

Reachability in Symmetric VASS

Łukasz Kamiński 

University of Warsaw, Poland

Sławomir Lasota 

University of Warsaw, Poland

Abstract

We investigate the reachability problem in *symmetric* vector addition systems with states (VASS), where transitions are invariant under a group of permutations of coordinates. One extremal case, the trivial groups, yields general VASS. In another extremal case, the symmetric groups, we show that the reachability problem can be solved in PSPACE, regardless of the dimension of input VASS (to be contrasted with Ackermannian complexity in general VASS). We also consider other groups, in particular alternating and cyclic ones. Furthermore, motivated by the open status of the reachability problem in data VASS, we estimate the gain in complexity when the group arises as a combination of the trivial and symmetric groups.

2012 ACM Subject Classification Theory of computation → Distributed computing models

Keywords and phrases vector addition systems, Petri nets, reachability problem, symmetry, permutation group

Digital Object Identifier 10.4230/LIPIcs.MFCS.2025.60

Related Version *Full Version*: <https://arxiv.org/abs/2506.23578>

Funding *Łukasz Kamiński*: Partially supported by NCN grant 2021/41/B/ST6/00535.

Sławomir Lasota: Partially supported by the ERC grant INFSYS, agreement no. 950398.

1 Introduction

Petri nets, equivalently presentable as vector addition systems with states (VASS), are an established model of concurrency with widespread applications. The central algorithmic problem for this model is the *reachability problem* which asks whether from a given initial configuration there exists a sequence of valid execution steps reaching a given final configuration. The decidability of the problem was established in 1981 by Mayr [19], and subsequently improved by Kosaraju [11] and Lambert [12]. The exponential space lower bound was shown by Lipton already in 1976 [18]. For around 40 years these were the only known complexity bounds, and the complexity of the problem remained for a long time as one of the hardest open questions in the verification of concurrent systems. Only in the last few years we have seen a sudden progress, namely an Ackermannian upper bound [17], a breakthrough non-elementary lower bound [3, 4], and finally the matching Ackermannian lower bound (independently [5] and [16]).

All the above-mentioned bounds are more fine-grained and apply to the problem parametrised by the dimension, where the input is restricted to d -dimensional VASS (d -VASS) for fixed $d \in \mathbb{N}$. Both the upper and lower bounds have been subsequently further improved [14, 2, 9], and currently the reachability problem for d -VASS is known to belong to the complexity class \mathcal{F}_d , and for $(2d + 3)$ -VASS it is known to be \mathcal{F}_d -hard. Here \mathcal{F}_d denotes the d th level of the hierarchy of complexity classes corresponding to Grzegorzczuk’s fast growing function hierarchy. For instance, the currently best lower bound of [2] shows the reachability problem for 9-VASS to be hard for the class $\mathcal{F}_3 = \text{TOWER}$. Subsequently, [6] proved the same complexity bound even for 8-VASS. Summing up, for sufficiently large dimensions the lower bound is prohibitive.



© Łukasz Kamiński and Sławomir Lasota;

licensed under Creative Commons License CC-BY 4.0

50th International Symposium on Mathematical Foundations of Computer Science (MFCS 2025).

Editors: Paweł Gawrychowski, Filip Mazowiecki, and Michał Skrzypczak; Article No. 60; pp. 60:1–60:17

Leibniz International Proceedings in Informatics



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

Symmetric VASS. In this paper we investigate subclasses of d -VASS which are *symmetric*, meaning their sets of transitions are invariant under certain permutations of the coordinates $\{1, \dots, d\}$. Our study is parametric in the choice of a group of permutations $G \leq S_d$ of coordinates, and our objective is to analyze the gain in complexity to be achieved when the input to the reachability problem is restricted to G -VASS, the subclass of those d -VASS whose sets of transitions are invariant under permutations $\sigma \in G$. In one extreme case, when G is the trivial permutation group I_d containing just the identity permutation, G -VASS are just general d -VASS. In another extreme case, when $G = S_d$ is the symmetric group of degree d , transitions of G -VASS are invariant under all permutations of coordinates. The two extreme cases may be combined. For instance, one can consider the permutation group $G = I_n \wr S_d \leq S_{nd}$ of degree nd containing all permutations that are identity inside each of the n -element blocks:

$$\{1, \dots, n\}, \quad \{n+1, \dots, 2n\}, \quad \dots, \quad \{n(d-1)+1, \dots, nd\},$$

but permute arbitrarily the whole blocks.

Our principal motivation is to exploit symmetry of a model in order to lower the prohibitively high complexity of the reachability problem in the general case. Another motivation comes from the model of *data VASS* [13, 15, 20] extending plain VASS with data. In terms of Petri nets, the extension allows tokens to carry data value coming from some infinite countable set, and allows transitions to test (dis)equalities between data values of involved tokens. When $G = I_n \wr S_d \leq S_{nd}$, the G -VASS are exactly data VASS of dimension n , where the set of data values is restricted to be finite¹, of size d . It is not known if the reachability problem is decidable for data VASS, and our current study may shed some more light on this hard open problem.

Contribution. Intuitively speaking, the larger the group G , the lower the complexity of the reachability problem, since whenever $G \leq H$, every H -VASS is automatically a G -VASS. The main contribution of this paper is twofold. First, we concentrate on the potentially easiest cases, namely symmetric groups S_d , and discover a huge complexity drop: regardless of dimension $d \geq 2$, the reachability problem is PSPACE-complete for VASS invariant under symmetric groups.² We also prove the same complexity for another potentially easy case, namely for the alternating groups A_d : again, regardless of dimension $d \geq 3$, the reachability problem is PSPACE-complete for VASS invariant under alternating groups. Both the PSPACE decision procedures are designed for a fixed dimension d , but work equally well when d is part of input. We thus get uniform PSPACE upper bound for VASS invariant under symmetric (resp. alternating) groups. At the other side, we prove that cyclic groups Z_d are hard, namely the reachability problem for d -VASS invariant under Z_d is as hard as in $\Theta(d)$ -VASS.

The case of trivial group $G = I_d$ coincides with general d -VASS. As our second contribution we investigate combinations of symmetric groups and trivial groups, providing one positive and one negative result. On one hand, when $G = I_n \wr S_d \leq S_{nd}$ (as discussed above) and $n \geq 2$, the reachability problem is as hard as in case of $(n-1)d$ -VASS. From the perspective of data VASS this may be interpreted as a bad news: the complexity grows with the increasing number d of data values. On the other hand, we investigate another combination of the

¹ The other way around, the n -dimensional data VASS can be defined as G -VASS, where $G = I_n \wr S_\omega$ combines the trivial group I_n with the symmetric group S_ω of infinite countable degree.

² In dimension 1 the permutation group G is irrelevant, and the problem is NL- or NP-complete, for input represented in unary or binary, respectively [10].

two groups, namely $G = S_d \wr I_n \leq S_{dn}$ containing all permutations that independently and arbitrarily permute each of d -element blocks:

$$\{1, \dots, d\}, \quad \{d+1, \dots, 2d\}, \quad \dots, \quad \{d(n-1)+1, \dots, dn\},$$

but preserve each of the blocks. In this case we provide an exponential-time reduction to n -VASS, i.e., this time the complexity is independent of the degree d of the symmetric group. The complexity in this case is thus (significantly) lower than for $I_n \wr S_d$. This is in agreement with $I_n \wr S_d$ being a subgroup of $S_d \wr I_n$, up to isomorphism of permutation groups.

2 Symmetric VASS

As usual, let \mathbb{Z}, \mathbb{N} denote integers and nonnegative integers, respectively. The value of the i th coordinate of a vector $\mathbf{w} \in \mathbb{Z}^d$ is written as $\mathbf{w}(i)$, namely $\mathbf{w} = (\mathbf{w}(1), \dots, \mathbf{w}(d))$, i.e., we identify a vector with a function $\mathbf{w}: [d] \rightarrow \mathbb{Z}$, where $[d] = \{1, \dots, d\}$. The *norm* of a vector $\mathbf{v} \in \mathbb{Z}^d$ is the maximum of absolute values of its coordinates, i.e., $\|\mathbf{v}\| = \max_{i \in [d]} |\mathbf{v}(i)|$. Then the *norm* of a set of vectors X is defined as $\|X\| = \max_{\mathbf{v} \in X} \|\mathbf{v}\|$. For $n \in \mathbb{Z}$, by \bar{n} we denote the constant vector $\bar{n} = (n, \dots, n) \in \mathbb{Z}^d$, in the dimension d to be always determined by the context. For a finite set X , by $|X| \in \mathbb{N}$ we denote its size.

