

New Pumping Technique for 2-Dimensional VASS

Wojciech Czerwiński 

University of Warsaw, Poland
wczerin@mimuw.edu.pl

Sławomir Lasota 

University of Warsaw, Poland
sl@mimuw.edu.pl

Christof Löding

RWTH Aachen, Germany
loeding@informatik.rwth-aachen.de

Radosław Piórkowski 

University of Warsaw, Poland
r.piorowski@mimuw.edu.pl

Abstract

We propose a new pumping technique for 2-dimensional vector addition systems with states (2-VASS) building on natural geometric properties of runs. We illustrate its applicability by reproving an exponential bound on the length of the shortest accepting run, and by proving a new pumping lemma for languages of 2-VASS. The technique is expected to be useful for settling questions concerning languages of 2-VASS, e.g., for establishing decidability status of the regular separability problem.

2012 ACM Subject Classification Theory of computation → Models of computation; Theory of computation → Formal languages and automata theory

Keywords and phrases vector addition systems with states, pumping, decidability

Digital Object Identifier 10.4230/LIPIcs.MFCS.2019.62

Related Version <https://arxiv.org/abs/1906.10494>

Funding This research was supported by Polish NCN grants “Separation problems in automata theory” (2016/21/D/ST6/01376), “Automatic analysis of concurrent systems” (2017/27/B/ST6/02093) and “Symbolic computations on first-order definable objects” (2016/21/B/ST6/01505).

1 Introduction

Vector addition systems [8] are a widely accepted model of concurrency equivalent to Petri nets. Another equivalent model, called vector addition systems with states (VASS) [7], is an extension of finite automata with integer counters, on which the transitions can perform operations of increment or decrement (but no zero tests), with the proviso that counter values are 0 initially and must stay non-negative along a run. The number of counters d defines the *dimension* of a VASS. For brevity, we call a VASS of dimension d a d -VASS. Formally, every transition of a d -VASS V has adjoined a vector $v \in \mathbb{Z}^d$ describing the effect of executing this transition on counter values; thus a transition is a triple $(q, v, q') \in Q \times \mathbb{Z}^d \times Q$, where Q is the set of control states of V . A finite *path*, i.e., a sequence of transitions of the form $\pi = (q_0, v_1, q_1), (q_1, v_2, q_2), \dots, (q_{n-1}, v_n, q_n)$, induces a run if the counter values stay non-negative, i.e., $v_1 + \dots + v_i \in \mathbb{N}^d$ for every i .

In this paper we concentrate on *pumping*, i.e., techniques exploiting repetitions of states in runs. Pumping is an ubiquitous phenomenon which typically provides valuable tools in proving short run properties, or showing language inexpressibility results. It seems to be particularly relevant in case of VASS, as even the core of the seminal decision procedure for



© Wojciech Czerwiński, Sławomir Lasota, Christof Löding, and Radosław Piórkowski; licensed under Creative Commons License CC-BY

44th International Symposium on Mathematical Foundations of Computer Science (MFCS 2019).

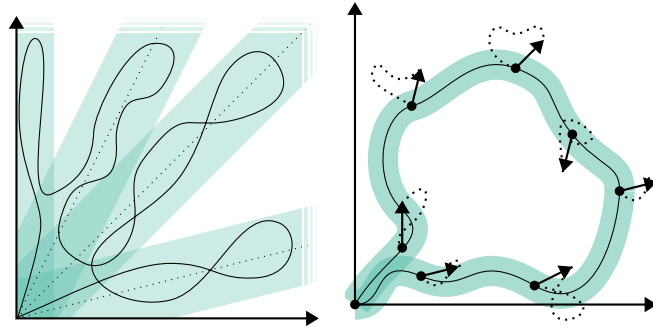
Editors: Peter Rossmanith, Pinar Heggernes, and Joost-Pieter Katoen; Article No. 62; pp. 62:1–62:14

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

the reachability problem in VASS by Mayr and Kosaraju [11, 9] is fundamentally based on pumping: briefly speaking, the decision procedure decomposes a VASS into a finite number of VASS, each of them admitting a property that every path can be pumped up so that it induces a run. Pumping techniques are used even more explicitly when dealing with subclasses of VASS of bounded dimension. The PSPACE upper bound for the reachability problem in 2-VASS [2] relies on various un-pumping transformations of an original run, leading to a simple run of at most exponential length, in the form of a short path with adjoined short disjoint cycles. A smart surgery on those simple runs was also used to obtain a stronger upper bound (NL) in case when the transition effects are represented in unary [6]. Un-pumping is also used in [3] to provide a quadratic bound on the length of the shortest run for 1-VASS, also known as one counter automata without zero tests, and for unrestricted one counter automata. See also [1, 10] for pumping techniques in one counter automata.



■ **Figure 1** Thin (above) and thick run (below). Points correspond to counter values, and control states along a run are ignored.

Contribution. The above-mentioned techniques are mostly oriented towards reachable sets, and henceforth may ignore certain runs as long as the reachable set is preserved. In consequence, they are not very helpful in solving decision problems formulated in terms of the whole language accepted by a VASS, like the regular separability problem (cf. the discussion below). Our primary objective is to design a pumping infrastructure applicable to *every* run of a 2-VASS. Therefore, as our main technical contribution we perform a thorough classification of runs, in the form of a dichotomy (see the illustrations on the right): for every run π of a 2-VASS, whose initial and final values of both counters are 0,

- either π is *thin*, by which we mean that the counter values along the run stay within belts, whose direction and width are all bounded polynomially in the number of states and the largest absolute value of vectors of the 2-VASS;
- or π is *thick*, by which we mean that a number of cycles is enabled along the run, the effect vectors of these cycles span (slightly oversimplifying) the whole plane, and furthermore the lengths of cycles and the extremal factors of π are all bounded polynomially in M and exponentially in n . (For the sake of simplicity some details are omitted here; the fully precise statement of the dichotomy is Theorem 3.1 in Section 3).

The dichotomy immediately entails a pumping lemma for 2-VASS by using, essentially, the pumping scheme of 1-VASS in case of thin runs, and the cycles enabled along a run in case of thick runs (cf. Theorem 4.1). As a more subtle application of the dichotomy, we derive an alternative proof of the exponential run property (shown originally in [2]), which immediately implies PSPACE-membership of the reachability problem (cf. Theorem 4.2).

Further applications. We envisage other possible applications of the dichotomy. One important case can be the regular separability problem: given two *labeled* 2-VASS V_1 and V_2 , decide if there is a regular language separating languages of V_1 and V_2 , i.e., including one of them and disjoint from the other. The problem is decidable in PSPACE for 1-VASS [5] while the decidability status for 2-VASS is still open. A cornerstone of the decision procedure of [5] is a well-behaved over-approximation of a language of a 1-VASS V by a sequence of regular languages $(V_n)_{n \in \mathbb{N}}$, where the precision of approximation increases with increasing n . In case of 1-VASS, the language V_n is obtained by abstraction of V modulo n ; on the other hand, as argued in [5], the very same approach necessarily fails for dimensions larger than 1. It seems that our dichotomy classification of runs of a 2-VASS prepares the ground for the right definition of abstraction V_n modulo n . Indeed, intuitively speaking, as long as the run stays within belts, 1-dimensional counting modulo n along the direction of a belt is sufficient; otherwise, a 2-dimensional abstraction modulo n can be applied as soon as a sufficient number of pumpable cycles has been identified along a run.

