \boldsymbol{v}$ and $k \le \|\boldsymbol{v}\|$, so Lemma 7.8 gives

$$\ell_U\big(p(\boldsymbol{x}), q_2(\boldsymbol{v})\big) \le (M + \|\boldsymbol{h}\|)^{\mathcal{O}(1)} + kM^{\mathcal{O}(1)} \le (M + \|\boldsymbol{v}\|)^{\mathcal{O}(1)}.$$

Take $\tau := \rho_1\rho_2$ and $\sigma := \rho\sigma_*$, so that $\pi = \tau\sigma$ with the boundary configuration $q_2(\boldsymbol{z}_2)$. Since $\boldsymbol{a}_+ \ge \mathbf{1}$ is a free period of H, its base is forward C-coverable in the component containing q_2 after adjoining the corresponding self-loop. This completes the proof. □

E.5 Proof of Lemma 7.10

Lemma 7.10. *Let U be a sequential sub-VASS of $[\![V]\!]_B$, let C be a transformation-area threshold for $[\![V]\!]_B$, and let H be a hybrid set at state $\boldsymbol{r}$ with a nonzero free period. Suppose that some member of H reaches a backward C-coverable target $\boldsymbol{q}(\boldsymbol{y})$ in U. Then there exists $\boldsymbol{h} \in H$ such that*

$$\|\boldsymbol{h}\|, \ell_V\big((\boldsymbol{r},\boldsymbol{h}),\boldsymbol{t}\big) \le \big(B + C + \textsc{size}(V,\boldsymbol{t}) + \textsc{norm}(H)\big)^{\mathcal{O}(1)}, \text{ where } \boldsymbol{t} = (\boldsymbol{q},\boldsymbol{y}).$$

Proof. Write $H = \boldsymbol{b} + P^{\langle \boldsymbol{c},\cdot\rangle \le K}$ and let $M := C + B + \textsc{size}(V,\boldsymbol{t}) + \textsc{norm}(H)$. Consider $U[P]$ with specified source state $\boldsymbol{r}$. Let V^+ be the original 5-VAS with the lifted periods $(0,0,\boldsymbol{a})$, $\boldsymbol{a} \in P$, adjoined as transitions. Note that $U[P]$ is a sequential sub-VASS of $[\![V^+]\!]_B$. Clearly, $\boldsymbol{r}(\boldsymbol{b}) \xrightarrow{*} \boldsymbol{q}(\boldsymbol{y})$ in $U[P]$, since some member of H reaches this target in U. A nonzero free period of H makes $\boldsymbol{r}(\boldsymbol{b})$ pumpable in $U[P]$. The original backward C-covering run of the target is still available. By Corollary 4.7, C remains a transformation-area threshold after adjoining the periods to V. Since $\textsc{size}(U) \le (B+1)^2\,\textsc{size}(V)$ and $|P| \le (\|P\|+1)^3$, both $U[P]$ and V^+ have size polynomial in M. By Lemma 7.3, there is a run in V^+

$$(\boldsymbol{r},\boldsymbol{b}) \xrightarrow{\rho} \boldsymbol{t}, \qquad |\rho| \le L := F_0(M)$$

for a fixed polynomial F_0.

Remove the added periods as in the proof of Lemma 6.17, with singleton target $\{\boldsymbol{y}\}$. This gives a run ρ' in V with

$$(\boldsymbol{r},\boldsymbol{h}) \xrightarrow{\rho'} \boldsymbol{t}, \qquad |\rho'| \le L, \qquad \|\boldsymbol{h}\| \le \|\boldsymbol{b}\| + L\,\|P\| \le M^{\mathcal{O}(1)}. \tag{25}$$

If $K \ge \|\boldsymbol{c}\|\,L$, the same cost estimate gives $\boldsymbol{h} \in H$. If $K < \|\boldsymbol{c}\|\,L$, choose $\boldsymbol{z} \in H$ such that $\boldsymbol{r}(\boldsymbol{z}) \xrightarrow{*} \boldsymbol{q}(\boldsymbol{y})$ in U, and let $P_0 := \{\boldsymbol{a} \in P : \boldsymbol{c}(\boldsymbol{a}) = 0\}$. Applying Lemma 6.16 to a representation of $\boldsymbol{z}$ in H gives

$$\boldsymbol{z} \in \widehat{H} := \widehat{\boldsymbol{b}} + P_0^* \subseteq H, \qquad \left\|\widehat{\boldsymbol{b}}\right\| \le M^{\mathcal{O}(1)}.$$

Since $\boldsymbol{z} \in \widehat{H}$ and P_0 contains a nonzero period, $\boldsymbol{r}(\widehat{\boldsymbol{b}})$ reaches $\boldsymbol{q}(\boldsymbol{y})$ and is pumpable in $U[P_0]$. The target remains backward C-coverable, so reapply Lemma 7.3 and remove the added periods as in (25). All remaining periods are free, so the selected member belongs to $\widehat{H} \subseteq H$, with its norm and distance to $\boldsymbol{t}$ bounded by $M^{\mathcal{O}(1)}$. □

E.6 Proof of Claim 7.13

Claim 7.13. *There exist indices $0 \le r, \ell \le \kappa_0$ and members $\boldsymbol{h} \in H_r$, $\boldsymbol{h}' \in J_\ell$ whose norms and reachability distance $\ell_V((\boldsymbol{p}, \boldsymbol{h}), (\boldsymbol{q}, \boldsymbol{h}'))$ are at most $M^{\mathcal{O}(1)}$.*

Proof of Claim. Let $\lambda_{\boldsymbol{d}}$ and $\mu_{\boldsymbol{d}'}$ count the added positive and negative period occurrences in ρ, and define

$$\boldsymbol{\xi} := \sum_{\boldsymbol{d} \in P} \lambda_{\boldsymbol{d}} \cdot \boldsymbol{d}, \qquad \boldsymbol{\xi}' := \sum_{\boldsymbol{d}' \in P'} \mu_{\boldsymbol{d}'} \cdot \boldsymbol{d}'.$$

Set $\boldsymbol{h} := \boldsymbol{b}_r + \boldsymbol{\xi}$ and $\boldsymbol{h}' := \boldsymbol{b}'_\ell + \boldsymbol{\xi}'$. Delete the added occurrences and start the remaining transitions at $(\boldsymbol{p}, \boldsymbol{h})$. At every position of the resulting path, the vector is increased by the unused positive periods and the already removed negative periods. Thus the remaining transitions form a run in V from $(\boldsymbol{p}, \boldsymbol{h})$ to $(\boldsymbol{q}, \boldsymbol{h}')$ of length at most L. Moreover,

$$\|\boldsymbol{h}\|, \|\boldsymbol{h}'\| \le \max\{\|\boldsymbol{b}\|, \|\boldsymbol{b}'\|\} + \kappa_0 \max\{\|\boldsymbol{\delta}\|, \|\boldsymbol{\delta}'\|\} + L \max\{\|P\|, \|P'\|\} \le M^{\mathcal{O}(1)}.$$