Vector addition systems with states. For $d \in \mathbb{N}_+ = \mathbb{N} \setminus \{0\}$, we define d -dimensional *vector addition systems with states*, denoted shortly as d -VASS, or simply VASS when the dimension d is irrelevant or clear from the context. A d -VASS $\mathcal{V} = (Q, T)$ consists of a finite set Q of states and a finite set of transitions $T \subseteq Q \times \mathbb{Z}^d \times Q$. A *configuration* c of \mathcal{V} consists of a state $q \in Q$ and a nonnegative vector $\mathbf{w} \in \mathbb{N}^d$, and is written as $c = q(\mathbf{w})$. A transition $t = (q, \mathbf{v}, q')$ induces *steps*

$$q(\mathbf{w}) \xrightarrow{t} q'(\mathbf{w}') \tag{1}$$

between configurations, where $\mathbf{w}' = \mathbf{w} + \mathbf{v}$. We refer to the vector $\mathbf{v} \in \mathbb{Z}^d$ as the *effect* of the transition (q, \mathbf{v}, q') or of an induced step. The *norm* of a transition $t = (q, \mathbf{v}, q')$ is defined as $\|t\| := \|\mathbf{v}\|$. Similarly, for a configuration $c = q(\mathbf{w})$, we define $\|c\| := \|\mathbf{w}\|$. The *norm* of a VASS \mathcal{V} is the maximal norm of its transitions, i.e., $\|\mathcal{V}\| := \max_{t \in T} \|t\|$.

A *run* π in \mathcal{V} is a sequence of steps with the proviso that the target configuration of every step matches the source configuration of the next one:

$$\pi = c_0 \xrightarrow{t_1} c_1 \longrightarrow \dots \xrightarrow{t_n} c_n. \tag{2}$$

We say that the run π is *from* c_0 *to* c_n , call c_0, c_n the source and the target configuration of the run, respectively, and write $c_0 \xrightarrow{\pi} c_n$. When the source configuration c_0 is clear from the context, we may identify a run (2) from c_0 with the sequence $t_1 \dots t_n$ of transitions fired. The *effect* of a run is the sum of effects of all steps, and its *length* is the number n of steps. We also write $c \longrightarrow c'$ if there is some run from c to c' . We often use the natural operation of concatenation of runs: if $\alpha: c \longrightarrow c'$ and $\beta: c' \longrightarrow c''$ then $(\alpha; \beta): c \longrightarrow c''$, assuming that the target configuration of the former run coincides with the source configuration of the latter one. For convenience, we allow ourselves to drop this assumption and use the convention that the latter run β is implicitly shifted so that its source matches the target of α , as long as the shifted β is still a run. Under this convention, we write α^n for n -fold concatenation of α .

In the sequel we refer to the *state graph* of a VASS $\mathcal{V} = (Q, T)$, a directed graph (Q, E) whose nodes are states Q , and edges $E \subseteq Q^2$ are those pairs (q, q') for which there is a transition $(q, \mathbf{v}, q') \in T$, for some $\mathbf{v} \in \mathbb{Z}^d$. This graph may contain self-loops.

Reachability problem. One of the most fundamental computational problems studied in the setting of VASS is the *reachability problem* (for the sake of this paper we prefer to formulate the problem for a *fixed* dimension d):

d -VASS-REACH: given a d -VASS \mathcal{V} together with two configurations, source s and target t , one asks if \mathcal{V} has a run from s to t .

Complexity of this problem in subclasses of *symmetric* VASS, to be defined below, is the main topic studied in this paper. An input to the problem is a triple (\mathcal{V}, s, t) . For simplicity, such a triple (\mathcal{V}, s, t) , or a pair (\mathcal{V}, s) with just the source state, we also call a VASS, hoping that this does not lead to confusion. We also use shorthands $\|\mathcal{V}, s\| := \max\{\|\mathcal{V}\|, \|s\|\}$ and $\|\mathcal{V}, s, t\| := \max\{\|\mathcal{V}\|, \|s\|, \|t\|\}$.

Symmetric VASS. Let S_d denote the symmetric group containing all permutations of $[d]$. Permutations of $[d]$ act on vectors $\mathbf{w} \in \mathbb{Z}^d$ by permuting dimensions, namely $\sigma \in S_n$ maps \mathbf{w} to the vector $\sigma(\mathbf{w})$ defined by $\sigma(\mathbf{w})(i) = \mathbf{w}(\sigma^{-1}(i))$, or equivalently $\sigma(\mathbf{w})(\sigma(i)) = \mathbf{w}(i)$. Treating a vector as a function $\mathbf{w}: [d] \rightarrow \mathbb{Z}$, we may write

$$\sigma(\mathbf{w}) = \mathbf{w} \circ \sigma^{-1}, \quad \text{or equivalently} \quad \sigma(\mathbf{w}) \circ \sigma = \mathbf{w}.$$

We naturally extend this action to transitions, $t = (q, \mathbf{v}, q') \mapsto \sigma(t) = (q, \sigma(\mathbf{v}), q')$, and likewise to steps, runs, etc., as expected. Given a group of permutations $G \leq S_d$, by a G -VASS, we mean any d -VASS $\mathcal{V} = (Q, T)$ whose transitions are invariant under the action of G , namely for every permutation $\sigma \in G$ and transition $t \in T$ we have $\sigma(t) \in T$. Thus, for every $G \leq S_d$, we identify a subclass of d -VASS admitting the invariance property. In the most restrictive case of $G = S_d$, we get the subclass of d -VASS invariant under all permutations of dimensions. In the less restrictive case of trivial group $I_d = \{\text{Id}_d\}$, we get the whole class of all d -VASS.

Reachability problem for symmetric VASS. In our investigations we always assume that a group $G \leq S_d$ (and hence also the dimension of VASS) is fixed.³ The VASS reachability problem, when its input is restricted to G -VASS for a fixed permutation group G , we call G -VASS-REACH. For the trivial group $G = I_d = \{\text{Id}_d\}$ we get d -VASS-REACH.

The set of transitions of a G -VASS is determined uniquely by representatives of orbits, i.e., by representatives of the equivalence relation on transitions: $t \equiv_G t'$ if $t = \sigma(t')$ for some $\sigma \in G$. Unfolding the action of G on transitions, $(q, \mathbf{v}, p) \equiv_G (q', \mathbf{v}', p')$ if and only if $q = q'$, $p = p'$, and $\mathbf{v} = \sigma(\mathbf{v}')$ for some $\sigma \in G$. Equivalence classes of \equiv_G are called G -orbits of transitions, or simply *orbits* when the group is clear from the context. When measuring the size of input we assume a succinct representation, where transitions T of a G -VASS are given by orbit representatives, i.e., by a subset $\tilde{T} \subseteq T$ containing one transition from every orbit. By the *size* of such a representation (Q, \tilde{T}) of a G -VASS \mathcal{V} we mean the bitsize of its description, e.g.,

$$|\mathcal{V}| = |\tilde{T}| \cdot (2 \cdot |Q| + d \cdot \lceil \log(2 \cdot |\tilde{T}| + 1) \rceil),$$

where $|\tilde{T}| = \max_{t \in \tilde{T}} \|t\|$. Note that $|\mathcal{V}|$ is independent of a choice of orbit representatives, and logarithmic in terms of $|\tilde{T}|$.

³ Nevertheless, all our subsequent PSPACE decision procedures are uniform with respect to dimension, namely they work within the same complexity bounds if dimension d is part of input.

Whenever $G' \leq G \leq S_d$, i.e., G' is a permutation subgroup of G , (a representation of) a G -VASS may be canonically transformed into (a representation of) a G' -VASS in two different ways. Let $\mathcal{V} = (Q, T)$ be a G -VASS. On one hand, \mathcal{V} is itself a G' -VASS, as its transitions are automatically invariant under the action of G' , but the representation of \mathcal{V} seen as a G' -VASS may be exponentially larger than its representation when seen as G -VASS (since the number of G' -orbits of T may be exponentially larger than the number of G -orbits of T). This yields an exponential (but not necessarily polynomial) reduction from G -VASS-REACH to G' -VASS-REACH. On the other hand, a *representation* of \mathcal{V} is automatically a representation of another G' -VASS $\mathcal{V}' = (Q, T')$, where the set of transitions $T' \subseteq T$ may be of exponentially smaller size than T , but the number of G' -orbits of T' is the same as the number of G -orbits of T . Clearly, every run of \mathcal{V}' is a run of \mathcal{V} , but the converse does not hold in general.

\mathbb{Z} -reachability. An important relaxation of reachability is \mathbb{Z} -reachability, where one considers a relaxed notion of configurations, namely \mathbb{Z} -configurations $Q \times \mathbb{Z}^d$ instead of $Q \times \mathbb{N}^d$. Contrary to configurations, values in \mathbb{Z} -configurations may drop below 0. We define \mathbb{Z} -runs exactly like runs, but using \mathbb{Z} -configurations instead of configurations. Each \mathbb{Z} -run, and hence also each run, induces a unique path in the state graph. Conversely, every path in the state graph gives rise to multiple different \mathbb{Z} -runs, depending on the choice of source vector, and of transitions witnessing consecutive edges of a path.

The \mathbb{Z} -reachability problem asks, given a G -VASS (\mathcal{V}, s, t) , if \mathcal{V} has a \mathbb{Z} -run from s to t . We prove that the problem lies in NP, regardless of the choice of dimension d of VASS and of the group $G \leq S_d$. This observation will be used later in proofs of main results.