As our approach builds on natural geometric properties of runs, we believe that it can be generalized to dimensions larger than 2. However, one should not expect efficient length bounds from this generalization itself, as already in dimension 3 the prefix of a run preceding the first pumpable cycle has non-elementary length (the length can be as large as tower of n exponentials in the composition of n copies of the Hopcroft and Pansiot example [7]).

2 Preliminaries

2-dimensional vector addition systems with states. We use standard symbols $\mathbb{Q}, \mathbb{Z}, \mathbb{N}$ for the sets of rationals, integers, and non-negative integers, respectively. Whenever convenient we use subscripts to specify subsets, e.g., $\mathbb{Q}_{\geq 0}$ for non-negative rationals. We refer to elements of \mathbb{Z}^2 briefly as *vectors*. *Non-negative* vectors are elements of \mathbb{N}^2 , and *positive* vectors are elements of $\mathbb{Z}_{>0}^2$. A vector with only non-negative coordinates and at least one positive coordinate is called *semi-positive*; it is either positive, or *vertical* of the form $(0, a)$, or *horizontal* of the form $(a, 0)$, for $a \in \mathbb{Z}_{>0}$.

A 2-dimensional *vector addition system with states* (2-VASS) V consists of a finite set of control states Q and a finite set of transitions $T \subseteq Q \times \mathbb{Z}^2 \times Q$. We refer to the vector v as the *effect* of a transition (p, v, q) . A *path* in V from control state p to q is a sequence of transitions $\pi = (q_0, v_1, q_1), (q_1, v_2, q_2), \dots, (q_{n-1}, v_n, q_n) \in T^*$ where $p = q_0$ and $q = q_n$; it is called a *cycle* whenever the starting and ending control states coincide ($q_0 = q_n$). The *effect* of a path is defined as $\text{eff}(\pi) = v_1 + \dots + v_n \in \mathbb{Z}^2$, and its *length* is n . A cycle is called non-negative, semi-positive or positive, if its effect is so.

A *configuration* of V is an element of $\text{Conf} = Q \times \mathbb{N}^2$. A transition $t = (p, v, q)$ is *enabled* in a configuration $c = (p', u)$ if $p = p'$ and $u + v \in \mathbb{N}^2$. Analogously, a path π is enabled in a configuration $c = (p', u)$ if $q_0 = p'$ and $u_i = u + v_1 + \dots + v_i \in \mathbb{N}^2$ for every i . In such case we say that π induces a *run* of the form


$$\rho = (c_0, t_1, c_1), (c_1, t_2, c_2), \dots, (c_{n-1}, t_n, c_n) \in (\text{Conf} \times T \times \text{Conf})^*$$

with intermediate configurations $c_i = (q_i, u_i)$, from the *source* configuration $\text{src}(\rho) = c_0$ to the *target* one $\text{trg}(\rho) = c_n$. If the source configuration c_0 is clear from the context, we do not distinguish between a path enabled in c_0 and a run with source c_0 , and simply say that the path *is* the run. A $(0, 0)$ -*run* is a run whose source and target are $(0, 0)$ -*configurations*, i.e., a configuration whose vector is $(0, 0)$.

We will sometimes relax the non-negativeness requirement on some coordinates: For $j \in \{1, 2\}$, we say that a path π is $\{j\}$ -enabled in a configuration $c = (p', u)$ if $q_0 = p'$ and $(u + v_1 + \dots + v_i)[j] \in \mathbb{N}$ for every i . We also say that π is \emptyset -enabled in c if just $q_0 = p'$.

The reversal of a 2-VASS $V = (Q, T)$, denoted $\text{rev}(V)$, is a 2-VASS with the same control states and with transitions $\{(q, -v, p) \mid (p, v, q) \in T\}$. We sometimes speak of the reversal $\text{rev}(\rho)$ of a run ρ of V , implicitly meaning a run in the reversal of V .

As the norm of $v = (v_1, v_2) \in \mathbb{Q}^2$, we take the largest of absolute values of v_1 and v_2 , $\|v\| := \max\{|v_1|, |v_2|\}$. By the norm of a configuration $c = (q, v)$ we mean the norm of its vector v , and by the norm $\|V\|$ of a 2-VASS V we mean the largest among norms of effects of transitions.



■ **Figure 2** Above $u_1 \circ u_2 \circ \dots \circ u_{11} \circ u_1$. Also, $u_4 \circ u_9$, but $u_4 \not\circ u_{11}$ and $u_{11} \circ u_4$. Pairs of vectors u_i, u_{i+6} are contralinear, for $i = 1, \dots, 5$.

Sequential cones. For a vector $v \in \mathbb{Z}^2$, define the half-line induced by v as $\ell_v := \mathbb{Q}_{\geq 0} \cdot v = \{\alpha v \mid \alpha \in \mathbb{Q}_{\geq 0}\}$. We call two vectors v, w *colinear* if $\ell_v = \ell_w$, and *contralinear* if $\ell_v = \ell_{-w}$. For two vectors $u, v \in \mathbb{Z}^2 \setminus \{(0, 0)\}$, define the angle $\angle[u, v] \subseteq \mathbb{Q}^2$ as the union of all half-lines which lie clock-wise between ℓ_u and ℓ_v , including the two half-lines themselves. In particular, $\angle[v, v] = \ell_v$. Analogously we define the sets $\angle[u, v]$, $\angle(u, v)$ and $\angle(u, v)$ which exclude one or both of the half-lines. We refer to an angle of the form $\angle[v, -v]$ as *half-plane*. We write $v \circ u$ when $u \in \angle(v, -v)$, i.e., u is oriented clock-wise with respect to v (see Figure 2 for an illustration). Note that \circ defines a total order on pairwise non-colinear non-negative vectors.

By the *cone* of a finite set of vectors $\{v_1, \dots, v_k\} \subseteq \mathbb{Z}^2$ we mean the set of all non-negative rational linear combinations of these vectors:

$$\text{CONE}(v_1, \dots, v_k) := \{\sum_{j=1}^k a_j v_j \in \mathbb{Q}^2 \mid a_1, \dots, a_k \in \mathbb{Q}_{\geq 0}\}.$$

We call the cone of a single vector $\text{CONE}(v) = \ell_v$ *trivial*, and the cone of zero vectors $\text{CONE}(\emptyset) = \{(0, 0)\}$ *degenerate*. Two non-zero vectors v_1 and v_2 can be in four distinct relations: (I) they are colinear, (II) they are contralinear, (III) $v_1 \circ v_2$ and hence $\text{CONE}(v_1, v_2) = \angle[v_1, v_2]$, (IV) $v_2 \circ v_1$ and hence $\text{CONE}(v_1, v_2) = \angle[v_2, v_1]$.