Both selected vectors meet their cost bounds if

$$K \ge \|\boldsymbol{c}\| L, \qquad K' \ge \|\boldsymbol{c}'\| L.$$

If either cost bound fails, we absorb the restricted periods on that side. We describe the source side, where $K < \|\boldsymbol{c}\| L$; the target side is symmetric. Apply Lemma 6.16 to the fixed representation of $\boldsymbol{z}$ to obtain

$$\widehat{H}_s := \widehat{\boldsymbol{b}} + s \cdot \boldsymbol{\delta} + P_0^* \subseteq H_s, \qquad \boldsymbol{z} \in \widehat{H}_0, \qquad \left\|\widehat{\boldsymbol{b}}\right\| \le \|\boldsymbol{b}\| + K \|P\|,$$

where $0 \le s \le \kappa$. The original connected members remain available, while monotonicity preserves the covering runs and κ. We repeat the augmented-run construction with the reduced period sets. The bound in (13) still applies. Each failure eliminates all restricted periods on at least one side, so at most two repetitions suffice to obtain members of the original families with norms and reachability distance at most $M^{\mathcal{O}(1)}$. ■

E.7 Proof of Lemma 7.15

Lemma 7.15. *Let $U = (Q, T)$ be a 2-VASS with $S = \textsc{size}(U)$, and let p, q be states of U.*

- *There exists a finite family $\mathcal{L}$ of LPSs Λ satisfying $|\Lambda| \le S^{\mathcal{O}(1)}$ and $\|\Lambda\| \le \|T\|$ such that*

$$\mathrm{Reach}_U^q(p(\boldsymbol{x})) = \bigcup_{\Lambda \in \mathcal{L}} \{\boldsymbol{y} : (\boldsymbol{x}, \boldsymbol{y}) \in \mathrm{Reach}_\Lambda\} \qquad \text{for every } \boldsymbol{x} \in \mathbb{N}^2.$$

  Moreover, whenever $p(\boldsymbol{x}) \xrightarrow{*} q(\boldsymbol{y})$, $\ell_U(p(\boldsymbol{x}), q(\boldsymbol{y})) \le \left(S + \|\boldsymbol{x}\| + \|\boldsymbol{y}\|\right)^{\mathcal{O}(1)}$.

- *Let $A = \boldsymbol{b}_A + P_A^{\langle \boldsymbol{c}_A, \cdot \rangle \le K_A}$ be a hybrid set at p. There is a finite family $\mathcal{G}$ of hybrid sets at q such that $\mathrm{Reach}_U^q(p(A)) = \bigcup_{G \in \mathcal{G}} G$, and every $G = \boldsymbol{b}_G + P_G^{\langle \boldsymbol{c}_G, \cdot \rangle \le K_G} \in \mathcal{G}$ satisfies*

$$\textsc{norm}(G) \le \left(S + \textsc{norm}(A)\right)^{\mathcal{O}(1)}, \qquad \|P_G\|, \|\boldsymbol{c}_G\|, |P_G| \le \left(S + \|P_A\| + \|\boldsymbol{c}_A\|\right)^{\mathcal{O}(1)}.$$

Proof. The LPS representation is [BEF+21, Theorem 3.1], and the length bound is [BEF+21, Theorem 3.2]. The hybrid-source description is [CFF+26, Lemma C.11]. The bound $\mathrm{NORM}(G) \le (S + \mathrm{NORM}(A))^{\mathcal{O}(1)}$ is already established in the cited proof. It remains to show that $\|P_G\|$, $\|\boldsymbol{c}_G\|$, and $|P_G|$ have polynomial bounds independent of the source base $\boldsymbol{b}_A$ and the cost bound K_A. Put $R := S + \|P_A\| + \|\boldsymbol{c}_A\| + 1$.

By [CFF+26, Lemma C.1 and Theorem C.2], the original construction consists of a short prefix, a zigzag middle, and a short suffix. The first and last stages involve LPSs with constantly many cycles. Their output periods and costs are determined by homogeneous solutions of linear Diophantine systems, whose coefficients are independent of $\|\boldsymbol{b}_A\|$ and K_A. In [CFF+26, Lemma C.7], the input base affects only the right-hand side, and a homogeneous solution with coefficients $\boldsymbol{\lambda}$ for the input periods has cost $\langle \boldsymbol{c}_{\mathrm{in}}, \boldsymbol{\lambda} \rangle$. Since these systems have a constant number of equations, Corollary 3.4 gives

$$\|P_{\mathrm{out}}\|, \|\boldsymbol{c}_{\mathrm{out}}\| \le (S + \|P_{\mathrm{in}}\| + \|\boldsymbol{c}_{\mathrm{in}}\| + 1)^{\mathcal{O}(1)}.$$

The one-dimensional reductions used in [CFF+26, Lemmas C.6, C.8, and C.9] do not increase period norms. By [CFF+26, Lemma C.4], each one-turn return produces a finite family of arithmetic sets, represented by hybrid sets with at most one period of cost zero or one and with period norms satisfying

$$D_{i+1} \le \max\{D_i, Q(S)\},$$

where D_i bounds the period norms and Q is a fixed polynomial. Consequently, the period bound does not accumulate over the polynomially many returns.

Applying the constant-cycle bound to the final short suffix yields $\|P_G\|, \|\boldsymbol{c}_G\| \le R^{\mathcal{O}(1)}$. We remove zero periods and retain the least cost for repeated periods. Finally, since $P_G \subseteq \mathbb{N}^2 \setminus \{\mathbf{0}\}$, $|P_G| \le (\|P_G\| + 1)^2 \le R^{\mathcal{O}(1)}$. This proves the required bounds. □

E.8 Proofs of the Claims in Lemma 7.14

Whenever we apply Lemma 7.15 to an LPS Λ, we use the 2-VASS induced by Λ, realizing its paths and cycles with fresh states. This VASS has polynomial size in $|\Lambda| + \|\Lambda\|$, and it has the same reachability relation as Λ between the designated source and target. Thus the resulting hybrid sets represent only target vectors of runs captured by the corresponding LPS.

