ON THE LENGTH OF STRONGLY MONOTONE DESCENDING CHAINS OVER $\mathbb{N}^d$

SYLVAIN SCHMITZ[1] AND LIA SCHÜTZE[2]

Abstract. A recent breakthrough by Künnemann, Mazowiecki, Schütze, Sinclair-Banks, and Węgrzycki (ICALP 2023) bounds the running time for the coverability problem in d-dimensional vector addition systems under unary encoding to $n^{2^{O(d)}}$, improving on Rackoff's $n^{2^{O(d \lg d)}}$ upper bound (*Theor. Comput. Sci.* 1978), and provides conditional matching lower bounds.

In this paper, we revisit Lazić and Schmitz' "ideal view" of the backward coverability algorithm (*Inform. Comput.* 2021) in the light of this breakthrough. We show that the controlled strongly monotone descending chains of downwards-closed sets over $\mathbb{N}^d$ that arise from the dual backward coverability algorithm of Lazić and Schmitz on d-dimensional unary vector addition systems also enjoy this tight $n^{2^{O(d)}}$ upper bound on their length, and that this also translates into the same bound on the running time of the backward coverability algorithm.

Furthermore, our analysis takes place in a more general setting than that of Lazić and Schmitz, which allows to show the same results and improve on the 2EXPSPACE upper bound derived by Benedikt, Duff, Sharad, and Worrell (LICS 2017) for the coverability problem in invertible affine nets.

Keywords. Vector addition system, coverability, well-quasi-order, order ideal, affine net

1. Introduction

Well-Quasi-Orders (wqo for short) are a notion from order theory [29, 41] that has proven very effective in many areas of mathematics, logic, combinatorics, and computer science in order to establish finiteness statements. For instance, in the field of formal verification, they provide the termination arguments for the generic algorithms for *well structured transition systems* [1, 22], notably the *backward coverability algorithm* for deciding safety properties [3, 1, 22].

In full generality, one cannot extract complexity bounds from wqo-powered termination proofs. Nevertheless, in an algorithmic setting, one can "instrument" wqos by considering so-called *controlled sequences* [41, 39], and new tight complexity upper bounds for wqo-based algorithms now appear on a regular basis [e.g., 40, 4, 6, 5, 26, for a few recent examples].

Those complexity upper bounds are however astronomically high, and sometimes actually way too high for the problem at hand. An emblematic illustration of this phenomenon is the backward coverability algorithm for vector addition systems (VAS), which was shown to run in double exponential time by Bozzelli and Ganty [13] based on an original analysis due to Rackoff [37]: the corresponding bounds over the wqo $\mathbb{N}^d$ are Ackermannian [20].

[1] Université Paris Cité, CNRS, IRIF, Paris, France
[2] Max Planck Institute for Software Systems (MPI-SWS), Kaiserslautern, Germany

Descending Chains. One way pioneered by Lazić and Schmitz [32] to close such complexity gaps while retaining some of the wide applicability of wqos and well structured transition systems is to focus on the descending chains of downwards closed sets over the wqo at hand. Indeed, one of the equivalent characterisations of wqos is the *descending chain condition* [29, 41], which guarantees that those descending chains are finite.

In themselves, descending chains are no silver bullet: e.g., the controlled descending chains over $\mathbb{N}^d$ are also of Ackermannian length [32, Thm. 3.10]. Nevertheless, these chains sometimes exhibit a form of "monotonicity," which yields vastly improved upper bounds. When applied to a *dual* version of the backward coverability algorithm in well structured transition systems, this allows to recover the same double exponential time upper bound as in [13, 37] for the VAS coverability problem, along with tight upper bounds for coverability in several VAS extensions. The same framework was also the key to establishing tight bounds for coverability in ν-Petri nets [31]. As a further testimony to the versatility of the approach, Benedikt, Duff, Sharad, and Worrell use it in [7] to derive original upper bounds for problems on invertible polynomial automata and invertible affine nets, in a setting that is not strictly speaking one of well structured transition systems.

Fine-grained Bounds for VAS Coverability. The coverability problem in VAS is well-known to be EXPSPACE-complete, thanks to Rackoff's 1978 upper bound matching a 1976 lower bound by Lipton [34]. The main parameter driving this complexity is the dimension of the system: the problem is in pseudo-polynomial time in fixed dimension d; more precisely, Rackoff's analysis yields a $n^{2^{O(d \lg d)}}$ deterministic time upper bound for d-dimensional VAS encoded in unary [38], by proving the same bound on the length of a covering execution of minimal length. Here, there is a discrepancy with the $n^{2^{\Omega(d)}}$ lower bound on the length of that execution in Lipton's construction—a discrepancy that was already highlighted as an open problem in the early 1980's by Mayr and Meyer [35], and settled in the specific case of reversible systems by Koppenhagen and Mayr [28]. The upper bounds of both Bozzelli and Ganty [13] and Lazić and Schmitz [32] on the complexity of the backward coverability algorithm inherit from Rackoff's $n^{2^{O(d \lg d)}}$ bound and suffer from the same discrepancy.