► **Lemma 1.** *For every permutation group G , the \mathbb{Z} -reachability problem for G -VASS is NP-complete.*

The lemma does not follow immediately from NP upper bound for \mathbb{Z} -reachability in VASS, as the input is represented more succinctly in our case. In the proof we use the integer analogue of Caratheodory's theorem [7] to get a small solution property: if there is a \mathbb{Z} -run between two given configurations then there is one using only a subset of transitions of polynomial size. By guessing this subset we reduce to the \mathbb{Z} -reachability problem in VASS, which is in NP.

3 Lower bound

This section contains a proof of PSPACE-hardness of the reachability problem in G -VASS, for any $G \leq S_d$ with $d \geq 2$. Later, we will indicate cases where this lower bound is tight, for instance $G = S_d$.

► **Lemma 2.** *For every $d \geq 2$ and $G \leq S_d$, G -VASS-REACH is PSPACE-hard.*

Proof. We show a reduction from the *bounded reachability* problem for 1-VASS, which is known to be PSPACE-complete [8]. The input to this problem consists of a 1-VASS (\mathcal{V}, s, f) and a number $M \in \mathbb{N}$ represented in binary, and we ask if \mathcal{V} has a run $s \rightarrow^* f$ with all configurations bounded by M . For simplicity, we assume that the vector of a configuration of a 1-VASS $q(n)$ is just a nonnegative integer $n \in \mathbb{N}$, and the effect of a transition (q, z, p) is just an integer $z \in \mathbb{Z}$. W.l.o.g. we assume that effects of transitions of a 1-VASS are nonzero and at most M .

Fix an arbitrary dimension $d \geq 2$. Given a 1-VASS (\mathcal{V}, s, f) and M , where $\mathcal{V} = (Q, T)$, we construct a G -VASS $\mathcal{V}' = (Q', T')$ of dimension d whose states $Q' := Q \cup \{\langle q, t \rangle \mid q \in Q, t \in T\}$ extend Q by additional auxiliary states, and whose transitions T' are defined below. The idea is to represent a configuration $q(n)$ of \mathcal{V} bounded by M ($n \leq M$), by a configuration

$$\overline{q(n)} := q(n, \dots, n, 2M - n)$$

of \mathcal{V}' , while pairs $\langle q, t \rangle$ are additional auxiliary states. We define the set T' of transitions by providing a set of orbit representatives $\tilde{T} \subseteq T'$. Below, we use $z \in \mathbb{N}_+$ to range over positive integers, and use the vectors

$$\mathbf{v}_z = (M + z, \dots, M + z, -M - z) \quad \mathbf{v}_{-z} = (-M - z, \dots, -M - z, M + z)$$

that have value $M + z$ (resp. $-M - z$) on all coordinates except coordinate d which has the opposite value. We also use $\mathbf{v} = (M, \dots, M, -M)$. The set \tilde{T} is defined as follows:

- For every $t = (p, z, q) \in T$, we put into \tilde{T} two transitions (recall that $z \in \mathbb{N}_+$)

$$\tilde{t} = (p, \mathbf{v}_z, \langle q, t \rangle) \quad t^- = (\langle q, t \rangle, -\mathbf{v}, q);$$

- For every $t = (p, -z, q) \in T$, we put into \tilde{T} two transitions

$$t^+ = (p, \mathbf{v}, \langle p, t \rangle) \quad \tilde{t} = (\langle p, t \rangle, \mathbf{v}_{-z}, q).$$

As all vectors $\mathbf{v}_z, \mathbf{v}_{-z}$ and \mathbf{v} are constant except for one coordinate, we notice that the size of $T' = \{\sigma(t) \mid t \in \tilde{T}, \sigma \in G\}$ is at most $|T'| \leq d \cdot |\tilde{T}|$. This is however not relevant for the reduction, as \mathcal{V}' is represented by \tilde{T} , and not by T' .

We argue that there is a run $s \rightarrow f$ in \mathcal{V} whose all configurations are bounded by M , if and only if $\bar{s} \rightarrow \bar{f}$ in \mathcal{V}' . For the “only if” direction we observe that every step $c \xrightarrow{t} c'$ in \mathcal{V} which is bounded by M is simulated by two steps $\bar{c} \rightarrow c' \rightarrow \bar{c}'$ in \mathcal{V}' , for some intermediate configuration c' with an auxiliary state. Indeed, when $t = (p, z, q)$, a step $p(n) \xrightarrow{t} q(n + z)$ of \mathcal{V} which is bounded by M ($n + z \leq M$), is simulated in \mathcal{V}' by the following two steps:

$$\overline{p(n)} \xrightarrow{\tilde{t}} \langle q, t \rangle (M + n + z, \dots, M + n + z, M - n - z) \xrightarrow{t^-} \overline{q(n + z)}. \quad (3)$$

In the other case, when $t = (p, -z, q)$, a step $p(n) \xrightarrow{t} q(n - z)$ of \mathcal{V} which is bounded by M ($n \leq M$), is simulated in \mathcal{V}' by the following two steps:

$$\overline{p(n)} \xrightarrow{t^+} \langle p, t \rangle (M + n, \dots, M + n, M - n) \xrightarrow{\tilde{t}} \overline{q(n - z)}. \quad (4)$$

For the “if” direction, we consider a configuration $\overline{q(n)} = q(n, \dots, n, 2M - n)$, and analyze all possible two-step runs in \mathcal{V}' starting from $\overline{q(n)}$: $\overline{q(n)} \rightarrow c' \rightarrow c''$. By definition of T' , we deduce that the only possible such two-step runs are of the form (3) or (4). Indeed, the first transition may be only \tilde{t} for some $t = (q, z, p) \in T$, or t^+ for some $t = (q, -z, p) \in T$, and then the next transition is uniquely determined by the auxiliary state. In particular, we observe that only transitions from \tilde{T} can be used, i.e., all other transitions $T' \setminus \tilde{T}$ are useless if one starts from a configuration of the form $\overline{p(n)}$. In consequence, if $\bar{s} \rightarrow \bar{f}$ in \mathcal{V}' , we have also a run $s \rightarrow f$ in \mathcal{V} whose all configurations are bounded by M , as required. Correctness of the reduction follows. \blacktriangleleft

► **Remark 3.** In particular, PSPACE-hardness for S_2 -VASS improves upon PSPACE-hardness for 2-VASS [1].

4 Transitive and fair groups

Within this section we fix an arbitrary *transitive* group $G \leq S_d$, i.e., we assume that for every $i, j \in [d]$ there exists $\sigma \in G$ such that $\sigma(i) = j$. The results of this section will be applied later to specific transitive groups G .

Pumpable VASS. Consider a G -VASS (\mathcal{V}, s, t) , where $s = q(\mathbf{w})$ and $t = q'(\mathbf{w}')$. We say that (\mathcal{V}, s) is *forward pumpable* if \mathcal{V} has a run $q(\mathbf{w}) \rightarrow q(\mathbf{w} + \mathbf{e})$ for some $\mathbf{e} \geq \bar{1}$. Symmetrically, (\mathcal{V}, t) is *backward pumpable* if $(\mathcal{V}^{\text{rev}}, t)$ is forward pumpable where \mathcal{V}^{rev} , the *reverse* of \mathcal{V} , is obtained by replacing each transition (q, \mathbf{v}, q') of \mathcal{V} by its reverse $(q', -\mathbf{v}, q)$. Equivalently, \mathcal{V} has a run $q'(\mathbf{w}' + \mathbf{e}') \rightarrow q'(\mathbf{w}')$ for some $\mathbf{e}' \geq \bar{1}$. Finally, (\mathcal{V}, s, t) is *pumpable* if (\mathcal{V}, s) is forward pumpable, and (\mathcal{V}, t) is backward pumpable. We prove that reachability reduces to \mathbb{Z} -reachability, when the group G is transitive and VASS are pumpable.