Claim 7.16. *If a cycle used by the chosen run has nonzero nonnegative effect, there exists a hybrid set H with a nonzero free period such that $\boldsymbol{z} \in H \subseteq \mathcal{S}_\Lambda$ and $\mathrm{NORM}(H) \le (D + \|\boldsymbol{x}\|)^{\mathcal{O}(1)}$.*

Proof of Claim. If a remaining cycle β has nonzero nonnegative effect $\boldsymbol{e}$, require one copy by replacing its starred expression with $\beta\beta^*$. This restricted scheme still contains the actual witness and has polynomial size. Choose a hybrid set containing its actual target by Lemma 7.15 and adjoin $\boldsymbol{e}$ as a free period. Every represented run executes an enabled copy of β. Further copies and the original suffix are enabled by monotonicity. The extra copies increase the exponent of the same starred block, so every member lies in $\mathcal{S}_\Lambda$, including when the original cycle has a nonzero internal drop. If G is the selected hybrid set, this construction is

$$H := G + \boldsymbol{e}^*, \qquad \boldsymbol{z} \in G \subseteq H, \qquad \boldsymbol{e} \ge \mathbf{0}, \quad \boldsymbol{e} \ne \mathbf{0}.$$

The added period has norm at most E and cost zero. Thus the first outcome has period norm and base norm polynomial in $D + \|\boldsymbol{x}\|$. ■

Claim 7.17. *There exists a run $\eta_-\beta^\kappa\eta_+$ from the encoded source to $\boldsymbol{z}$, with vector $\boldsymbol{a}$ immediately before β^κ, satisfying the following properties:*

1. *β is a cycle with $\mathrm{EFF}(\beta) = -u \cdot \boldsymbol{e}_j + v \cdot \boldsymbol{e}_k$, where $1 \le u, v \le E$.*
2. *The transition-effect sequences of η_- and η_+ are captured by LPSs Λ_- and Λ_+, respectively, obtained by cutting Λ after the selected cycle β, with that cycle included in Λ_-.*
3. *Every vector $\boldsymbol{w}$ along η_+ satisfies $\boldsymbol{w}(k) \ge \kappa v$.*

Proof of Claim. After deleting zero-effect cycles, every cycle with positive effect on coordinate k has negative effect on coordinate j. The R-large hypothesis gives

$$\boldsymbol{z}(k) > \|\boldsymbol{x}\| + R \ge \boldsymbol{x}(k) + \kappa E.$$

If every cycle increasing k occurred fewer than κ times, the total k-increase would be at most κE, counting all fixed paths. Choose the last block with at least κ copies of a cycle β increasing k, and write

$$\beta^\ell = \beta^{\ell-\kappa}\beta^\kappa, \qquad \mathrm{EFF}(\beta) = -u \cdot \boldsymbol{e}_j + v \cdot \boldsymbol{e}_k, \quad \ell \ge \kappa, \quad u, v \ge 1.$$

From any position after this block, the remaining k-effect is at most $\kappa(E - |\beta|\,\|T\|)$: later positive-effect cycles occur fewer than κ times, complete nonpositive-effect cycles contribute at most zero, and fixed paths and a possible partial cycle use only the scheme outside the selected block. Consequently every vector $\boldsymbol{w}$ on that suffix satisfies

$$\boldsymbol{w}(k) \ge \boldsymbol{z}(k) - \kappa(E - |\beta|\,\|T\|) > \kappa|\beta|\,\|T\| \ge \kappa v. \tag{26}$$

Let $\boldsymbol{a}$ be the vector just before the last κ copies of β. The prefix scheme Λ_- ending there includes the preceding starred copies of β, and Λ_+ is the unchanged suffix scheme after this block. Both are obtained from Λ without unrolling the remaining repetitions, so their sizes are polynomial in D and independent of κ. The actual prefix and suffix give η_- and η_+, and $u, v \le |\beta|\,\|T\| \le E$. ∎

Claim 7.18. *There is a translated family $\{H_s\}_{s=0}^{\kappa}$ with $\boldsymbol{z} \in H_0$ and direction $\boldsymbol{\delta} = u \cdot \boldsymbol{e}_j - v \cdot \boldsymbol{e}_k$. Its periods and costs have norm at most $D^{\mathcal{O}(1)}$, and its base has norm at most $(D + \|\boldsymbol{x}\| + \kappa)^{\mathcal{O}(1)}$. Every member is reachable from the encoded source at the fixed encoded target state.*

Proof of Claim. Let $\boldsymbol{d} := \mathrm{DROP}(\beta)$. Apply Lemma 7.15 to choose a hybrid set $A = \boldsymbol{b}_A + P_A^{\langle \boldsymbol{c}_A, \cdot \rangle \le K_A}$ containing $\boldsymbol{a}$ and consisting of vectors reachable through Λ_-. The validity of the actual β^κ gives the guards

$$\boldsymbol{a}(j) \ge g_j := \max\{\kappa u, \boldsymbol{d}(j) + (\kappa - 1)u\}, \qquad \boldsymbol{a}(k) \ge g_k := \boldsymbol{d}(k). \tag{27}$$

Fix its actual coefficients $\boldsymbol{\lambda}_*$. Freeze occurrences from the fixed representation of $\boldsymbol{a}$ into the base until $\overline{\boldsymbol{b}}_A(j) \ge g_j$ and $\overline{\boldsymbol{b}}_A(k) \ge g_k$. Choose $\boldsymbol{0} \le \boldsymbol{\theta} \le \boldsymbol{\lambda}_*$ and set

$$\overline{\boldsymbol{b}}_A := \boldsymbol{b}_A + P_A\boldsymbol{\theta}, \qquad \overline{K}_A := K_A - \langle \boldsymbol{c}_A, \boldsymbol{\theta} \rangle.$$

Let $\overline{A} := \overline{\boldsymbol{b}}_A + P_A^{\langle \boldsymbol{c}_A, \cdot \rangle \le \overline{K}_A}$. Since the full representation of $\boldsymbol{a}$ satisfies both inequalities and all periods are nonnegative, a greedy choice needs at most $g_j + g_k$ occurrences. Hence

$$\|\boldsymbol{\theta}\|_1 \le g_j + g_k, \qquad \overline{K}_A \ge 0, \qquad \left\|\overline{\boldsymbol{b}}_A\right\| \le \|\boldsymbol{b}_A\| + (g_j + g_k)\,\|P_A\|\,.$$

Thus $\boldsymbol{a} \in \overline{A} \subseteq A$. The base of $\overline{A}$ has norm polynomial in $D + \|\boldsymbol{x}\| + \kappa$, while its periods and costs satisfy the same bounds as those of A. For $\boldsymbol{w} \in \overline{A}$ and $0 \le t < \kappa$, these guards give

$$\big(\boldsymbol{w} + t \cdot \mathrm{EFF}(\beta)\big)(j) = \boldsymbol{w}(j) - tu \ge \boldsymbol{d}(j) + (\kappa - 1 - t)u \ge \boldsymbol{d}(j).$$

The other coordinate satisfies $\big(\boldsymbol{w} + t \cdot \mathrm{EFF}(\beta)\big)(k) = \boldsymbol{w}(k) + tv \ge \boldsymbol{d}(k)$. Hence every member enables β^q for each $0 \le q \le \kappa$, including all internal transitions of the cycle.