► **Lemma 2.1.** *Every cone either equals the whole plane \mathbb{Q}^2 , or is included in some half-plane.*

Proof. Assume, w.l.o.g. that the vectors v_1, \dots, v_k are non-zero and include no colinear pair. Suppose there is a contralinear pair v_i, v_j among v_1, \dots, v_k . If all other vectors v_h satisfy $v_i \circ v_h \circ v_j$ then $\text{CONE}(v_1, \dots, v_k)$ is included in the half-plane $\angle[v_i, v_j]$. Otherwise $\text{CONE}(v_1, \dots, v_k)$ is the whole plane.

Now suppose there is no contralinear pair among v_1, \dots, v_k . If some three v_i, v_j, v_h of them satisfy $v_i \circ v_j \circ v_h \circ v_i$ then $\text{CONE}(v_1, \dots, v_k)$ includes the three angles $\angle[v_i, v_j]$, $\angle[v_j, v_h]$ and $\angle[v_h, v_i]$, the union of which is the whole plane. Otherwise, the relation \circ is

transitive and hence defines a (strict) total order on $\{v_1, \dots, v_k\}$. The minimal and maximal element v_i and v_j w.r.t. the order satisfy $v_i \circ v_j$, and hence $\text{CONE}(v_1, \dots, v_k) = \angle[v_i, v_j]$ is included in the half-plane $\angle[v_i, -v_i]$. ◀


The *sequential cone* of vectors $v_1, \dots, v_k \in \mathbb{Z}^2$ imposes additional non-negativeness conditions, namely for every i , the partial sum $a_1v_1 + \dots + a_iv_i$ must be non-negative (this is required later, when pumping cycles in a run whose effects are v_1, \dots, v_k in that order):

$$\text{SEQCONE}(v_1, \dots, v_k) := \{\sum_{j=1}^k a_j v_j \in \mathbb{Q}_{\geq 0}^2 \mid a_1, \dots, a_k \in \mathbb{Q}_{\geq 0}, \forall_i \sum_{j=1}^i a_j v_j \in \mathbb{Q}_{\geq 0}^2\}.$$

Note that v_1 may be assumed w.l.o.g. to be semi-positive, but other vectors v_i are not necessarily non-negative; and that every sequential cone is a subset of the non-negative orthant $\mathbb{Q}_{\geq 0}^2$. Importantly, contrarily to cones, the order of vectors v_1, \dots, v_k matters for sequential cones. In fact, sequential cones are just convenient syntactic sugar for cones of pairs of non-negative vectors:

► **Lemma 2.2.** *For all vectors v_1, \dots, v_k , the sequential cone $\text{SEQCONE}(v_1, \dots, v_k)$, if not degenerate, equals $\text{CONE}(u, v)$, for two non-negative vectors u, v , and each of them either belongs to $\{v_1, \dots, v_k\}$, or is horizontal, or vertical.*

Proof. We proceed by induction on k . For $k = 1$ we have $\text{SEQCONE}(v_1) = \ell_{v_1} = \text{CONE}(v_1, v_1)$. Let v_0 and h_0 denote some fixed vertical and horizontal vector, respectively. For the induction step we assume $\text{SEQCONE}(v_1, \dots, v_{k-1}) = \text{CONE}(u, v)$ for non-negative vectors u, v ; and compute the value of $\text{SEQCONE}(v_1, \dots, v_k)$, separately in each of the following distinct cases (assume w.l.o.g. $u \circ v$):




$$\text{SEQCONE}(v_1, \dots, v_k) = \begin{cases} \text{CONE}(v_k, v) & \text{if } v_k \in \angle[v_0, u] \\ \text{CONE}(u, v) & \text{if } v_k \in \angle[u, v] \\ \text{CONE}(u, v_k) & \text{if } v_k \in \angle(v, h_0] \\ \text{CONE}(u, h_0) & \text{if } v_k \in \angle(h_0, -u] \\ \text{CONE}(v_0, h_0) & \text{if } v_k \in \angle(-u, -v) \\ \text{CONE}(v_0, v) & \text{if } v_k \in \angle[-v, v_0]. \end{cases}$$

◀

3 Thin-Thick Dichotomy

The main result of this section (cf. Theorem 3.1 below) classifies $(0, 0)$ -runs in a 2-VASS into *thin* and *thick* ones. Throughout this section we consider an arbitrary fixed 2-VASS $V = (Q, T)$.

Let $n = |Q|$ and $M = \|V\|$.



■ **Figure 3** Thin run within belts $B_{v,W}$.

Thin runs. The belt of direction $v \in \mathbb{N}^2$ and width W is the set

$$\mathcal{B}_{v,W} = \{u \in \mathbb{N}^2 \mid \text{dist}(u, \ell_v) \leq W\},$$

where $\text{dist}(u, \ell_v)$ denotes the Euclidean distance between the point u and the half-line ℓ_v . For $A \in \mathbb{N}$, we call $\mathcal{B}_{v,W}$ an A -belt if $\|v\| \leq A$ and $W \leq A$. We say that a run ρ of V is A -thin if for every configuration c in ρ there exists an A -belt B such that $c \in Q \times B$.

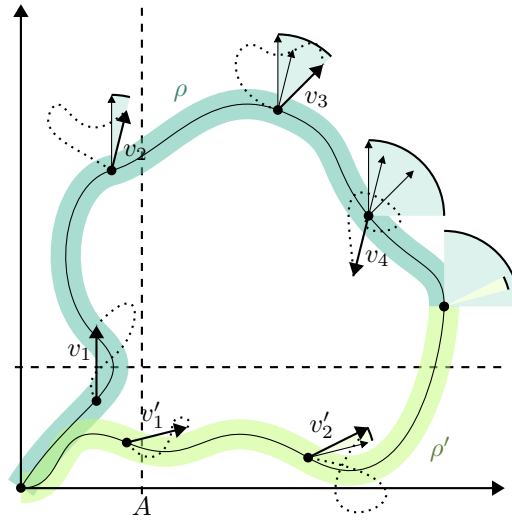
Thick runs. Let $A \in \mathbb{N}$. Four cycles $\pi_1, \pi_2, \pi_3, \pi_4 \in T^*$ are A -sequentially enabled in a run ρ if their lengths are at most A , and the run ρ factors into $\rho = \rho_1 \rho_2 \rho_3 \rho_4 \rho_5$ so that (denote by v_1, v_2, v_3, v_4 the effects of $\pi_1, \pi_2, \pi_3, \pi_4$, respectively):

- The effect v_1 is semi-positive, the cycle π_1 is enabled in $c_1 := \text{trg}(\rho_1)$, and both coordinates are bounded by A along ρ_1 .
- If v_1 is positive then π_2 is \emptyset -enabled in $c_2 := \text{trg}(\rho_2)$. Otherwise (let j be the coordinate s.t. $v_1[j] = 0$) π_2 is $\{j\}$ -enabled in $c_2 := \text{trg}(\rho_2)$, and j th coordinate is bounded by A along ρ_2 .
- The cycle π_i is \emptyset -enabled in $c_i := \text{trg}(\rho_i)$, for $i = 3, 4$.

We also say that the four vectors v_1, v_2, v_3, v_4 are A -sequentially enabled in ρ , quantifying the cycles existentially. A $(0,0)$ -run τ is called A -thick if it partitions into $\tau = \rho \rho'$ so that

1. some vectors v_1, v_2, v_3, v_4 are A -sequentially enabled in ρ ,
2. some vectors v'_1, v'_2, v'_3, v'_4 are A -sequentially enabled in $\text{rev}(\rho')$,
3. $\text{SEQCONE}(v_1, v_2, v_3, v_4) \cap \text{SEQCONE}(v'_1, v'_2, v'_3, v'_4)$ is non-trivial.