► **Lemma 4.** *If G is transitive then every pumpable G -VASS (\mathcal{V}, s, t) admitting a \mathbb{Z} -run from s to t , admits a run from s to t .*

Proof. Let $G \leq S_d$. We start by proving the following equality for all $i, j \in [d]$:

$$|\{\sigma \in G \mid \sigma(i) = j\}| = \frac{|G|}{d}. \quad (5)$$

When $i = j$, the equality follows immediately from the Orbit-Stabilizer Theorem. Otherwise, suppose $i \neq j$ and take any permutation $\sigma_{ij} \in G$ such that $\sigma_{ij}(i) = j$ (it exists as G is transitive). By post-composing with σ_{ij} we get a bijection $\sigma \mapsto \sigma_{ij} \circ \sigma$ between permutations $\sigma \in G$ that satisfy $\sigma(i) = i$, and permutations $\sigma \in G$ that satisfy $\sigma(i) = j$. In consequence, we get the equality (5):

$$|\{\sigma \in G \mid \sigma(i) = j\}| = |\{\sigma \in G \mid \sigma(i) = i\}| = \frac{|G|}{d}.$$

Consider a G -VASS \mathcal{V} and two configurations $s = q(\mathbf{v})$ and $t = q'(\mathbf{v}')$, and suppose \mathcal{V} has a \mathbb{Z} -run γ from $q(\mathbf{v})$ to $q'(\mathbf{v}')$, and two runs

$$\alpha: q(\mathbf{v}) \rightarrow q(\mathbf{v} + \mathbf{e}) \quad \alpha': q'(\mathbf{v}' + \mathbf{e}') \rightarrow q'(\mathbf{v}'),$$

for some $\mathbf{e}, \mathbf{e}' \geq \bar{1}$. We argue that \mathcal{V} has a run $q(\mathbf{v}) \rightarrow q'(\mathbf{v}')$. W.l.o.g. we can assume that α is *increasing*, i.e., every coordinate of $\mathbf{v} + \mathbf{e}$ is greater than the greatest coordinate of \mathbf{v} (otherwise we replace α by its n -fold concatenation $\alpha^n: q(\mathbf{v}) \rightarrow q(\mathbf{v} + n\mathbf{e})$, for sufficiently large $n \in \mathbb{N}$). Symmetrically, we can assume that α' is *decreasing*, i.e., every coordinate of $\mathbf{v}' + \mathbf{e}'$ is greater than the greatest coordinate of \mathbf{v}' .


Let $G = \{\sigma_1, \dots, \sigma_m\}$, where $\sigma_1 = \text{Id}_d$ is the identity permutation on $[d]$, and let

$$\Delta := \sum_{i=1}^m \sigma_i(\mathbf{e}), \quad \Delta' := \sum_{i=1}^m \sigma_i(\mathbf{e}').$$

As α is increasing and $\sigma_1(\alpha) = \alpha$, the concatenation of runs

$$\tilde{\alpha} = \sigma_1(\alpha); \sigma_2(\alpha); \dots; \sigma_m(\alpha)$$

is forcedly a run from $q(\mathbf{v})$ to $q(\mathbf{v} + \Delta)$. Symmetrically, as α' is decreasing and $\sigma_1(\alpha') = \alpha'$, the concatenation of runs $\tilde{\alpha}' = \sigma_m(\alpha'); \dots; \sigma_2(\alpha'); \sigma_1(\alpha')$ is forcedly a run from $q'(\mathbf{v}' + \Delta')$



■ **Figure 1** A sequential VASS.

to $q'(\mathbf{v}')$. We argue that $\Delta(j) = \Delta(k)$ for every $j, k \in [d]$, i.e., the effect of $\tilde{\alpha}$ is a constant vector. Indeed, by point-wise expansion of the equality

$$\Delta = \sum_{i=1}^m \sigma_i(\mathbf{e}) = \sum_{i=1}^m \mathbf{e} \circ \sigma_i,$$

and using the equality (5) to deduce that for every ℓ , the value of ℓ th coordinate $\mathbf{e}(\ell)$ of \mathbf{e} appears the same number of times on both sides of the equality below, we compute:

$$\Delta(j) = \sum_{\ell=1}^d \sum_{i: \sigma_i(\ell)=j} \mathbf{e}(\ell) \stackrel{(5)}{=} \sum_{\ell=1}^d \sum_{i: \sigma_i(\ell)=k} \mathbf{e}(\ell) = \Delta(k).$$

Therefore, $\Delta = \bar{n}$ for some $n \in \mathbb{N}_+$. Analogously we prove that $\Delta' = \bar{n'}$ for some $n' \in \mathbb{N}_+$. Let $\tilde{\Delta} = \overline{(n \cdot n')}$. Since $\tilde{\Delta} = n' \cdot \Delta = n \cdot \Delta'$, we may construct two runs that pump forward and backward by the same vector $\tilde{\Delta}$:

$$\beta = \tilde{\alpha}^{(n')}: q(\mathbf{v}) \longrightarrow q(\mathbf{v} + \tilde{\Delta}) \quad \beta' = (\tilde{\alpha}')^n: q'(\mathbf{v}' + \tilde{\Delta}) \longrightarrow q'(\mathbf{v}').$$

Thus, for sufficiently large $m \in \mathbb{N}_+$, the m -fold iterations of β and β' are enough to lift the \mathbb{Z} -run γ into a run, i.e., to ensure that $\beta^m; \gamma; (\beta')^m$ is a run from $q(\mathbf{v})$ to $q'(\mathbf{v}')$. ◀

Fair groups. Given $d \in \mathbb{N}_+$ and a transitive permutation group $G \leq S_d$, we say that G is *fair* if there is a polynomial P such that for every $R \in \mathbb{N}$ and every G -VASS (\mathcal{V}, s) with $S := \|\mathcal{V}, s\|$ and $N := \|\mathcal{V}\|$, the implication (1) \implies (2) holds, where

- (1) \mathcal{V} has a run $s \longrightarrow p(\mathbf{w})$ for some $p \in Q$ and \mathbf{w} of norm $\|\mathbf{w}\| > P(N) \cdot (S + R)$,
- (2) \mathcal{V} has a run $s \longrightarrow p(\mathbf{w})$ for some $p \in Q$ and $\mathbf{w} \geq \overline{S + R + 1}$.

In words: whenever (\mathcal{V}, s) has a run whose target vector exceeds $P(N) \cdot (S + R)$ on *some* coordinate, (\mathcal{V}, s) is guaranteed to have a run whose target vector is greater than $S + R$ on *all* coordinates; and the property holds for any chosen $R \in \mathbb{N}$.

► **Lemma 5.** *If G is transitive and fair, then G -VASS-REACH is in PSPACE.*

Proof. We start with some preparatory definitions. A VASS is called *strongly connected* if its state graph is so. A VASS $\mathcal{V} = (Q, T)$ is called *sequential*, if it can be partitioned into a number of strongly connected VASS $\mathcal{V}_1 = (Q_1, T_1), \dots, \mathcal{V}_k = (Q_k, T_k)$ with pairwise disjoint state spaces, called *components* of \mathcal{V} , and $k - 1$ transitions $u_i = (q_i, \mathbf{v}_i, q'_i)$, for $i \in [k - 1]$, where $q_i \in Q_i$ and $q'_i \in Q_{i+1}$, called *bridges* of \mathcal{V} (see Fig. 1). For every non-sequential VASS (\mathcal{V}, s, t) there is a finite set $W = \{\mathcal{V}^1, \dots, \mathcal{V}^\ell\}$ such that each \mathcal{V}^j is a sequential VASS, and $s \longrightarrow t$ in \mathcal{V} if and only if $s \longrightarrow t$ in \mathcal{V}^j for some $j \in [\ell]$. The set W , called *sequential decomposition* of \mathcal{V} in the sequel, may be computed based on the decomposition of the state graph of \mathcal{V} into strongly connected components.

Assume an input G -VASS (\mathcal{V}, s, t) , where $\mathcal{V} = (Q, T)$, $s = q(\mathbf{v})$ and $t = p(\mathbf{w})$, represented by orbit representatives $\tilde{T} \subseteq T$. The PSPACE decision procedure to check if $s \longrightarrow t$, to be described below, proceeds in several steps.

As the first step, the decision procedure decomposes the state graph of \mathcal{V} into strongly connected components, and nondeterministically picks up one sequential VASS from the sequential decomposition of (\mathcal{V}, s, t) . From now on we may thus assume, w.l.o.g., that \mathcal{V} is

sequential with k components $\mathcal{V}_1, \dots, \mathcal{V}_k$, the configuration s belongs to the first component, and t to the last one. The number k of components is bounded by the number $|Q|$ of states of \mathcal{V} . Let $R := |Q| \cdot \|\mathcal{V}\|$. Thus, for every two states p, p' from the same component of \mathcal{V} , there is a path from p to p' that decreases each counter by at most R .

For $M \in \mathbb{N}_+$, we say that a configuration $p(\mathbf{w})$ is *bounded* by M if $\|\mathbf{w}\| \leq M$. We also say that the first component (\mathcal{V}_1, s) is bounded by M if every configuration c reachable in this component from s (i.e., c satisfying $s \rightarrow c$) is bounded by M . Symmetrically, we say that the last component (\mathcal{V}_k, t) is *reverse bounded* by M if every configuration c reverse reachable from t (i.e., c satisfying $c \rightarrow t$) in this component is bounded by M .