Constructing the Translated Family. Define the shifted source hybrid set

$$A^- := \overline{A} - \kappa u \cdot \boldsymbol{e}_j, \qquad \boldsymbol{a}^- := \boldsymbol{a} - \kappa u \cdot \boldsymbol{e}_j \in A^-.$$

Its base is nonnegative by (27). By (26), subtracting $\kappa v \cdot \boldsymbol{e}_k$ from η_+ gives a valid run from $\boldsymbol{a}^-$ to $\boldsymbol{z} - \kappa v \cdot \boldsymbol{e}_k$. Apply Lemma 7.15 with source A^- to select an output hybrid set $G = \boldsymbol{b}_G + P^{\langle \boldsymbol{c}, \cdot \rangle \le K_G}$ containing that target. Its base has norm polynomial in $D + \|\boldsymbol{x}\| + \kappa$, while Lemma 7.15 bounds its periods and costs without using the shifted source base. Define

$$H_s := G + su \cdot \boldsymbol{e}_j + (\kappa - s)v \cdot \boldsymbol{e}_k, \tag{28}$$

where $0 \le s \le \kappa$. Then $H_s = \boldsymbol{b} + s \cdot \boldsymbol{\delta} + P^{\langle \boldsymbol{c}, \cdot \rangle \le K_G}$, where $\boldsymbol{b} := \boldsymbol{b}_G + \kappa v \cdot \boldsymbol{e}_k$ and $\boldsymbol{\delta} := u \cdot \boldsymbol{e}_j - v \cdot \boldsymbol{e}_k$. Every base is nonnegative, $\boldsymbol{z} \in H_0$, and the periods, costs, and cost bound are those of G.

To verify the prefix witnesses, take $\boldsymbol{g} \in G$ reached by a run ρ_+ through Λ_+ from some $\boldsymbol{w}^- \in A^-$. The vector $\boldsymbol{w} := \boldsymbol{w}^- + \kappa u \cdot \boldsymbol{e}_j \in \overline{A}$ has a prefix ρ_- through Λ_- and satisfies the guards in (27). Execute $\kappa - s$ copies of β from $\boldsymbol{w}$ and translate the suffix witness by the nonnegative vector $su \cdot \boldsymbol{e}_j + (\kappa - s)v \cdot \boldsymbol{e}_k$. This gives

$$\boldsymbol{x} \xrightarrow{\rho_-} \boldsymbol{w} \xrightarrow{\beta^{\kappa - s}} \boldsymbol{w}^- + su \cdot \boldsymbol{e}_j + (\kappa - s)v \cdot \boldsymbol{e}_k \xrightarrow{\rho_+} \boldsymbol{g} + su \cdot \boldsymbol{e}_j + (\kappa - s)v \cdot \boldsymbol{e}_k \in H_s.$$

Thus every member of the family is reachable in the bounded-coordinate encoding. ∎

E.9 Proof of Lemma 7.19

Lemma 7.19 (R-large: terminal case)**.** *There exists $D \le R \le D^{\mathcal{O}(1)}$ such that the following holds. Let $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{\pi_0} \boldsymbol{r}(\boldsymbol{z}) \xrightarrow{\pi_1} \boldsymbol{q}(\boldsymbol{y})$ be a run in U whose target $\boldsymbol{q}(\boldsymbol{y})$ is backward C-coverable, where $\boldsymbol{s} = (\boldsymbol{p}, \boldsymbol{x})$ and $\boldsymbol{t} = (\boldsymbol{q}, \boldsymbol{y})$. Assume that π_0 is R-large and $i \in \mathrm{BD}(\pi_0, D)$. If the configuration $\boldsymbol{r}(\boldsymbol{w})$ is forward C-coverable whenever $\boldsymbol{w}(i) \ge \boldsymbol{z}(i)$ and $\boldsymbol{w}(j) \ge D$ for all $j \ne i$, then*

$$\ell_V(\boldsymbol{s}, \boldsymbol{t}) \le \big(D + \|\boldsymbol{x}\| + \|\boldsymbol{t}\|\big)^{\mathcal{O}(1)}.$$

Proof. By Lemma 7.14, a fixed polynomial $f(D)$ bounds the periods of every extracted translated family for all κ. Use the constant target family $J_\ell = \{\boldsymbol{y}\}$ for $0 \le \ell \le \kappa$, with $P' = \varnothing$, $K' = 0$, and $\boldsymbol{\delta}' = \boldsymbol{0}$. Every J_ℓ is backward C-coverable. As $\|P'\| = 0$, let g be a fixed nondecreasing polynomial such that the threshold in Lemma 7.12 satisfies $\kappa_0 \le g(D + \|P\|)$. Choose

$$\max\{2D, g(D + f(D))\} \le \kappa \le D^{\mathcal{O}(1)},$$

and then choose the sufficient R from Lemma 7.14:

$$D \le R \le (D + \kappa)^{\mathcal{O}(1)} \le D^{\mathcal{O}(1)}.$$

These choices are uniform over all suffixes.

Apply Lemma 7.14. If $\boldsymbol{z} \in H$ and H has a nonzero free period, use Lemma 7.10 to choose $\boldsymbol{h} \in H$ with polynomial norm and distance to $\boldsymbol{t}$. Together with the prefix bound in (17), this gives

$$\ell_V(\boldsymbol{s}, \boldsymbol{t}) \le \ell_V\big(\boldsymbol{s}, (\boldsymbol{r}, \boldsymbol{h})\big) + \ell_V\big((\boldsymbol{r}, \boldsymbol{h}), \boldsymbol{t}\big) \le (D + \|\boldsymbol{x}\| + \|\boldsymbol{t}\|)^{\mathcal{O}(1)}.$$

Otherwise $\boldsymbol{z} \in H_0$ for a translated family from Lemma 7.14. Our choice gives $\kappa \ge \kappa_0$ and $\kappa \ge 2D$. For the extracted direction $u \cdot \boldsymbol{e}_j - v \cdot \boldsymbol{e}_k$ with $u, v \ge 1$, choose $a := \lceil \max\{D - \boldsymbol{b}(j), 0\}/u \rceil \le D$. Then $\boldsymbol{b}_a(j) \ge D$, while $\boldsymbol{b}_\kappa \ge \boldsymbol{0}$ gives $\boldsymbol{b}_a(k) \ge (\kappa - a)v \ge D$. Since $\boldsymbol{\delta}$ is nonzero precisely on coordinates j, k, the definition of a translated family implies that every period has zero effect on coordinate i. Together with $\boldsymbol{z} \in H_0$, this gives $\boldsymbol{b}_a(i) = \boldsymbol{z}(i)$. The covering hypothesis makes H_a forward C-coverable. Together with the given suffix and J_0, this proves C-joinability. By Lemma 7.12, choose $\boldsymbol{h} \in H_r$ with polynomial norm and distance to $\boldsymbol{t}$, since the selected member of the constant target family is $\boldsymbol{y}$. Append its run to the prefix supplied by (17), obtaining

$$\ell_V(\boldsymbol{s}, \boldsymbol{t}) \le (D + \|\boldsymbol{x}\| + \|\boldsymbol{t}\|)^{\mathcal{O}(1)}.$$