Figure 4 illustrates the geometric ideas underlying these three conditions for A -thick runs. Concerning condition 1, a cycle π_1 depicted by a dotted line, with vertical effect v_1 , can be used to increase the second (vertical) coordinate arbitrarily, which justifies the relaxed requirement that a cycle π_2 with effect v_2 is only $\{1\}$ -enabled. Note that the norm of the configuration enabling π_1 , as well as the first coordinate of the configuration enabling π_2 , are bounded by A . Concerning condition 2, a cycle π'_1 with positive effect v'_1 can be used to increase both coordinates arbitrarily; therefore a cycle π'_2 with effect v'_2 is only required to be \emptyset -enabled, and no coordinate of the configuration enabling π'_2 is required to be bounded by A . In the illustrated example, vectors v'_3 and v'_4 are not needed; formally, one can assume $v'_2 = v'_3 = v'_4$ and $\rho'_3 = \rho'_4 = \varepsilon$. Condition 3 ensures that the cycles



■ **Figure 4** Thick run. Blue angles denote sequential cones $\text{SEQCONE}(v_1, v_2)$, $\text{SEQCONE}(v_1, v_2, v_3)$ and $\text{SEQCONE}(v_1, v_2, v_3, v_4)$, respectively, and green angle denotes $\text{SEQCONE}(v'_1, v'_2)$.

π_1, \dots, π_4 and π'_1, \dots, π'_4 can be pumped such that the pumped versions of ρ and ρ' are still connected. In the illustrated example, observe that $\text{SEQCONE}(v_1, v_2) \cap \text{SEQCONE}(v'_1) = \emptyset$. Intuitively, both coordinates in the target of ρ can be increased arbitrarily using v_1 and v_2 , and similarly both coordinates of the target of $\text{rev}(\rho')$ can be increased arbitrarily using v'_1 , but “directions of increase” are non-crossing. Adding v_3 and v'_2 is not sufficient, as still $\text{SEQCONE}(v_1, v_2, v_3) \cap \text{SEQCONE}(v'_1, v'_2) = \emptyset$. When vector v_4 is adjoined, condition 3 holds as $\text{SEQCONE}(v_1, v_2, v_3, v_4) = \mathbb{Q}_{\geq 0}^2$. Finally, the four vectors are really needed here, e.g., vector v_3 can not be omitted as $\text{SEQCONE}(v_1, v_2, v_4) = \text{SEQCONE}(v_1, v_2)$.

Here is the main result of this section:

► **Theorem 3.1** (Thin-Thick Dichotomy). *There is a polynomial p such that every $(0, 0)$ -run in a 2-VASS V is either $p(nM)^n$ -thin or $p(nM)^n$ -thick.*

For the proof of Theorem 3.1 we need the following core fact (for space reasons, the proof is only in the full version of the paper):

► **Lemma 3.2** (Non-negative Cycle Lemma). *There is a polynomial P such that every run ρ in V from a $(0, 0)$ -configuration to a target configuration of norm larger than $P(nM)^n$, contains a configuration enabling a semi-positive cycle of length at most $P(nM)$.*

Proof of Theorem 3.1. Let P be the polynomial from Lemma 3.2. The polynomial p required in Theorem 3.1 can be chosen arbitrarily as long as $p(x) \geq \sqrt{2} \cdot (P(x) + (x + 1)^3) \cdot x$ for all x ; note that the following inequality follows:

$$p(nM)^n \geq \sqrt{2} \cdot ((P(nM))^n + (nM + 1)^3) \cdot nM. \tag{1}$$

In the sequel we deliberately confuse configurations $c = (q, v)$ with *their* vectors v : whenever convenient, we use c to denote the vector v , hoping that this does not lead to any confusion.

Let τ be a $(0, 0)$ -run of V which is not $p(nM)^n$ -thin, i.e., τ contains therefore a configuration t which lies outside of all the $p(nM)^n$ -belts. We need to demonstrate points 1–3 in the

definition of thick run. To this aim we split τ into $\tau = \rho\rho'$ where $\text{trg}(\rho) = t = \text{src}(\rho')$, and are going to prove the following two claims (a) and (a'). Let $D := P(nM)^n + (nM + 1)^3$. For $x, y \in \mathbb{Q}^2$, let $\text{dist}(x, y)$ denote their Euclidean distance.

- (a) Some vectors v_1, v_2, v_3, v_4 are $P(nM)^n$ -sequentially enabled in ρ , and the sequential cone $\text{SEQCONE}(v_1, v_2, v_3, v_4)$ contains a point $u \in \mathbb{Q}_{\geq 0}^2$ with $\|u - t\| \leq D$.
- (a') Some vectors v'_1, v'_2, v'_3, v'_4 are $P(nM)^n$ -sequentially enabled in $\text{rev}(\rho')$, and the sequential cone $\text{SEQCONE}(v'_1, v'_2, v'_3, v'_4)$ contains a point $u \in \mathbb{Q}_{\geq 0}^2$ with $\|u - t\| \leq D$.

In simple words, instead of proving point 3, we prove that both sequential cones contain a point v which is sufficiently close to t .

▷ **Claim 3.3.** The conditions (a) and (a') guarantee that τ is thick.

Indeed, points 1–2 in the definition of thick run are immediate as $P(nM) \leq p(nM)$. For point 3, observe that the inequality (1) implies $p(nM)^n \geq \sqrt{2} \cdot D$, which guarantees that the circle $\{u \in \mathbb{Q}_{\geq 0}^2 \mid \text{dist}(u, t) \leq \sqrt{2} \cdot D\}$ does not touch any half-line ℓ_w induced by a non-negative vector w with $\|w\| \leq p(nM)^n$. In consequence, neither does the square $X := \{u \in \mathbb{Q}_{\geq 0}^2 \mid \|u - t\| \leq D\}$ inscribed in the circle, and hence X lies between two consecutive half-lines ℓ_w induced by a non-negative vector w with $\|w\| \leq p(nM)^n$. Hence, as $\text{SEQCONE}(v_1, v_2, v_3, v_4)$ contains some point of X , by Lemma 2.2 it includes the whole X , and likewise $\text{SEQCONE}(v'_1, v'_2, v'_3, v'_4)$. In consequence, the whole X is included in $\text{SEQCONE}(v_1, v_2, v_3, v_4) \cap \text{SEQCONE}(v'_1, v'_2, v'_3, v'_4)$ which entails point 3. Claim 3.3 is thus proved.

As condition (a') is fully symmetric to (a), we focus exclusively on proving condition (a), i.e., on constructing sequentially enabled vectors v_1, v_2, v_3, v_4 .