As the second step, the decision procedure iteratively removes either the first or the last component from (\mathcal{V}, s, t) until the following two conditions are met, for $N := \|\mathcal{V}\|$ and $S = \|\mathcal{V}, s, t\|$:

- the first component is *not bounded* by $P(N) \cdot (S + R)$, and
- the last component is *not reverse bounded* by $P(N) \cdot (S + R)$,

where P is the polynomial given by the definition of a fair group. In particular, the first condition implies (1) for (\mathcal{V}_1, s) , and the second one implies (1) for $(\mathcal{V}_k^{\text{rev}}, t)$. Suppose the first component is *bounded* by $P(N) \cdot (S + R)$ (the symmetric case, when the last component is reverse bounded by $P(N) \cdot (S + R)$, is treated in the same way). One iteration of the removal procedure proceeds as follows, depending on whether $k = 1$ or $k > 1$:

Case I: $k = 1$. If this component is the only one, the procedure checks in PSPACE if $s \rightarrow t$, and terminates. The check is done by a standard nondeterministic walk through the graph of configurations reachable from s , which starts at the source configuration s , and in every iteration moves to a nondeterministically chosen successor of the current configuration. It terminates when t is reached (and answers positively), or when the counter of so far visited configurations exceeds

$$|Q| \cdot (P(N) \cdot (S + R) + 1)^d,$$

the total number of configurations bounded by $P(N) \cdot (S + R)$. The memory size, namely the bitsize of a single configuration, and of the value of the counter, are both polynomially bounded, and hence the whole procedure works in (nondeterministic) PSPACE.

Case II: $k > 1$. Otherwise, let $u_1 = (p, \mathbf{u}, p')$ be the first bridge. We observe that if \mathcal{V} has a run $s \rightarrow t$, then its last configuration $p(\mathbf{w})$ in the first component is necessarily bounded by $P(N) \cdot (S + R)$. In consequence, since the bridge transition may increase the norm by at most N , the first configuration $p'(\mathbf{w}')$ of the run in the second component, where $\mathbf{w}' = \mathbf{w} + \mathbf{u}$, is forcedly bounded by $P(N) \cdot (S + R) + N$. Relying on this observation, the decision procedure nondeterministically chooses a vector \mathbf{w} with $\|\mathbf{w}\| \leq P(N) \cdot (S + R)$, and checks in PSPACE if $q(\mathbf{v}) \rightarrow p(\mathbf{w})$ in the first component, similarly as above. Then the procedure removes the first component from \mathcal{V} , and takes $s' = p'(\mathbf{w}') = p'(\mathbf{w} + \mathbf{u})$ as the source configuration, thus obtaining a new G -vass (\mathcal{V}', s', t) with one less component than (\mathcal{V}, s, t) .

This removal is iteratively repeated until either the procedure terminates, or produces a (\mathcal{V}', s', t) with the first component *not bounded* by $P(N') \cdot (S' + R)$, and the last component *not reverse bounded* by $P(N') \cdot (S' + R)$, where $N' := \|\mathcal{V}'\|$ and $S' := \|\mathcal{V}', s', t'\|$. The bound on the norm N' of the vass cannot increase, and so $N' \leq N$. However, in each iteration the bound on norm S' of the new source vector \mathbf{w}' increases, compared to the bound S on the norm of the previous source vector \mathbf{v} , by a multiplicative factor $\mathcal{O}(P(N) \cdot |Q|)$, namely

$$S' \leq P(N) \cdot (S + R) + N \leq (P(N) \cdot (1 + |Q|) + 1) \cdot S,$$

and therefore the bitsize of the source vector increases by (bounding roughly) at most $\mathcal{O}(\log(P(N) \cdot |Q|)) \cdot d$. As the number of iterations is bounded by $|Q|$, the final bitsize is polynomial in the input size, and therefore all the iterations are doable in PSPACE. Note that due to iterative increase of S' , the first (resp. last) component not bounded (resp. not reverse bounded) by $P(N') \cdot (S' + R)$ in some iteration may become bounded (resp. reverse bounded) in subsequent iterations.

Finally, the procedure arrives at a sequential G -VASS (\mathcal{V}, s, t) , $s = q(\mathbf{v})$, $t = p(\mathbf{w})$, whose first component is *not bounded* by $P(N) \cdot (S + R)$, and whose last component is *not reverse bounded* by $P(N) \cdot (S + R)$. As G is fair, there is a configuration c reachable from s in the first component, and a configuration c' reverse reachable from t in the last component (i.e., t is reachable from c'), both with vectors $\geq \overline{S + R + 1}$:

$$q(\mathbf{v}) \longrightarrow c \quad c' \longrightarrow p(\mathbf{w}).$$

As R is large enough so that every two states in the same component are connected by a path that decreases each counter by at most R , we have some runs

$$c \longrightarrow q(\bar{\mathbf{v}}) \quad p(\bar{\mathbf{w}}) \longrightarrow c'$$

in the first and the last component, respectively, with $\bar{\mathbf{v}}, \bar{\mathbf{w}} \geq \overline{S + 1}$. As $\|s\|, \|t\| \leq S$, this means that \mathcal{V}' is pumpable. Since G is transitive, Lemma 4 applies. In the last step, the decision procedure checks if there is a \mathbb{Z} -run from s to t , which is doable in NP due to Lemma 1. \blacktriangleleft

5 Symmetric group

In this section we prove PSPACE-completeness of the reachability problem in case of the symmetric groups S_d , $d \geq 2$. Our PSPACE upper bound works not only for arbitrary fixed dimension $d \geq 2$, but even in the uniform setting when d is part of input. When $d = 1$, the problem is known to be NP-complete [10].

The upper bound is shown using Lemma 5, which requires proving that S_d satisfies its assumptions. The symmetric groups are obviously transitive, it is thus sufficient to show:

► **Lemma 6.** S_d is fair, for every $d \in \mathbb{N}_+$.


Proof. Consider an S_d -VASS $\mathcal{V} = (Q, T)$ and a configuration $s = q(\mathbf{v})$. Let $S := \|\mathcal{V}, s\|$, let $R \in \mathbb{N}$ be any constant, and $B := 3S + 2R$. We demonstrate that S_d is fair, and that this property is witnessed by the constant polynomial $P(x) = 3d$. Towards this we prove a slightly stronger fact: we assume that \mathcal{V} has a run $\pi: q(\mathbf{v}) \longrightarrow p(\mathbf{w})$ with $\|\mathbf{w}\| > B \cdot d$, and construct a run $\pi': q(\mathbf{v}) \longrightarrow p(\mathbf{w}')$ whose target vector satisfies $\mathbf{w}' \geq \overline{S + R + 1}$.

W.l.o.g. assume that $\mathbf{w}(d) > B \cdot d$. For $i = 1, \dots, d - 1$ we inductively construct runs $\pi_i: q(\mathbf{v}) \rightarrow p(\mathbf{w}_i)$ satisfying

$$\mathbf{w}_i(d) > B \cdot (d - i) \quad \mathbf{w}_i(j) > S + R \quad (\text{for all } j \leq i). \quad (6)$$

We start with $\pi_0 := \pi$. Assuming π_{i-1} has already been constructed, we construct π_i as follows. If $\mathbf{w}_{i-1}(i) > S + R$ then we take $\pi_i := \pi_{i-1}$. Otherwise, we split π_{i-1} into $\pi_{i-1} = \alpha; \beta$,

$$q(\mathbf{v}) \xrightarrow{\alpha} r(\mathbf{u}) \quad r(\mathbf{u}) \xrightarrow{\beta} p(\mathbf{w}_{i-1}),$$



where β is the longest suffix of the run in which the value of coordinate d does not drop below $B \cdot (d - i)$. Therefore, $\mathbf{u}(d) < B \cdot (d - i) + S$, as otherwise β would be longer. Let $\Delta = \mathbf{w}_{i-1} - \mathbf{u}$ be the difference between the target and source of β . We observe that the value of coordinate d is increased by β by at least

$$\Delta(d) > B \cdot (d - i + 1) - B \cdot (d - i) - S = 2S + 2R,$$

while the value of coordinate i is increased by at most $\Delta(i) \leq \mathbf{w}_{i-1}(i) \leq S + R$. Therefore, $\Delta(d) - \Delta(i) > S + R$. In consequence, β contains some k steps

$$q_1(\mathbf{u}_1) \xrightarrow{t_1} q'_1(\mathbf{u}'_1) \quad \dots \quad q_k(\mathbf{u}_k) \xrightarrow{t_k} q'_k(\mathbf{u}'_k)$$

induced by transitions $t_1 = (q_1, \mathbf{e}_1, q'_1), \dots, t_k = (q_k, \mathbf{e}_k, q'_k)$ with effects

$$\mathbf{e}_1 = \mathbf{u}'_1 - \mathbf{u}_1 \quad \dots \quad \mathbf{e}_k = \mathbf{u}'_k - \mathbf{u}_k,$$

such that $\mathbf{e}_\ell(d) > \mathbf{e}_\ell(i)$ for $\ell \in [k]$, that is, each of the steps increases the difference between the values of coordinate d and i , and moreover

$$\mathbf{e}_1(d) + \dots + \mathbf{e}_k(d) > \mathbf{e}_1(i) + \dots + \mathbf{e}_k(i) + S + R, \quad (7)$$

that is, the steps jointly increase the difference between the values of coordinate d and i by more than $S + R$. We may choose a *minimal* set of transitions t_1, \dots, t_k satisfying (7). Hence, knowing that one step can change the difference between the values of coordinate d and i by at most by $2S$, we may safely assume

$$\mathbf{e}_1(d) + \dots + \mathbf{e}_k(d) \leq \mathbf{e}_1(i) + \dots + \mathbf{e}_k(i) + 3S + R, \quad (8)$$

that is, the steps jointly increase the difference between the values of coordinate d and i by at most $3S + R$. Let $\sigma = \begin{pmatrix} i & d \\ d & i \end{pmatrix} \in S_n$ be the permutation that swaps coordinates d and i and preserves the others. We define a run $\pi_i: q(\mathbf{v}) \rightarrow p(\mathbf{w}_i)$ by replacing each transition t_ℓ in π_{i-1} with the transition $\sigma(t_\ell)$. Due to the inequality (8), the coordinate d never drops below

$$\mathbf{u}(d) - (3S + R) > B \cdot (d - i) - (3S + R) > 0$$

in π_i , and $\mathbf{w}_i(d) > B \cdot (d - i + 1) - (3S + R) > B \cdot (d - i)$. Furthermore, due to the inequality (7) we get $\mathbf{w}_i(i) > S + R$, because the coordinate i was increased by more than $S + R$. Other coordinates are not affected, i.e., their values are the same in π_{i-1} and π_i . Therefore, the run π_i satisfies the condition (6), as required.