□

F Proofs for Section 8

F.1 Proof of Lemma 8.1

Lemma 8.1. *Suppose $\boldsymbol{p}(\boldsymbol{x})$ is pumpable in U and $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{*} \boldsymbol{q}(\boldsymbol{y})$ in U. For $\boldsymbol{s} = (\boldsymbol{p}, \boldsymbol{x})$ and $\boldsymbol{t} = (\boldsymbol{q}, \boldsymbol{y})$,*

$$\ell_V(\boldsymbol{s}, \boldsymbol{t}) \le \big(D + \|\boldsymbol{s}\| + \|\boldsymbol{t}\|\big)^{\mathcal{O}(1)}.$$

Proof. Let $M := D + \|\boldsymbol{s}\| + \|\boldsymbol{t}\|$ and apply Lemma 7.9. If the extracted prefix reaches $\boldsymbol{q}(\boldsymbol{y})$, its pointwise length bound at $\boldsymbol{y}$ proves the claim. Otherwise, it gives a forward C-coverable hybrid set $G = \boldsymbol{a} + P^{\langle \boldsymbol{c}, \cdot \rangle \le K}$ at state $\boldsymbol{r}$, with $\boldsymbol{z} \in G$ and $\boldsymbol{r}(\boldsymbol{z}) \xrightarrow{*} \boldsymbol{q}(\boldsymbol{y})$ in the suffix of U, such that

$$\textsc{norm}(G) \le M^{\mathcal{O}(1)}, \qquad \ell_U\big(\boldsymbol{p}(\boldsymbol{x}), \boldsymbol{r}(\boldsymbol{g})\big) \le (M + \|\boldsymbol{g}\|)^{\mathcal{O}(1)} \quad \text{for all } \boldsymbol{g} \in G. \tag{29}$$

Let $\widehat{P} := \{(0, 0, \boldsymbol{d}) : \boldsymbol{d} \in P\} \subseteq \mathbb{N}^5$. Consider the augmented VAS $V[\widehat{P}]$ and the suffix of U augmented with a self-loop of effect $\boldsymbol{d}$ at $\boldsymbol{r}$ for every $\boldsymbol{d} \in P$. The base $\boldsymbol{r}(\boldsymbol{a})$ is forward C-coverable there and reaches $\boldsymbol{q}(\boldsymbol{y})$ by first generating $\boldsymbol{z}$. By Corollary 4.7 and Theorem 7.1, there is a run ρ from $(\boldsymbol{r}, \boldsymbol{a})$ to $\boldsymbol{t}$ in $V[\widehat{P}]$ of length $L \le M^{\mathcal{O}(1)}$. Remove the added periods and handle the cost bounds as in the proof of Lemma 6.17, with singleton target $\{\boldsymbol{y}\}$. If absorption is needed, it keeps $\boldsymbol{z}$ in the selected hybrid set, increases the base, and leaves the free periods unchanged, so the base remains forward C-coverable by monotonicity. After at most one repetition of the preceding construction, this yields $\boldsymbol{g} \in G$ with norm and distance to $\boldsymbol{t}$ polynomial in M. Prepend its prefix witness from (29). □

F.2 Proof of Lemma 8.2

Lemma 8.2. *Let $H = \boldsymbol{a} + P^{\langle \boldsymbol{c}, \cdot \rangle \le K}$ and J be hybrid sets at the source state $\boldsymbol{p}$ and target state $\boldsymbol{q}$ of U, respectively. Write P_0 for the free periods of H. Suppose that $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{*} \boldsymbol{q}(\boldsymbol{y})$ in U for some $\boldsymbol{x} \in H$ and $\boldsymbol{y} \in J$. If $\boldsymbol{p}(\boldsymbol{a})$ is forward C-coverable or pumpable in $U[P_0]$, then there exist $\boldsymbol{h} \in H$ and $\boldsymbol{h}' \in J$ such that*

$$\|\boldsymbol{h}\|, \ \|\boldsymbol{h}'\|, \quad \ell_V\big((\boldsymbol{p}, \boldsymbol{h}), (\boldsymbol{q}, \boldsymbol{h}')\big) \le \big(D + \textsc{norm}(H) + \textsc{norm}(J)\big)^{\mathcal{O}(1)}. \tag{21}$$

Proof. Let $M := D + \mathrm{NORM}(H) + \mathrm{NORM}(J)$ and write $J = \boldsymbol{b} + P'^{\langle \boldsymbol{c}', \cdot \rangle \le K'}$. Fix representations of the connected members $\boldsymbol{x} \in H$ and $\boldsymbol{y} \in J$, together with a witnessing run in U. In $W := U[P, P']$, the positive period loops at $\boldsymbol{p}$ generate $\boldsymbol{x}$ from $\boldsymbol{a}$, and the negative period loops at $\boldsymbol{q}$ reduce $\boldsymbol{y}$ to $\boldsymbol{b}$. Thus $\boldsymbol{p}(\boldsymbol{a}) \xrightarrow{*} \boldsymbol{q}(\boldsymbol{b})$ in W. Let $V^{\pm}$ be V with the lifted transitions $(0, 0, \boldsymbol{d})$ for $\boldsymbol{d} \in P$ and $(0, 0, -\boldsymbol{d}')$ for $\boldsymbol{d}' \in P'$ adjoined. The source condition remains valid in W, and Corollary 4.7 preserves C for $V^{\pm}$. Since the period sets are subsets of $\mathbb{N}^3$, their cardinalities and $\mathrm{SIZE}(V^{\pm})$ are polynomial in M. The structural threshold for the augmented instance can therefore be chosen at most $(B + C + \mathrm{SIZE}(V^{\pm}))^{\mathcal{O}(1)} \le M^{\mathcal{O}(1)}$. Apply Theorem 7.1 in the forward coverable case and Lemma 8.1 in the pumpable case to obtain

$$(\boldsymbol{p}, \boldsymbol{a}) \xrightarrow{\rho} (\boldsymbol{q}, \boldsymbol{b}) \quad \text{in } V^{\pm}, \qquad |\rho| \le L \le M^{\mathcal{O}(1)}.$$

Remove the added periods and handle the cost bounds as in the proof of Lemma 6.17. Each absorption preserves the fixed connected member on that side, increases its base, and leaves its free periods unchanged. Hence monotonicity preserves the source's forward C-coverability or pumpability in the system augmented with its free periods, so the corresponding short-run bound remains applicable. Each side is absorbed at most once, and the same polynomial estimates give members $\boldsymbol{h} \in H$ and $\boldsymbol{h}' \in J$ satisfying (21). □