Vector t lies outside of $p(nM)^n$ -belts, hence outside of all the $P(nM)^n$ -belts, therefore its norm $\|t\| > P(nM)^n$. Relying on Lemma 3.2, let c_1 be the first configuration in the run ρ which enables a semi-positive cycle π_1 of length bounded by $P(nM)$, and let $v_1 = \text{eff}(\pi_1)$. We start with the following obvious claim (let v_0 be some vertical vector, e.g. $v_0 = (0, 1)$):

▷ **Claim 3.4.** $\text{SEQCONE}(v_0)$ contains a point $u \in \mathbb{Q}_{\geq 0}^2$ such that $\|u - c_1\| \leq P(nM)^n + nM$.

Indeed, due to Lemma 3.2 we may assume $\|c_1\| \leq P(nM)^n + M$ and hence $u = (0, 0)$ does the job.

Recall that the relation \circlearrowleft defines a total order on pairwise non-colinear non-negative vectors.


▷ **Claim 3.5.** We can assume w.l.o.g. that $v_1 \circlearrowleft t$.

Indeed, if v_1 and t were colinear then $t \in \text{CONE}(v_1)$ and hence condition (a) would hold.

Split ρ into the prefix ending in c_1 and the remaining suffix: $\rho = \rho_1\sigma$, where $\text{trg}(\rho_1) = c_1 = \text{src}(\sigma)$. As the next step we will identify a configuration c_2 in σ which satisfies Claim 3.6 (which will serve later as the basis of induction) and enables a cycle π_2 with effect v_2 (as stated in Claim 3.7).

▷ **Claim 3.6.** $\text{SEQCONE}(v_0, v_1)$ contains a point $u \in \mathbb{Q}_{\geq 0}^2$ such that $\|u - c_2\| \leq P(nM)^n + 2nM$.

The proof of Claim 3.6 depends on whether v_1 is positive. If v_1 is so, we simply duplicate the first cycle: $c_2 := c_1$ and $\pi_2 := \pi_1$, and use Claim 3.4. Otherwise v_1 is vertical due to Claim 3.5. If $t[1] \leq W = P(nM)^n + (n + 1)M$ then condition (a) holds immediately as $\text{SEQCONE}(v_1) = \ell_{v_1}$ contains a point $u \in \mathbb{Q}_{\geq 0}^2$ with $\|u - t\| \leq P(nM)^n + (n + 1)M \leq D$. Therefore suppose $t[1] > P(nM)^n + (n + 1)M$, and define the sequence d_1, \dots, d_m of configurations as follows (cf. Figure 5): let $d_1 := c_1$, and let d_{i+1} be the first configuration in



■ **Figure 5** Proof of Claim 3.6.

σ with $d_{i+1}[1] > d_i[1]$. Recall that $d_1[1] \leq P(nM)^n + M$, and observe that $d_{i+1}[1] \leq d_i[1] + M$. Thus by the pigeonhole principle $m > n$ and hence for some $i < j \leq n + 1$ the configurations d_i and d_j must have the same control state. The infix σ_{ij} of the path σ from d_i to d_j is thus a cycle, enabled in d_i , whose effect is positive on the first (horizontal) coordinate. Let $c_2 := d_i$. As $c_2[1] \leq P(nM)^n + (n + 1)M$, $\text{SEQCONE}(v_0, v_1) = \ell_{v_0}$ contains necessarily a point $u \in \mathbb{Q}_{\geq 0}^2$ such that $\|u - c_2\| \leq P(nM)^n + (n + 1)M$, which proves Claim 3.6.

▷ **Claim 3.7.** The configuration $c_2 \{1\}$ -enables a cycle π_2 of length bounded by $p(nM)^n$, such that the first coordinate of $\text{eff}(\pi_2)$ is positive.

Recalling the proof of the previous claim, observe that the first (horizontal) coordinate in the infix σ_{ij} is bounded by $P(nM)^n + (n + 1)M$, and think of the second (vertical) coordinate as irrelevant. Let π_2 be the path inducing σ_{ij} . For bounding the length of π_2 , as long as π_2 contains a cycle α with vertical effect $(0, w)$, remove α from π_2 . This process ends yielding a cycle π_2 of length at most $(P(nM)^n + (n + 1)M) \cdot n$, and hence at most $p(nM)^n$ (by the inequality (1)), which is $\{1\}$ -enabled in c_2 , but not necessarily enabled. Let $v_2 := \text{eff}(\pi_2)$.


▷ **Claim 3.8.** We can assume w.l.o.g. that $v_2 \circlearrowleft t$.

Indeed, if $v_1 = v_2$ then Claim 3.5 does the job; otherwise v_1 is vertical and then $t \circlearrowleft v_2$ (or t colinear with v_2) would imply $t \in \text{SEQCONE}(v_1, v_2)$, hence condition (a) would hold again.

Split σ further into the prefix ending in c_2 and the remaining suffix: $\sigma = \rho_2 \sigma'$, where $\text{trg}(\rho_2) = c_2 = \text{src}(\sigma')$. If σ' contains a configuration which \emptyset -enables a simple cycle whose effect w belongs to $\angle[t, -v_2)$ then $t \in \text{SEQCONE}(v_2, w)$ and hence condition (a) holds. We aim at achieving this objective incrementally (cf. Figure 6).

For $i \geq 2$, let c_{i+1} be the first configuration in σ' after c_i that \emptyset -enables a simple cycle π_{i+1} with effect $v_{i+1} \in \angle(v_i, -v_i)$. As discussed above, if $v_{i+1} \in \angle[t, -v_i)$ for some i then $t \in \text{SEQCONE}(v_i, v_{i+1})$ and hence condition (a) holds. Assume therefore that the sequence v_1, \dots, v_m so defined satisfies $v_{i+1} \in \angle(v_i, t)$ for all $i \geq 2$. Let $c_{m+1} := t$. As vectors v_3, \dots, v_m are pairwise different, semi-positive and, being effects of simple cycles, have norms at most nM , we know that $m \leq (nM + 1)^2 + 1$.

▷ **Claim 3.9.** For every $i = 1, \dots, m$, $\text{SEQCONE}(v_0, v_i)$ contains a point $u \in \mathbb{Q}_{\geq 0}^2$ such that $\|u - c_{i+1}\| \leq P(nM)^n + (i + 1)nM$.



■ **Figure 6** Incremental construction of v_1, \dots, v_m .

Proof. By induction on i . The induction base is exactly Claim 3.6. For the induction step, we are going to show that $\text{SEQCONE}(v_0, v_i)$ contains a vector u such that $\|u - c_{i+1}\| \leq P(nM)^n + (i+1)nM$. Decompose the infix of σ' which starts in c_i and ends in c_{i+1} into simple cycles, plus the remaining path $\bar{\rho}$ of length at most n . The norm of the effect \bar{v} of $\bar{\rho}$ is hence bounded by nM , and we have

$$c_{i+1} = c_i + s + \bar{v},$$

where s is the sum of effects of all the simple cycles. By the definition of v_{i+1} , the effects of all the simple cycles belong to the half-plane $\angle[-v_i, v_i]$, and hence there belongs s . By induction assumption there is $u' \in \text{SEQCONE}(v_0, v_{i-1})$ such that $\|u' - c_i\| \leq P(nM)^n + inM$. As $v_{i-1} \circlearrowleft v_i$, we also have $u' \in \text{SEQCONE}(v_0, v_i)$. Consider the point

$$u := u' + s$$

which necessarily belongs to the half-plane $\angle[-v_i, v_i]$ but not necessarily to $\text{SEQCONE}(v_0, v_i) = \angle[-v_i, v_i] \cap \mathbb{Q}_{\geq 0}^2$. Ignoring this issue, by routine calculations we get

$$\|u - c_{i+1}\| = \|u' + s - c_i - s - \bar{v}\| \leq \|u' - c_i\| + \|\bar{v}\| \leq \|u' - c_i\| + nM \leq P(nM)^n + (i+1)nM$$

as required for the induction step. Finally, if $u \notin \mathbb{Q}_{\geq 0}^2$, translate u towards c_{i+1} until it enters the non-negative orthant $\mathbb{Q}_{\geq 0}^2$; clearly, the translation can only decrease the value of $\|u - c_{i+1}\|$. ◀