Finally, we arrive at a run $\pi' = \pi_{d-1}: q(\mathbf{v}) \rightarrow p(\mathbf{w}_{d-1})$ with $\mathbf{w}_{d-1}(i) \geq S + R + 1$ for every $i \in [d]$, and this completes the proof. ◀

▶ **Theorem 7.** S_d -VASS-REACH is PSPACE-complete, for every $d \geq 2$.

Proof. Membership in PSPACE follows immediately by Lemmas 5 and 6, and PSPACE-hardness follows by Lemma 2. ◀

6

 Combining groups

In this section we investigate groups arising as combinations of smaller permutation groups. The way of combining two permutation groups $G \leq S_g$ and $H \leq S_h$ will be their *wreath product* $G \wr H \leq S_{gh}$. The set $[gh] = \{1, \dots, gh\}$ we conveniently split into h blocks:

$$\{1, \dots, g\} \quad \{g+1, \dots, 2g\} \quad \dots \quad \{g(h-1)+1, \dots, gh\},$$

and write $\langle i, j \rangle$ to denote $gj + i \in [gh]$, for $i \in [g]$ and $j \in [h]$. Given a h -tuple $(\sigma_1, \dots, \sigma_h)$ of permutations from G and one permutation δ from H , we define the permutation $(\sigma_1, \dots, \sigma_h) \wr \delta \in S_{g \cdot h}$ that applies, for each j , the permutation σ_j on the block $\{g(j-1)+1, \dots, gj\}$ independently, plus permutes the blocks according to δ , namely

$$(\sigma_1, \dots, \sigma_h) \wr \delta : \langle i, j \rangle \longmapsto \langle \sigma_j(i), \delta(j) \rangle.$$

The *wreath product*⁴ of G and H consists of all so constructible permutations of $[gh]$:

$$G \wr H := \{(\sigma_1, \dots, \sigma_h) \wr \delta \mid \sigma_1, \dots, \sigma_h \in G, \delta \in H\}.$$

In the sequel we are mostly interested in combining trivial permutation groups I_n with symmetric groups S_d , and relating the complexity of the reachability problem for the combined group to the complexity of the problem for component groups. As we show in Theorems 8 and 10 below, it turns out that the complexity significantly depends on the order of composing the groups. In the case of the group $G = I_n \wr S_d$ which allows for arbitrary permutations of d blocks (of size n each) but disallows permutations inside blocks, the complexity of G -VASS-REACH is at least as high as the complexity of $((n-1)d)$ -VASS-REACH (Theorem 8). On the other hand, in the dual case of the group $G = S_d \wr I_n$, which disallows permutations of n blocks (of size d each) but allows arbitrary simultaneous independent permutations inside blocks, we prove (in Theorem 10 below) that the complexity of G -VASS-REACH drops to the complexity of n -VASS-REACH. Summing up, in the former case the impact of symmetric group S_d is essentially nullified as both degrees n and d contribute significantly to the complexity of the combined group, while in the latter case the impact of symmetric group S_d is preserved, as the degree d of the symmetric group is irrelevant for complexity. We notice that the difference of complexities is in agreement with the fact that up to the isomorphism $\langle i, j \rangle \mapsto \langle j, i \rangle$ of permutation groups, $I_n \wr S_d$ is a subgroup of $S_d \wr I_n$.

Trivial symmetry inside blocks. The first result we formulate slightly more generally, with an arbitrary group $G \leq S_d$ in place of the symmetric group S_d . The reachability problem for wreath product of the trivial permutation group I_n with any $G \leq S_d$, regardless of the choice of G (which seems to be quite surprising), is at least as hard as the reachability problem for VASS of dimension $(n-1)d$:

► **Theorem 8.** $((n-1)d)$ -VASS-REACH reduces in polynomial time to $(I_n \wr G)$ -VASS-REACH, for every $n \geq 2$, $d \geq 1$ and $G \leq S_d$.

Proof. Given a VASS $\mathcal{V} = (Q, T)$ of dimension $(n-1)d$, we define a $(I_n \wr G)$ -VASS $\mathcal{V}' = (Q', T')$ simulating \mathcal{V} , whose states extend Q by a number of auxiliary states:

$$Q' = Q \cup \{\langle q, i, t \rangle \mid q \in Q, i \in [d-1], t \in T\}.$$

⁴ The operation is definable more concisely, as semidirect product of h -fold direct product of G with H .

A configuration $q(\mathbf{w}) \in Q \times \mathbb{N}^{(n-1)d}$ of \mathcal{V} , will be simulated by a configuration $\overline{q(\mathbf{w})} = q(\overline{\mathbf{w}}) \in Q \times \mathbb{N}^{nd}$ of \mathcal{V}' whose vector $\overline{\mathbf{w}}$ extends \mathbf{w} , on the last n th coordinate of every block, by

$$\overline{\mathbf{w}}(n, j) = j - 1 \quad (\text{for } j \in [d]) \quad (9)$$

(we write $\overline{\mathbf{w}}(n, j)$ instead of $\overline{\mathbf{w}}(\langle n, j \rangle)$, for readability). We define transitions T' of \mathcal{V}' by its orbit representatives $\tilde{T} \subseteq T'$. For every vector $\mathbf{v} \in \mathbb{Z}^{(n-1)d}$ we define d vectors $\mathbf{v}_1, \dots, \mathbf{v}_d \in \mathbb{Z}^{nd}$ as follows:

$$\mathbf{v}_k(i, j) = \begin{cases} -d + 1 & i = n, j = k \\ 1 & i = n, j \neq k \\ \mathbf{v}(i, j) & i < n, j = k \\ 0 & i < n, j \neq k. \end{cases}$$

Thus, the value $\mathbf{v}_k(n, k)$ on the last coordinate of k th block equals $-d + 1$, while on the last coordinate of other blocks equals 1. Furthermore, \mathbf{v}_k coincides with \mathbf{v} on the remaining $n - 1$ coordinates of k th block, while it is 0 on the remaining $n - 1$ coordinates of other blocks. For every transition $t = (p, \mathbf{v}, q) \in T$, we put to the set \tilde{T} the following d transitions:

$$\begin{aligned} t_d &= (p, \mathbf{v}_d, \langle q, d - 1, t \rangle) & t_{d-1} &= (\langle q, d - 1, t \rangle, \mathbf{v}_{d-1}, \langle q, d - 2, t \rangle) \\ &\dots & & \\ t_2 &= (\langle q, 2, t \rangle, \mathbf{v}_2, \langle q, 1, t \rangle) & t_1 &= (\langle q, 1, t \rangle, \mathbf{v}_1, q). \end{aligned}$$

We need to prove that there is a run $p(\mathbf{w}) \rightarrow q(\mathbf{v})$ in \mathcal{V} if and only if there is a run $\overline{p(\mathbf{w})} \rightarrow \overline{q(\mathbf{v})}$ in \mathcal{V}' . The “only if” direction follows due to observation that every step $p(\mathbf{w}) \xrightarrow{t} q(\mathbf{v})$ in \mathcal{V} is simulated by a sequence of transitions $\alpha_t = t_d, t_{d-1}, \dots, t_1$, namely

$$\overline{p(\mathbf{w})} \xrightarrow{\alpha_t} \overline{q(\mathbf{v})} \quad (10)$$

in \mathcal{V}' , as every transition t_k updates coordinates $1 \dots n - 1$ of block k , and the effect of α_t on last n th coordinate of every block is 0.