F.3 Proof of Lemma 8.3

Lemma 8.3. *Let $H, J \subseteq \mathbb{N}^3$ be hybrid sets with $H \cap J \neq \emptyset$. There is $\boldsymbol{h} \in H \cap J$ such that*

$$\|\boldsymbol{h}\| \le \big(D + \mathrm{NORM}(H) + \mathrm{NORM}(J)\big)^{\mathcal{O}(1)}.$$

Proof. Let $M := D + \mathrm{NORM}(H) + \mathrm{NORM}(J)$ and write

$$H = \boldsymbol{b} + P^{\langle \boldsymbol{c}, \cdot \rangle \le K}, \qquad J = \boldsymbol{b}' + P'^{\langle \boldsymbol{c}', \cdot \rangle \le K'}, \qquad \boldsymbol{z} \in H \cap J.$$

Fix coefficients $\boldsymbol{\lambda}_*, \boldsymbol{\mu}_*$ representing $\boldsymbol{z}$ in both hybrid sets and choose a componentwise-minimal solution below $(\boldsymbol{\lambda}_*, \boldsymbol{\mu}_*)$ of

$$\begin{pmatrix} P & -P' \end{pmatrix} \begin{pmatrix} \boldsymbol{\lambda} \\ \boldsymbol{\mu} \end{pmatrix} = \boldsymbol{b}' - \boldsymbol{b}, \qquad (\boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathbb{N}^{|P| + |P'|}.$$

Any smaller solution would also lie below the fixed pair, so this solution is globally minimal. Nonnegative costs give

$$\langle \boldsymbol{c}, \boldsymbol{\lambda} \rangle \le \langle \boldsymbol{c}, \boldsymbol{\lambda}_* \rangle \le K, \qquad \langle \boldsymbol{c}', \boldsymbol{\mu} \rangle \le \langle \boldsymbol{c}', \boldsymbol{\mu}_* \rangle \le K'.$$

Let $m := |P| + |P'|$ and $A_0 := \max\{\|P\|, \|P'\|\}$. By Theorem 3.3, the three-row system satisfies

$$\|\boldsymbol{\lambda}\|_1 + \|\boldsymbol{\mu}\|_1 \le \big(1 + mA_0 + \|\boldsymbol{b}' - \boldsymbol{b}\|\big)^3 \le M^{\mathcal{O}(1)}.$$

Thus $\boldsymbol{h} := \boldsymbol{b} + P\boldsymbol{\lambda} = \boldsymbol{b}' + P'\boldsymbol{\mu}$ has the required properties. □

F.4 Proof of Lemma 8.4

Lemma 8.4. *Let $\mathcal{J} = \{J_\ell\}_{0\le\ell\le\kappa}$ be a translated family at the target state $\boldsymbol{r}'$ of U with*

$$J_\ell = \boldsymbol{b}' + \ell \cdot \boldsymbol{\delta}' + P'^{\langle \boldsymbol{c}', \cdot\rangle \le K'}.$$

Suppose that J_a is backward C-coverable in U for some $0 \le a \le \min\{D, \kappa\}$. Suppose that $\boldsymbol{p}(\boldsymbol{x}) \xrightarrow{*} \boldsymbol{r}'(\boldsymbol{z}')$ in U for some $\boldsymbol{z}' \in J_0$. There exists $\kappa_0 \le \left(D + \|P'\|\right)^{\mathcal{O}(1)}$ such that, if $\kappa \ge \kappa_0$, then some $0 \le \ell \le \kappa_0$ and $\boldsymbol{h}' \in J_\ell$ satisfy

$$\left\|\boldsymbol{h}'\right\|, \quad \ell_V(\boldsymbol{s}, (\boldsymbol{r}', \boldsymbol{h}')) \le \left(D + \|\boldsymbol{s}\| + \textsc{norm}(J_0) + \left\|\boldsymbol{\delta}'\right\|\right)^{\mathcal{O}(1)}, \qquad \boldsymbol{s} = (\boldsymbol{p}, \boldsymbol{x}). \tag{22}$$

Proof. We follow Theorem 7.1, keeping $\mathcal{J}$ fixed. By Lemma 7.14, the periods of every translated family extracted from a source-side prefix have norm at most $D^{\mathcal{O}(1)}$, independently of the prescribed translation range. Thus Lemma 7.12 allows us to choose a uniform sufficient joining threshold $2D \le \kappa_0 \le (D + \|P'\|)^{\mathcal{O}(1)}$ before choosing R. Assume $\kappa \ge \kappa_0$, restrict $\mathcal{J}$ to indices at most κ_0, and choose $D \le R \le (D + \kappa_0)^{\mathcal{O}(1)}$ by Lemma 7.14. The member J_a remains available because $a \le D \le \kappa_0$, and its backward covering run stays in the terminal component of every suffix considered below.

Let $M := D + \textsc{norm}(J_0) + \|\boldsymbol{\delta}'\|$. For $k \ge 1$, let $L_k(N)$ bound both the norm of the selected member of $\mathcal{J}$ and its reachability distance in V from a source of norm at most N, uniformly over the relevant sequential subinstances with at most k components. Choose a fixed nondecreasing polynomial F large enough for the bounds below.

Base Case: $k = 1$. If $\boldsymbol{p}(\boldsymbol{x})$ is forward C-coverable, apply Lemma 8.2 with $H = \{\boldsymbol{x}\}$ and $J = J_0$. Otherwise, Lemma 7.2 captures the entire witnessing run to $\boldsymbol{r}'(\boldsymbol{z}')$ in a 2-VASS of size at most $D\,\textsc{size}(U)$, since there is no later component. By Lemma 7.15, its reachable vectors at the encoded target state include a hybrid set H containing $\boldsymbol{z}'$ with $\textsc{norm}(H) \le (D + \|\boldsymbol{x}\|)^{\mathcal{O}(1)}$, after restoring the bounded coordinate. Since $\boldsymbol{z}' \in H \cap J_0$, Lemma 8.3 selects a small $\boldsymbol{h}' \in H \cap J_0$. Apply Lemma 7.15 to reach this member in the same encoding and lift the run to V. Both cases give $L_1(N) \le F(M + N)$.

For $k \ge 1$, assume that L_k has been constructed, and consider a witnessing run in a subinstance with at most $k + 1$ components and $\|\boldsymbol{s}\| \le N$. When the source is not forward C-coverable, apply Lemma 7.2 and write the witness as $\pi_0\pi_1$, where π_0 ends at $\boldsymbol{r}(\boldsymbol{z})$ and some coordinate i stays below D throughout π_0. We distinguish four terminal cases and one inductive case.