Applying the claim to $i = m$, and knowing that $m \leq (nM + 1)^2 + 1$, we get some point $u \in \text{SEQCONE}(v_0, v_m)$ such that $\|u - t\| \leq P(nM)^n + ((nM + 1)^2 + 1) \cdot nM \leq P(nM)^n + (nM + 1)^3$. Furthermore, relying on the assumptions that t lies outside of all $p(nM)^n$ -belts and that $v_1 \circlearrowleft t$ we prove, similarly as in the proof of Claim 3.3, that $v_1 \circlearrowleft u$ and hence the point u belongs also to $\text{SEQCONE}(v_1, v_m)$. This completes the proof of Theorem 3.1. ◀

4 Dichotomy in Action

This section illustrates applicability of Theorem 3.1. As before, we use symbols n and M for the number of control states, and the norm of a 2-VASS, respectively. As the first corollary we provide a pumping lemma for 2-VASS: in case of thin runs apply, essentially, pumping schemes of 1-VASS, and in case of thick runs use the cycles enabled along a run. As another application, we derive an alternative proof of the exponential run property for 2-VASS.

► **Theorem 4.1** (Pumping). *There is a polynomial p such that every $(0,0)$ -run τ in a 2-VASS of length greater than $p(nM)^n$ factors into $\tau = \tau_0 \tau_1 \dots \tau_k$ ($k \geq 1$), so that for some non-empty cycles $\alpha_1, \dots, \alpha_k$ of length at most $p(nM)^n$, the path $\tau_0 \alpha_1^i \tau_1 \alpha_2^i \dots \alpha_k^i \tau_k$ is a $(0,0)$ -run for every $i \in \mathbb{N}$. Furthermore, the lengths of τ_0 and τ_k are also bounded by $p(nM)^n$.*

► **Theorem 4.2** (Exponential run). *There is a polynomial p such that for every $(0,0)$ -run τ in a 2-VASS, there is a $(0,0)$ -run of length bounded by $p(nM)^n$ with the same source and target as τ .*

We fix from now on a 2-VASS $V = (Q, T)$ and the polynomial p of Theorem 3.1. Let $A = p(nM)^n$. Both proofs proceed separately for thin and thick runs τ . The former (fairly standard) case is treated in the full version of the paper, so assume below τ to be A -thick. The polynomials required in Theorems 4.1 and 4.2 can be read out from the constructions.

We rely on the standard tool, cf. Prop. 2 in [4] (the norm of a system of inequalities is the largest absolute value of its coefficient, and likewise we define the norm of a solution):

► **Lemma 4.3.** *Let \mathcal{U} be a system of d linear inequalities of norm M with k variables. Then the smallest norm of a non-negative-integer solution of \mathcal{U} is in $\mathcal{O}(k \cdot M)^d$.*

Consider a split $\tau = \rho\rho'$, where $\rho = \rho_1 \rho_2 \rho_3 \rho_4 \rho_5$ and $\rho' = \rho'_5 \rho'_4 \rho'_3 \rho'_2 \rho'_1$, as well as cycles π_1, \dots, π_4 and π'_1, \dots, π'_4 given by the definition of thick run. Let v_1, \dots, v_4 and v'_1, \dots, v'_4 be the respective effects of π_1, \dots, π_4 and π'_1, \dots, π'_4 . For $j = 1, \dots, 4$ let $c_j = \text{trg}(\rho_j)$ and for $j = 2, \dots, 4$ let $e_j \in \mathbb{N}^2$ be the minimal non-negative vector such that the configuration $c_j + e_j$ enables cycle π_j . We define the following system \mathcal{U} of linear inequalities with 6 variables a_1, a_2, a_3, a_4, x, y (max is understood point-wise):

$$a_1 v_1 \geq e_2 \tag{2}$$

$$a_1 v_1 + a_2 v_2 \geq \max(e_2, e_3) \tag{3}$$

$$a_1 v_1 + a_2 v_2 + a_3 v_3 \geq \max(e_3, e_4) \tag{4}$$

$$a_1 v_1 + a_2 v_2 + a_3 v_3 + a_4 v_4 = (x, y) \geq e_4 \tag{5}$$

(Observe that when $v_1[j] = 0$, i.e., in case when v_1 is vertical or horizontal, $e_j = 0$ and therefore one of the two first inequalities is always satisfied, namely $a_1 v_1[j] \geq e_2[j]$.) Likewise, we have a system of inequalities \mathcal{U}' with 6 variables $a'_1, a'_2, a'_3, a'_4, x', y'$. Observe that the sequential cone $\text{SEQCONE}(v_1, v_2, v_3, v_4)$ contains exactly (projections on (x, y) of) non-negative rational solutions of the modified system $\mathcal{U}^{(0,0)}$ obtained by replacing all the right-hand sides with $(0,0)$. Likewise we define $\mathcal{U}'^{(0,0)}$. Finally, we define the compound system \mathcal{C} by enhancing the union of \mathcal{U} and \mathcal{U}' with two additional equalities (likewise we define the system $\mathcal{C}^{(0,0)}$)

$$(x, y) = (x', y'). \tag{6}$$

► **Claim 4.4.** \mathcal{C} admits a non-negative integer solution $(a_1, a_2, a_3, a_4, x, y, a'_1, a'_2, a'_3, a'_4, x', y')$.

Proof. The system $\mathcal{C}^{(0,0)}$ admits a non-negative rational solution as the intersection of the cones $\text{SEQCONE}(v_1, v_2, v_3, v_4)$ and $\text{SEQCONE}(v'_1, v'_2, v'_3, v'_4)$ is non-empty by assumption. As intersection of cones is stable under multiplications by non-negative rationals, the solution can be scaled up arbitrarily, to yield a non-negative integer one, and even a non-negative integer solution of the stronger system \mathcal{C} . ◀

► **Claim 4.5.** For every non-negative integer solution of \mathcal{C} , for the cycles defined as $\alpha_j := \pi_j^{a_j}$ and $\alpha'_j := (\pi'_j)^{a'_j}$, for $j = 1, 2, 3, 4$, the following path is a $(0,0)$ -run:

$$\rho_1 \alpha_1 \rho_2 \alpha_2 \rho_3 \alpha_3 \rho_4 \alpha_4 \rho_5 \rho'_5 \alpha'_4 \rho'_4 \alpha'_3 \rho'_3 \alpha'_2 \rho'_2 \alpha'_1 \rho'_1.$$