For the “if” direction, we observe that due to the auxiliary states in Q' , every run $\overline{p(\mathbf{w})} \rightarrow \overline{q(\mathbf{v})}$ in \mathcal{V}' , starting and ending in a state from Q , necessarily splits into *segments* of length d , where each segment is of the form

$$\sigma_d(t_d), \sigma_{d-1}(t_{d-1}), \dots, \sigma_1(t_1), \quad (11)$$

for some $\sigma_1, \dots, \sigma_d \in I_n \wr S_d$ and $t = (p, \mathbf{v}, q) \in T$. In particular, the segment starts in $p \in Q$ and ends in $q \in Q$. It is thus enough to argue that whenever $p(\overline{\mathbf{w}}) \rightarrow q(\mathbf{v}')$ in \mathcal{V}' along a single segment (11), there is a vector \mathbf{v} such that $\mathbf{v}' = \overline{\mathbf{v}}$ and $p(\mathbf{w}) \rightarrow q(\mathbf{v})$ in \mathcal{V} (*).

A single segment updates coordinates $1 \dots n - 1$ in all the d blocks, in the order determined by $\sigma_1, \dots, \sigma_d$. The crucial observation is that this order is necessarily $d, d - 1, \dots, 1$. Indeed, by the definition of the vector $\overline{\mathbf{w}}$ in (9), the first update is in the d th block as the last n th coordinate is $\geq d - 1$ only in this block. As a result of this update, the $(d - 1)$ th block is the only one with value $\geq d - 1$ on the last n -th coordinate, and therefore the second update necessarily happens in this block. And so on, until the last update in the first block. In consequence, $\sigma_i(t_i) = t_i$ for all $i \in [d]$. Therefore, the segment (11) is necessarily of the form (10), and the implication (*) is shown. In particular, we have shown that only transitions from \tilde{T} can be used in \mathcal{V}' when starting from a configuration $\overline{p(\mathbf{w})}$. ◀

Trivial symmetry between blocks. Now we investigate the other case of wreath product of the symmetric group of degree d with the trivial permutation group of degree n , and show the complexity of the reachability problem in this case to be essentially the same as of the reachability problem for VASS of dimension n .

We start with a definition and a lemma, both stated for any $G \leq S_n$ in place of I_n . A configuration $q(\mathbf{w})$ of a $(S_d \wr G)$ -VASS \mathcal{V} is *B-balanced* if

$$|\mathbf{w}(i, j) - \mathbf{w}(i', j)| \leq B \quad (12)$$

for every $i, i' \in [d]$ and $j \in [n]$. In words: in every block, the maximal difference of values is at most B . A run $s \rightarrow t$ of \mathcal{V} is *B-balanced* if all its configurations are so. Finally, (\mathcal{V}, s, t) is *B-balanceable* if an existence of a run $s \rightarrow t$ implies an existence of a run $s \rightarrow t$ which is *B-balanced*.

► **Lemma 9.** *There is a polynomial P such that for every $d, n \geq 1$ and $G \leq S_n$, every $(S_d \wr G)$ -VASS (\mathcal{V}, s, t) is $P(\|\mathcal{V}, s, t\|)$ -balanceable.*

Proof. Let (\mathcal{V}, s, t) be an arbitrary $(S_d \wr G)$ -VASS, and let $N := \|\mathcal{V}, s, t\|$, $B_0 := 4N^2$, and $B := B_0 + 2N + 8N^3$. Suppose \mathcal{V} has a run $s \xrightarrow{\pi} t$. We iteratively modify this run until it becomes *B-balanced*.

By definition $\|s\|, \|t\| < B_0 \leq B$, and therefore the source configuration s and the target one t are *B-balanced*. We make π balanced in n stages. For each $j \in [n]$, in an arbitrary order, we modify the run π to make it balanced inside the j th block, i.e., to make all configurations of π satisfy (12) where i, i' range over $[d]$ but j is fixed. As the modification does not affect the other blocks, the final run is thus balanced.

From now on we concentrate on one stage, for a fixed $j \in [n]$. For a vector $\mathbf{u} \in \mathbb{Z}^{dn}$ and $i, i' \in [d]$, we use the notation $\Delta_{ii'}(\mathbf{u}) := \mathbf{u}(i, j) - \mathbf{u}(i', j)$ for the difference between the values of \mathbf{u} on coordinates $\langle i, j \rangle$ and $\langle i', j \rangle$. The stage amounts to iteratively modifying the run until it finally becomes balanced inside the j th block. Each such individual modification we call a *microstage*. Below we describe a single microstage, and argue that it decreases a certain nonnegative rank, which ensures termination of the whole stage.

Microstage. Suppose the run π is not balanced in the j th block, namely π contains a configuration $p(\mathbf{v})$ such that $\Delta_{ii'}(\mathbf{v}) > B$ for some $i, i' \in [d]$. Consider the longest infix

$$q(\mathbf{w}) \xrightarrow{\alpha} p(\mathbf{v}) \xrightarrow{\beta} q'(\mathbf{w}') \quad (13)$$

of π , such that all its configurations $p'(\mathbf{v}')$ satisfy $\Delta_{ii'}(\mathbf{v}') \geq B_0$. Since one step may change the difference between the values of coordinates (i, j) and (i', j) by at most by $2N$, we notice that $\Delta_{ii'}(\mathbf{w}) < B_0 + 2N$, as otherwise α would be longer. Likewise, $\Delta_{ii'}(\mathbf{w}') < B_0 + 2N$. In consequence $B - \Delta_{ii'}(\mathbf{w}) \geq 8N^3$ and $B - \Delta_{ii'}(\mathbf{w}') \geq 8N^3$, and therefore α contains at least $4N^2$ steps that increase the difference between the values of coordinates $\langle i, j \rangle$ and $\langle i', j \rangle$ and similarly β contains at least $4N^2$ steps that decrease this difference. As there are at most $2N$ different values by which a single step can increase (or decrease) this difference, by the pigeonhole principle α contains some $2N$ steps that increase this difference by the same value a , and likewise β contains some $2N$ steps that decrease this difference by the same value b . We pick any b among the $2N$ steps in α , executing transitions t_1, \dots, t_b , say, and replace them by $\sigma(t_1), \dots, \sigma(t_b)$, where $\sigma \in S_d \wr G$ is a permutation that swaps (i, j) and (i', j) . Likewise, we choose any a among the $2N$ steps in β , executing transitions t'_1, \dots, t'_a , say, and replace them by $\sigma(t'_1), \dots, \sigma(t'_a)$. Importantly, this modification of the run does

not change its target configuration. The only coordinate whose value decreases in some configurations along the run is (i, j) , but it never drops below zero since the decrease is never larger than $4N^2$, and $B_0 \geq 4N^2$. Therefore, the modification produces a run.

We claim that the following nonnegative rank decreases as a result of the microstage (where the first sum ranges over configurations $r(\mathbf{u})$ along the run):

$$\sum_{r(\mathbf{u})} \sum_{\ell < \ell' \in [d]} |\Delta_{\ell\ell'}(\mathbf{u})|,$$

on the basis of the following three observations. First, the value $|\Delta_{ii'}(\mathbf{u})|$ never increases, and decreases for some \mathbf{u} , namely for $p(\mathbf{v})$ in (13), the value $|\Delta_{ii'}(\mathbf{v})| = |\Delta_{i'i}(\mathbf{v})|$ decreases by $ab > 0$ (but $\Delta_{ii'}(\mathbf{v})$ remains positive). Second, in all cases when the value $|\Delta_{ii'}(\mathbf{u})|$ changes (i.e., decreases) in the microstage, this is due to adding opposite values to $\mathbf{u}(i, j)$ and $\mathbf{u}(i', j)$. Therefore, for every $\ell \notin \{i, i'\}$ and every configuration $r(\mathbf{u})$ in π , the value of $|\Delta_{i\ell}(\mathbf{u})| + |\Delta_{i'\ell}(\mathbf{u})|$ does not increase. Finally, for $\ell, \ell' \notin \{i, i'\}$, the value of $\Delta_{\ell\ell'}(\mathbf{u})$ is preserved. In consequence, the rank of the whole run decreases, as required. ◀

► **Theorem 10.** $(S_d \wr I_n)$ -VASS-REACH reduces in exponential time to n -VASS-REACH, for every $d, n \geq 1$.

Proof. Let $\mathcal{V} = (Q, T)$ be $(S_d \wr I_n)$ -VASS and s, t its configurations. Let $B := P(N)$, where P is the polynomial of Lemma 9 and $N = ||\mathcal{V}, s, t||$. Let $\mathcal{B} := [B] \cup \{0\}$. Relying on Lemma 9 we construct an n -VASS $\mathcal{V}' = (Q', T')$, where $Q' = Q \times \mathcal{B}^{nd}$, that simulates B -balanced runs of \mathcal{V} . Specifically, every j th block of d coordinates of \mathcal{V} is simulated by a single coordinate j of \mathcal{V}' whose value is invariantly equal to the largest value in the j th block. The simulation is possible due to the auxiliary component \mathcal{B}^{nd} of states that keeps the difference between the value of every coordinate (i, j) of \mathcal{V} , and the largest value inside the j th block (to which (i, j) belongs). The transitions T' of \mathcal{V}' simulate transitions of T , and accordingly update the auxiliary component of states. We claim that \mathcal{V} has a B -balanced run $s \rightarrow t$ if and only if \mathcal{V}' has a run $s' \rightarrow t'$, where s', t' are obtained from s, t by computing the auxiliary component of states, and replacing the values of each block of coordinates by the largest value in this block. The reduction is thus completed. ◀

7 Other groups

We complete Section 5 by showing that alternating groups admit similar complexity drop as symmetric groups. We also demonstrate that sole transitivity is not enough to achieve significant drop, namely the complexity of the reachability problem in case of cyclic groups is comparable to general VASS.