Terminal Cases.

1. *The source is forward C-coverable.* Apply Lemma 8.2 with $H = \{\boldsymbol{x}\}$ and $J = J_0$.

2. *The bounded-coordinate prefix reaches $\boldsymbol{r}'(\boldsymbol{z}')$.* Here $\pi_1 = \varepsilon$. Use Lemmas 7.15 and 8.3 to select and reach a small member of $H \cap J_0$, exactly as in the base case.

3. *An R-large prefix yields a nonzero free period.* Apply Lemma 7.14 with range κ_0 to obtain a hybrid set H with $\boldsymbol{z} \in H$ and a nonzero free period. Its base is pumpable after adjoining the free periods. The unchanged suffix connects H to J_0, so Lemma 8.2 selects small members $\boldsymbol{h} \in H$ and $\boldsymbol{h}' \in J_0$ connected by a short run in V. Prepend the run to $\boldsymbol{h}$ supplied by (17).

4. *An R-large prefix yields a translated family.* The other outcome of Lemma 7.14 gives $\mathcal{H} = \{H_r\}_{r=0}^{\kappa_0}$ with $\boldsymbol{z} \in H_0$. Since $\kappa_0 \ge 2D$, the nonnegative bases and the stopping condition of Lemma 7.2

give a forward C-coverable H_b for some $b \le D$, by the same argument as in Appendix E.9. Together with the unchanged suffix from $\boldsymbol{r}(\boldsymbol{z})$ to $\boldsymbol{r}'(\boldsymbol{z}')$ and the backward C-coverability of J_a, this proves C-joinability of $\mathcal{H}$ and $\mathcal{J}$ in that suffix subinstance. Apply Lemma 7.12 to select small members $\boldsymbol{h} \in H_r$ and $\boldsymbol{h}' \in J_\ell$ connected by a short run in V, and prepend the prefix from (17).

In every terminal case, both $\|\boldsymbol{h}'\|$ and the concatenated length are at most $F(M+N)$.

Inductive Case. The remaining case has an R-small prefix π_0 ending at a later component entry $\boldsymbol{r}(\boldsymbol{z})$. For $\boldsymbol{u} := (\boldsymbol{r}, \boldsymbol{z})$, the same estimate as in (19) gives $\|\boldsymbol{u}\| \le N + R$. By Lemma 7.15, replace π_0 by a run in V to this same configuration of length at most $F(M+N+R)$. The unchanged suffix has at most k components, reaches the same $\boldsymbol{z}' \in J_0$, and preserves the backward covering run for J_a. Apply the induction hypothesis with the same target family and prepend the shortened prefix. The selected-member norm and total length therefore satisfy the recurrence in (20), with F in place of P. Since $k \le (B+1)^2$ and $R \le M^{\mathcal{O}(1)}$, the same additive-growth estimate as in Theorem 7.1 proves (22), with degree independent of k. □

F.5 Proof of Proposition 6.5

Proposition 6.5. *Let V be a 5-VAS and let $(U, \boldsymbol{p}(\boldsymbol{x}), \boldsymbol{q}(\boldsymbol{y}))$ be a reachable instance specified above. Then*

$$\ell_V(\boldsymbol{s}, \boldsymbol{t}) \le \big(B + \mathrm{SIZE}(V, \boldsymbol{s}, \boldsymbol{t})\big)^{\mathcal{O}(1)}. \tag{4}$$

Proof. Choose C by Theorem 4.2 and D uniformly for sequential subinstances in both orientations by Lemma 7.2. The period bound in Lemma 7.14 is independent of the prescribed range, so Lemma 8.4 gives a uniform sufficient cutoff $\kappa_0 \le D^{\mathcal{O}(1)}$ for all extracted families in the reversed system. Choose $\max\{2D, \kappa_0\} \le \kappa \le D^{\mathcal{O}(1)}$, then choose $D \le R \le (D+\kappa)^{\mathcal{O}(1)}$ by the dichotomy. By Corollary 4.7, the same C is valid under reversal. We follow the additive-growth induction in the proof of Lemma 8.4, now with the fixed target $\boldsymbol{q}(\boldsymbol{y})$ and a witnessing run π. If the source is not forward C-coverable, use Lemma 7.2 to write $\pi = \pi_0\pi_1$, with π_0 ending at $\boldsymbol{r}(\boldsymbol{z})$.

Terminal Cases.

1. *The source is forward C-coverable.* Apply Theorem 7.1.
2. *The bounded-coordinate prefix reaches the target.* If $\pi_1 = \varepsilon$, shorten π_0 in its 2-VASS encoding by Lemma 7.15 and lift the run to V.
3. *An R-large prefix yields a nonzero free period.* The dichotomy gives a hybrid set H containing $\boldsymbol{z}$, whose base is pumpable after adjoining the free periods. Apply Lemma 8.2 to H and $\{\boldsymbol{y}\}$ in the suffix subinstance, and prepend the run to its selected member supplied by (17).
4. *An R-large prefix yields a translated family.* Write $\mathcal{H} = \{H_r\}_{r=0}^{\kappa}$ for the extracted family, with $\boldsymbol{z} \in H_0$. Since $\kappa \ge 2D$, the stopping condition gives a forward C-coverable H_b for some $b \le D$, as in Appendix E.9. In the reversed suffix, H_b is backward C-coverable and $\boldsymbol{q}(\boldsymbol{y})$ reaches $\boldsymbol{r}(\boldsymbol{z})$ with $\boldsymbol{z} \in H_0$. Thus Lemma 8.4 applies to this fixed target family and selects a polynomial-norm $\boldsymbol{h} \in H_r$ with a short run from $\boldsymbol{t}$ to $(\boldsymbol{r}, \boldsymbol{h})$ in $-V$. Reverse that run and prepend the prefix from (17).

The first two cases cover a single-component instance. The extracted hybrid norms and translation directions are polynomially bounded in $D + \|\boldsymbol{s}\|$, so every terminal case gives a run of length at most $(D + \|\boldsymbol{s}\| + \|\boldsymbol{t}\|)^{\mathcal{O}(1)}$.

The remaining R-small case is identical to the inductive step of Lemma 8.4: the new source norm is at most $\|\boldsymbol{s}\| + R$, and the unchanged suffix has fewer components. Since there are at most $(B+1)^2$ components and $D, R \leq (B + \text{SIZE}(V))^{\mathcal{O}(1)}$, the same recurrence (20) proves (4). □

References

[BEF+21] Michael Blondin, Matthias Englert, Alain Finkel, Stefan Göller, Christoph Haase, Ranko Lazić, 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.