Proof. The first two inequalities (2) enforce that the first cycle π_1 is repeated sufficiently many a_1 times so that π_2 is enabled in configuration $\text{trg}(\rho_1 \alpha_1 \rho_2)$. Then the next two inequalities (3) enforce that π_1 and π_2 are jointly repeated sufficiently many a_1, a_2 times so that π_2 is still enabled after its last repetition (which guarantees that every of intermediate repetitions of π_2 is also enabled), and that π_3 is enabled in configuration $\text{trg}(\rho_1 \alpha_1 \rho_2 \alpha_2 \rho_3)$. Likewise for (4). Finally, the inequalities (5) enforce that π_1, \dots, π_4 are jointly repeated sufficiently many times so that π_4 is still enabled after its last repetition. Analogous argument, but in the reverse order, applies for the repetitions of π'_4, \dots, π'_1 . Finally, equalities (6) ensure that the total effect of $\alpha_1, \dots, \alpha_4$ is precisely compensated by the total effect of $\text{rev}(\alpha'_1), \dots, \text{rev}(\alpha'_4)$. ◀

Proof of Theorem 4.1. Consider a solution of \mathcal{C} . In particular the sum $\text{eff}(\alpha_1) + \dots + \text{eff}(\alpha_j)$, as well as $\text{eff}(\text{rev}(\alpha'_1)) + \dots + \text{eff}(\text{rev}(\alpha'_j))$, is necessarily non-negative for every $j = 1, \dots, 4$. Therefore, as a direct corollary of Claim 4.5, for every $i \in \mathbb{N}$ the path

$$\rho_1 \alpha_1^i \rho_2 \alpha_2^i \rho_3 \alpha_3^i \rho_4 \alpha_4^i \rho_5 \rho'_5 (\alpha'_4)^i \rho'_4 (\alpha'_3)^i \rho'_3 (\alpha'_2)^i \rho'_2 (\alpha'_1)^i \rho'_1$$

is also a $(0, 0)$ -run. For bounding the lengths of cycles we use Claim 4.4 and apply Lemma 4.3 to \mathcal{C} , to deduce that \mathcal{C} admits a non-negative integer solution of norm polynomial in $A = p(nM)^n$. This, together with the bounds on lengths of cycles π_1, \dots, π_4 and π'_1, \dots, π'_4 in the definition of A -thick run, entails required bounds on the lengths of the pumpable cycles. Finally, the lengths of the extremal factors ρ_1 and ρ'_1 can be also bounded: if ρ_1 (resp. ρ'_1) is long enough it must admit a repetition of configuration, we add one more cycle determined by the first (resp. last) such repetition, thus increasing k from 8 to 10. ◀

For proving Theorem 4.2 we will need a slightly more elaborate pumping. By the definition of thick run, both coordinates are bounded by A along ρ_1 and ρ'_1 . W.l.o.g. assume that no configuration repeats in each of the two runs, and hence their lengths are bounded by A^2 .

Let \mathcal{C}_δ denote the union of \mathcal{U} and \mathcal{U}' enhanced, this time, by the two equalities

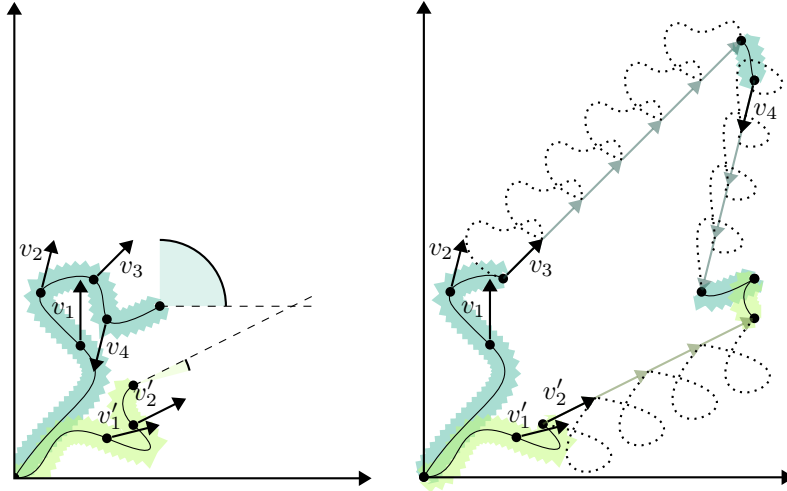
$$(x, y) + (\delta_x, \delta_y) = (x', y').$$

The two additional variables δ_x, δ_y describe, intuitively, possible differences between the total effect of $\pi_1^{a_1}, \dots, \pi_4^{a_4}$ and the total effect of $\text{rev}(\pi'_1)^{a'_1}, \dots, \text{rev}(\pi'_4)^{a'_4}$. The projection of any solution of \mathcal{C}_δ on variables (δ_x, δ_y) we call below a *shift*.

▷ **Claim 4.6.** For some non-negative integer m bounded polynomially with respect to A , all the four vectors $(0, m)$, $(m, 0)$, $(0, -m)$ and $(-m, 0)$ are shifts.

Proof. We reason analogously as in the proof of Claim 4.4, but this time we rely on the assumption that intersection of the cones $\text{SEQCONE}(v_1, v_2, v_3, v_4)$ and $\text{SEQCONE}(v'_1, v'_2, v'_3, v'_4)$ is non-trivial, and hence contains, for some $v \in \mathbb{Q}_{>0}^2$ and $a \in \mathbb{Q}_{>0}$, the points v and $v + (0, a)$. By scaling we obtain an integer point $v' \in \mathbb{N}^2$ and a non-negative integer $m_1 \in \mathbb{N}$ so that v' and $v' + (0, m_1)$ both belong to the intersection of cones. Therefore the vector $(0, m_1)$ is a shift. Likewise we obtain three other non-negative integers $m_2, m_3, m_4 \in \mathbb{N}$ such that $(m_2, 0)$, $(0, -m_3)$ and $(-m_4, 0)$ are all shifts. Each of the integers m_1, \dots, m_4 can be bounded polynomially in A using Lemma 4.3. As shifts are stable under multiplication by non-negative integers, it is enough to take as m the least common multiple of the four integers. ◀

Proof of Theorem 4.2. We use m from the last claim to modify all factors of τ except for ρ_1 and ρ'_1 , in order to reduce their lengths to at most $n \cdot m^2$. W.l.o.g. assume m to be larger than A (take a sufficient multiplicity of m otherwise); this assumption allows us to



■ **Figure 7** Contracted paths $\tilde{\rho}, \tilde{\rho}'$ (left) and reconstructed $(0,0)$ -run $\bar{\tau} = \bar{\rho} \bar{\rho}'$ (right).