Alternating group. Let $A_d \leq S_d$ denote the alternating group of degree d , i.e., the group of all even permutations of $[d]$. The group is generated by 3 -cycles, permutations of the form $\sigma_{ijk} = \begin{pmatrix} i & j & k \\ j & k & i \end{pmatrix}$, for pairwise distinct $i, j, k \in [d]$. When $d = 2$ the alternating group is trivial, A_2 -VASS = 2-VASS, and the reachability problem is PSPACE-complete [1]. In the sequel we thus consider degrees $d \geq 3$.

► **Lemma 11.** A_d is fair, for $d \geq 3$.

The argument elaborates on the proof of Lemma 6, and proceeds in two steps. First, assuming a run π with one sufficiently large coordinate at the end, by applying a permutation from A_d to a properly chosen suffix of π we construct a run with *two* large coordinates at the end. Second, we further modify the run to increase all coordinates above a required threshold, similarly to the proof of Lemma 6.

► **Theorem 12.** A_d -VASS-REACH is PSPACE-complete, for every $d \geq 3$.

Proof. Membership in PSPACE follows immediately by Lemmas 5 and 11, while PSPACE-hardness follows by Lemma 2. ◀

Cyclic groups. Let Z_n denote the cyclic group of degree n , i.e., the group $Z_n = \{\sigma^i \mid i \in \{0, 1, \dots, n-1\}\}$ generated by the single cyclic shift $\sigma = \begin{pmatrix} 1 & 2 & \dots & n \\ 2 & 3 & \dots & 1 \end{pmatrix}$.

► **Theorem 13.** n -VASS-REACH reduces in polynomial time to (Z_{2n+8}) -VASS-REACH, for every $n \geq 1$.

In the proof, an n -VASS is simulated using n consecutive counters of a Z_{2n+8} -VASS. The remaining $n+8$ counters are used to enforce that only one transition from each Z_{2n+8} -orbit can be applied, which assures faithfulness of the simulation.

8 Final remarks

In this paper we have investigated the impact of symmetry in VASS on the complexity of the reachability problem. The permutation groups considered in this paper split clearly into two groups. On one side there are “easy” groups, which includes both symmetric and alternating groups featuring PSPACE-completeness, but also the combinations $S_d \wr I_n$ of symmetric groups and trivial ones, where the complexity is independent of the degree d of the symmetric group (dependence on the degree n of the trivial group is unavoidable). On the other side there are “hard” groups, which includes both trivial and cyclic groups, which not differ significantly in complexity, but also the dual combinations $I_n \wr S_d$ of symmetric groups and trivial ones, where the complexity is dependent on both d and n . We end up with a research question left for further research: can one classify *all* finite permutation groups into “easy” and “hard” ones?

The model of data VASS is definable in our setting as $(I_n \wr S_\omega)$ -VASS. Therefore, “hardness” of $I_n \wr S_d$ may suggest that the reachability problem for data VASS is harder than for plain VASS. On the other hand, “easiness” of $S_d \wr I_n$ suggests that the reachability problem for the subclass of data VASS invariant under independent data permutations in each dimension (definable as $(S_\omega \wr I_n)$ -VASS in our setting) is not significantly harder than for plain VASS (and, in particular, decidable). Our initial results seem to confirm both these suggestions.

References

- 1 Michael Blondin, Alain Finkel, Stefan Göller, Christoph Haase, and Pierre McKenzie. Reachability in Two-Dimensional Vector Addition Systems with States Is PSPACE-Complete. In *Proceedings of LICS 2015*, pages 32–43, 2015. doi:10.1109/LICS.2015.14.
- 2 Wojciech Czerwinski, Ismaël Jecker, Slawomir Lasota, Jérôme Leroux, and Lukasz Orlikowski. New lower bounds for reachability in vector addition systems. In *Proc. FSTTCS 2023*, volume 284 of *LIPICs*, pages 35:1–35:22. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPICs.FSTTCS.2023.35.

- 3 Wojciech Czerwinski, Sławomir Lasota, Ranko Lazic, Jérôme Leroux, and Filip Mazowiecki. The reachability problem for Petri nets is not elementary. In *Proceedings of STOC 2019*, pages 24–33. ACM, 2019. doi:10.1145/3313276.3316369.
- 4 Wojciech Czerwinski, Sławomir Lasota, Ranko Lazic, Jérôme Leroux, and Filip Mazowiecki. The Reachability Problem for Petri Nets Is Not Elementary. *Journal of the ACM*, 68(1):7:1–7:28, 2021. doi:10.1145/3422822.
- 5 Wojciech Czerwinski and Łukasz Orlikowski. Reachability in Vector Addition Systems is Ackermann-complete. In *Proceedings of FOCS 2021*, pages 1229–1240. IEEE, 2021. doi:10.1109/FOCS52979.2021.00120.
- 6 Wojciech Czerwinski and Łukasz Orlikowski. Lower bounds for the reachability problem in fixed dimensional vasses. In *Proceedings of LICS 2022*, pages 40:1–40:12. ACM, 2022. doi:10.1145/3531130.3533357.
- 7 Friedrich Eisenbrand and Gennady Shmonin. Carathéodory bounds for integer cones. *Operations Research Letters*, 34(5):564–568, 2006. doi:10.1016/J.ORL.2005.09.008.
- 8 John Fearnley and Marcin Jurdziński. Reachability in two-clock timed automata is PSPACE-complete. *Information and Computation*, 243:26–36, 2015. doi:10.1016/J.IC.2014.12.004.
- 9 Yuxi Fu, Qizhe Yang, and Yangluo Zheng. Improved algorithm for reachability in d-vass. In *Proc. ICALP 2024*, volume 297 of *LIPICs*, pages 136:1–136:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICs.ICALP.2024.136.
- 10 Christoph Haase, Stephan Kreutzer, Joël Ouaknine, and James Worrell. Reachability in succinct and parametric one-counter automata. In *Proceeding of CONCUR 2009*, volume 5710, pages 369–383. Springer, 2009. doi:10.1007/978-3-642-04081-8_25.
- 11 S. Rao Kosaraju. Decidability of reachability in vector addition systems (preliminary version). In *Proc. STOC 1982*, pages 267–281, 1982. doi:10.1145/800070.802201.
- 12 Jean-Luc Lambert. A structure to decide reachability in Petri nets. *Theor. Comput. Sci.*, 99(1):79–104, 1992. doi:10.1016/0304-3975(92)90173-D.
- 13 Sławomir Lasota. Decidability border for Petri nets with data: WQO dichotomy conjecture. In *Proc. Petri Nets 2016*, volume 9698 of *Lecture Notes in Computer Science*, pages 20–36. Springer, 2016. doi:10.1007/978-3-319-39086-4_3.
- 14 Sławomir Lasota. Improved Ackermannian Lower Bound for the Petri Nets Reachability Problem. In *Proceedings of STACS 2022*, volume 219 of *LIPICs*, pages 46:1–46:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICs.STACS.2022.46.
- 15 Ranko Lazic, Thomas Christopher Newcomb, Joël Ouaknine, A. W. Roscoe, and James Worrell. Nets with tokens which carry data. In *Proc. ICATPN 2007*, volume 4546 of *Lecture Notes in Computer Science*, pages 301–320. Springer, 2007. doi:10.1007/978-3-540-73094-1_19.
- 16 Jérôme Leroux. The Reachability Problem for Petri Nets is Not Primitive Recursive. In *Proceedings of FOCS 2021*, pages 1241–1252. IEEE, 2021. doi:10.1109/FOCS52979.2021.00121.
- 17 Jérôme Leroux and Sylvain Schmitz. Reachability in Vector Addition Systems is Primitive-Recursive in Fixed Dimension. In *Proceedings of LICS 2019*, pages 1–13. IEEE, 2019. doi:10.1109/LICS.2019.8785796.
- 18 Richard J. Lipton. The Reachability Problem Requires Exponential Space. Technical report, Yale University, 1976.
- 19 Ernst W. Mayr. An algorithm for the general Petri net reachability problem. In *Proc. STOC 1981*, pages 238–246, 1981. doi:10.1145/800076.802477.
- 20 Fernando Rosa-Velardo and David de Frutos-Escrig. Decidability and complexity of Petri nets with unordered data. *Theor. Comput. Sci.*, 412(34):4439–4451, 2011. doi:10.1016/J.TCS.2011.05.007.