[CCM+24] Dmitry Chistikov, Wojciech Czerwiński, Filip Mazowiecki, Łukasz Orlikowski, Henry Sinclair-Banks, and Karol Węgrzycki. The tractability border of reachability in simple vector addition systems with states. In *65th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2024, Chicago, IL, USA, October 27-30, 2024*, pages 1332–1354. IEEE, 2024.

[CFF+26] Weijun Chen, Bo Fu, Yuxi Fu, Huan Long, Chengfeng Xue, Qizhe Yang, and Yangluo Zheng. 3-VASS reachability is in EXPSPACE. *CoRR*, abs/2607.14983, 2026.

[CFZ26] Weijun Chen, Yuxi Fu, and Yangluo Zheng. Improving reachability in vector addition systems through pumpability. In Ana Sokolova and Patrick Totzke, editors, *37th International Conference on Concurrency Theory, CONCUR 2026, Liverpool, UK, September 1-4, 2026*, volume 391 of *LIPIcs*, pages 24:1–24:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2026. Full version: arXiv:2604.24095.

[CGO+26] Wojciech Czerwiński, Roland Guttenberg, Łukasz Orlikowski, Henry Sinclair-Banks, and Yangluo Zheng. Exploring VASS parameterised by geometric dimension. In Sayan Bhattacharya, Danupon Nanongkai, Michael Benedikt, and Gabriele Puppis, editors, *53rd International Colloquium on Automata, Languages, and Programming, ICALP 2026, Royal Holloway, University of London, Egham, United Kingdom, July 7-10, 2026*, volume 374 of *LIPIcs*, pages 177:1–177:25. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2026.

[CJL+23] Wojciech Czerwiński, Ismaël Jecker, Sławomir Lasota, Jérôme Leroux, and Łukasz Orlikowski. New Lower Bounds for Reachability in Vector Addition Systems. In Patricia Bouyer and Srikanth Srinivasan, editors, *43rd IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2023, IIIT Hyderabad, Telangana, India, December 18-20, 2023*, volume 284 of *LIPIcs*, pages 35:1–35:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.

[CJLO25] Wojciech Czerwiński, Ismaël Jecker, Sławomir Lasota, and Łukasz Orlikowski. Reachability in 3-VASS Is Elementary. In Keren Censor-Hillel, Fabrizio Grandoni, Joël Ouaknine, and Gabriele Puppis, editors, *52nd International Colloquium on Automata, Languages, and Programming, ICALP 2025, Aarhus, Denmark, July 8-11, 2025*, volume

334 of *LIPIcs*, pages 153:1–153:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025.

[CO21] Wojciech Czerwiński and Łukasz 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.

[Esp96] Javier Esparza. Decidability and complexity of Petri net problems - an introduction. In Wolfgang Reisig and Grzegorz Rozenberg, editors, *Lectures on Petri Nets I: Basic Models, Advances in Petri Nets, the volumes are based on the Advanced Course on Petri Nets, held in Dagstuhl, September 1996*, volume 1491 of *Lecture Notes in Computer Science*, pages 374–428. Springer, 1996.

[FYZ24] Yuxi Fu, Qizhe Yang, and Yangluo Zheng. Improved Algorithm for Reachability in d-VASS. In Karl Bringmann, Martin Grohe, Gabriele Puppis, and Ola Svensson, editors, *51st International Colloquium on Automata, Languages, and Programming, ICALP 2024, Tallinn, Estonia, July 8-12, 2024*, volume 297 of *LIPIcs*, pages 136:1–136:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.

[GS80] V. S. Grinberg and S. V. Sevast'yanov. Value of the Steinitz constant. *Functional Analysis and Its Applications*, 14:125–126, 1980.

[HKOW09] 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.

[HP79] John E. Hopcroft and Jean-Jacques Pansiot. On the Reachability Problem for 5-Dimensional Vector Addition Systems. *Theor. Comput. Sci.*, 8:135–159, 1979.

[KL27] Łukasz Kamiński and Sławomir Lasota. Reachability in 3-VAS. In Vesa Halava and Jarkko Kari, editors, *Reachability Problems*, pages 114–123, Cham, 2027. Springer Nature Switzerland.

[KM69] Richard M. Karp and Raymond E. Miller. Parallel program schemata. *J. Comput. Syst. Sci.*, 3(2):147–195, 1969.

[Kos82] S. Rao Kosaraju. Decidability of Reachability in Vector Addition Systems (Preliminary Version). In Harry R. Lewis, Barbara B. Simons, Walter A. Burkhard, and Lawrence H. Landweber, editors, *Proceedings of the 14th Annual ACM Symposium on Theory of Computing, May 5-7, 1982, San Francisco, California, USA*, pages 267–281. ACM, 1982.

[Lam92] Jean-Luc Lambert. A Structure to Decide Reachability in Petri Nets. *Theor. Comput. Sci.*, 99(1):79–104, 1992.

[Ler21] 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.

[LS04] 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.

[LS19] 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.

[May81] Ernst W. Mayr. An algorithm for the general Petri net reachability problem. In Ronald L. Rivest, Georgia I. Davida, Walter A. Burkhard, and Richard J. Lipton, editors, *Proceedings of the 13th Annual ACM Symposium on Theory of Computing, May 11-13, 1981, Milwaukee, Wisconsin, USA*, pages 238–246. ACM, 1981.

[Mur89] Tadao Murata. Petri nets: Properties, analysis and applications. *Proc. IEEE*, 77(4):541–580, 1989.

[Pot91] 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.

[Rac78] Charles Rackoff. The Covering and Boundedness Problems for Vector Addition Systems. *Theor. Comput. Sci.*, 6:223–231, 1978.

[Sch16] Sylvain Schmitz. Complexity Hierarchies beyond Elementary. *ACM Trans. Comput. Theory*, 8(1):3:1–3:36, 2016.

[Ste13] Ernst Steinitz. Bedingt konvergente Reihen und konvexe Systeme. *Journal für die reine und angewandte Mathematik*, 1913(143):128–176, 1913.

[vL74] Jan van Leeuwen. A partial solution to the reachability-problem for vector-addition systems. In Robert L. Constable, Robert W. Ritchie, Jack W. Carlyle, and Michael A. Harrison, editors, *Proceedings of the 6th Annual ACM Symposium on Theory of Computing, April 30 - May 2, 1974, Seattle, Washington, USA*, pages 303–309. ACM, 1974.

[Zhe25] Yangluo Zheng. Reachability in Vector Addition System with States Parameterized by Geometric Dimension. In Patricia Bouyer and Jaco van de Pol, editors, *36th International Conference on Concurrency Theory, CONCUR 2025, Aarhus, Denmark, August 26-29, 2025*, volume 348 of *LIPIcs*, pages 38:1–38:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025.

[ZZ94] Richard Zurawski and MengChu Zhou. Petri Nets and Industrial Applications: A Tutorial. *IEEE Trans. Ind. Electron.*, 41(6):567–583, 1994.