proceed uniformly, irrespectively whether v_1 is positive or not. Observe that any path longer than $n \cdot m^2$ must contain two configurations with the same control state whose vectors are coordinate-wise congruent modulo m . As long as this happens, we remove the infix; note that this operation changes the effect of the whole path by a multiplicity of m on every coordinate. If this operation is performed on factors $\rho_2, \rho_3, \rho_4, \rho_5, \rho'_5, \rho'_4, \rho'_3, \rho'_2$, the paths ρ, ρ' are transformed into contracted paths (see the left picture in Figure 7) of the form:

$$\tilde{\rho} = \rho_1 \tilde{\rho}_2 \tilde{\rho}_3 \tilde{\rho}_4 \tilde{\rho}_5, \quad \tilde{\rho}' = \tilde{\rho}'_5 \tilde{\rho}'_4 \tilde{\rho}'_3 \tilde{\rho}'_2 \sigma_1,$$

each of total length at most $5n \cdot m^2$. Importantly, their effects $\text{eff}(\tilde{\rho})$ and $\text{eff}(\tilde{\rho}')$ are bounded polynomially in A , and their difference is (coordinate-wise) divisible by m :

$$\text{eff}(\tilde{\rho}) - \text{eff}(\text{rev}(\tilde{\rho}')) = (am, bm) \quad \text{for some integers } a, b \in \mathbb{Z} \text{ polynomial in } A.$$

Our aim is now to pump up the cycles π_1, \dots, π_4 and $\text{rev}(\pi'_1), \dots, \text{rev}(\pi'_4)$ (see the right picture in Figure 7), to finally end up with the paths of the form

$$\bar{\rho} = \rho_1 \pi_1^{a_1} \tilde{\rho}_2 \pi_2^{a_2} \tilde{\rho}_3 \pi_3^{a_3} \tilde{\rho}_4 \pi_4^{a_4} \tilde{\rho}_5, \quad \bar{\rho}' = \tilde{\rho}'_5 (\pi'_4)^{a'_4} \tilde{\rho}'_4 (\pi'_3)^{a'_3} \tilde{\rho}'_3 (\pi'_2)^{a'_2} \tilde{\rho}'_2 (\pi'_1)^{a'_1} \rho'_1, \quad (7)$$

such that $\bar{\tau} = \bar{\rho} \bar{\rho}'$ is a $(0,0)$ -run. In other words, we aim at $\text{eff}(\bar{\rho}) = \text{eff}(\text{rev}(\bar{\rho}'))$. We are going to use Lemma 4.3 twice. For $j = 2, \dots, 5$ let $c_j := \text{eff}(\rho_1 \tilde{\rho}_2 \dots \tilde{\rho}_j) \in \mathbb{Z}^2$, and let f_j be the minimal non-negative vector such that the configuration $c_{j-1} + f_j$ enables $\tilde{\rho}_j$. For $j = 2, \dots, 4$ let $e_j \in \mathbb{N}^2$ be the minimal non-negative vector such that the configuration $c_j + e_j$ enables π_j . Finally, let e_5 be the minimal non-negative vector such that $c_5 + e_5 \geq (0,0)$. Analogously to the system \mathcal{U} (2)–(5), we define the system $\tilde{\mathcal{U}}$ of linear inequalities:

$$\begin{aligned} a_1 m v_1 &\geq \max(e_2, f_2) \\ a_1 m v_1 + a_2 m v_2 &\geq \max(e_2, e_3, f_3) \\ a_1 m v_1 + a_2 m v_2 + a_3 m v_3 &\geq \max(e_3, e_4, f_4) \\ a_1 m v_1 + a_2 m v_2 + a_3 m v_3 + a_4 m v_4 &\geq \max(e_4, e_5, f_5) \end{aligned}$$

In words, $\tilde{\mathcal{U}}$ requires that every prefix of $\bar{\rho}$ is enabled in the source $(0,0)$ -configuration, and that the number of repetitions of every cycle π_i is divisible by m . Clearly $\tilde{\mathcal{U}}$ has a non-negative integer solution, as v_1 is either positive, or vertical or horizontal in which case

v_2 is positive on the relevant coordinate. Likewise we define a system of inequalities \tilde{U}' that requires that every prefix of $\text{rev}(\bar{\rho}')$ is enabled in the target $(0, 0)$ -configuration. Consider some fixed solutions of \tilde{U} and \tilde{U}' bounded, by the virtue of Lemma 4.3, polynomially in A . We have thus two fixed runs $\bar{\rho}$ and $\text{rev}(\bar{\rho}')$ of the form (7), with source vector $(0, 0)$; the number of repetitions of each cycles is divisible by m , and the difference of their effects is (coordinate-wise) divisible by m :

$$\text{eff}(\bar{\rho}) - \text{eff}(\text{rev}(\bar{\rho}')) = (am, bm) \quad \text{for some integers } a, b \in \mathbb{Z} \text{ polynomial in } A.$$

As shifts are closed under addition, by Claim 4.6 we know that (am, bm) is a shift. Substituting (am, bm) for (δ_x, δ_y) in the system \mathcal{C}_δ yields a system which admits, again by Lemma 4.3, a solution bounded polynomially in A . We use such a solution to increase the numbers of repetitions of respective cycles a_1, \dots, a_4 and a'_4, \dots, a'_1 in $\bar{\rho}$ and $\bar{\rho}'$, respectively. This turns the path $\bar{\tau} = \bar{\rho}\bar{\rho}'$ into a $(0, 0)$ -run of length bounded polynomially in A . ◀

References

- 1 Mohamed Faouzi Atig, Dmitry Chistikov, Piotr Hofman, K. Narayan Kumar, Prakash Saivasan, and Georg Zetsche. The complexity of regular abstractions of one-counter languages. In *Proc. LICS'16*, pages 207–216, 2016.
- 2 Michael Blondin, Alain Finkel, Stefan Göller, Christoph Haase, and Pierre McKenzie. Reachability in Two-Dimensional Vector Addition Systems with States Is PSPACE-Complete. In *Proc. LICS'15*, pages 32–43, 2015.
- 3 Dmitry Chistikov, Wojciech Czerwinski, Piotr Hofman, Michal Pilipczuk, and Michael Wehar. Shortest Paths in One-Counter Systems. In *Proc. of FOSSACS'16*, pages 462–478, 2016.
- 4 Dmitry Chistikov and Christoph Haase. The Taming of the Semi-Linear Set. In *Proc. ICALP'16*, pages 128:1–128:13, 2016.
- 5 Wojciech Czerwinski and Slawomir Lasota. Regular separability of one counter automata. In *LICS'17*, pages 1–12, 2017.
- 6 Matthias Englert, Ranko Lazic, and Patrick Totzke. Reachability in Two-Dimensional Unary Vector Addition Systems with States is NL-Complete. In *Proc. of LICS '16*, pages 477–484, 2016.
- 7 John E. Hopcroft and Jean-Jacques Pansiot. On the Reachability Problem for 5-Dimensional Vector Addition Systems. *Theor. Comput. Sci.*, 8:135–159, 1979. doi:10.1016/0304-3975(79)90041-0.
- 8 Richard M. Karp and Raymond E. Miller. Parallel Program Schemata. *J. Comput. Syst. Sci.*, 3(2):147–195, 1969.
- 9 S. Rao Kosaraju. Decidability of Reachability in Vector Addition Systems (Preliminary Version). In *STOC'82*, pages 267–281, 1982.
- 10 Michel Latteux. Langues à un Compteur. *J. Comput. Syst. Sci.*, 26(1):14–33, 1983.
- 11 Ernst W. Mayr. An Algorithm for the General Petri Net Reachability Problem. In *STOC'81*, pages 238–246, 1981.