This was the situation until Künnemann, Mazowiecki, Schütze, Sinclair-Banks, and Węgrzycki showed an $n^{2^{O(d)}}$ upper bound on the length of minimal covering executions of unary encoded d-dimensional VAS, matching Lipton's lower bound [30, Thm. 3.3]. This directly translates into a deterministic algorithm with the same upper bound on the running time [30, Cor. 3.4]. Furthermore, assuming the exponential time hypothesis, Künnemann et al. also show that there does not exist a deterministic $n^{o(2^d)}$ time algorithm deciding coverability in unary encoded d-dimensional VAS [30, Thm. 4.2].

Thinness. The improved upper bound relies on the notion of a *thin* vector in $\mathbb{N}^d$ [30, Def. 3.6] (somewhat reminiscent of the "extractors" of Leroux [33]). The proof of [30, Thm. 3.3] works by induction on the dimension d. By splitting a covering execution of minimal length at the first non-thin configuration, Künnemann et al. obtain a prefix made of distinct thin configurations (which must then be of bounded length), and a suffix starting from a configuration with some components high enough to be disregarded, hence that can be treated as an execution in a VAS of lower dimension, on which the induction hypothesis applies.

Contributions. In this paper, we show that the improved $n^{2^{O(d)}}$ upper bound of Künnemann et al. [30] also applies to the number of iterations of the backward coverability algorithm for d-dimensional VAS encoded in unary (see Theorem 4.2). In order to do so, one could reuse the approach of Bozzelli and Ganty [13] to lift the improved bound from the length of minimal covering executions to the running time of the backward coverability algorithm, but here we aim for the generality of the framework of [32].

Our main contribution is thus to show in § 3 that the upper bounds on the length of strongly monotone controlled descending chains of downwards closed sets over $\mathbb{N}^d$—which include those constructed during the running of the backward coverability algorithm for VAS—can be improved similarly (see Theorem 3.6) when focusing on a suitably generalised notion of thinness. As a byproduct, we observe that thinness is an inherent property of such chains (see Corollary 3.7), rather than an *a priori* condition that—almost magically—yields the improved bound.

We apply our results to the coverability problem in vector addition systems in § 4.2—thus providing as promised an alternative to applying Bozzelli and Ganty's approach to Künnemann et al.'s results—and show that the backward coverability algorithm runs in time $n^{2^{O(d)}}$ (see Corollary 4.5) and is therefore conditionally optimal by [30, Thm. 4.2].

As a further demonstration of the versatility of our results, we show in § 4.3 how to apply them to invertible affine nets, a generalisation of vector addition systems introduced by Benedikt et al. [7], and a good showcase for our techniques. We obtain the same bounds for their coverability problem as in the case of vector addition systems (see Theorem 4.11 and Corollary 4.12), and thereby improve on the 2EXPSPACE upper bound of [7] by showing that the problem is actually EXPSPACE-complete (see Corollary 4.13). Along the way, we will see that the improved upper bounds also apply for other VAS extensions, for which Rackoff's proof scheme had been successfully adapted (see remarks 4.4 and 4.15), namely strictly increasing affine nets [12], branching VAS [16], and alternating VAS [15].

2. Well-Quasi-Orders and Ideals

We start by introducing the necessary background on well-quasi-orders, descending chains, and order ideals.

Well-Quasi-Orders. A *quasi-order* $(X, \leq)$ comprises a set X and a transitive reflexive relation $\leq \subseteq X \times X$. For a subset $S \subseteq X$, its *downward closure* is the set of elements smaller or equal to some element in S, i.e., ${\downarrow}S \stackrel{\text{def}}{=} \{x \in X \mid \exists y \in S \,.\, x \leq y\}$. When $S = \{y\}$ is a singleton, we note ${\downarrow}y$ for this set. A subset $S \subseteq X$ is *downwards-closed* if $S = {\downarrow}S$. A *well-quasi-order* is a quasi-order $(X, \leq)$ such that all the *descending chains*

$$D_0 \supsetneq D_1 \supsetneq D_2 \supsetneq \cdots \tag{1}$$

of downwards-closed subsets $D_k \subseteq X$ are finite [29, 41].

Conversely, the *upward closure* of a subset $S \subseteq X$ is ${\uparrow}S \stackrel{\text{def}}{=} \{x \in X \mid \exists y \in S.y \leq x\}$, and S is *upwards-closed* if $S = {\uparrow}S$. The complement $X \setminus D$ of a downwards-closed set D is upwards-closed (and conversely), hence wqos have the *ascending chain condition* for chains $U_0 \subsetneq U_1 \subsetneq \cdots$ of upwards-closed sets: they are necessarily

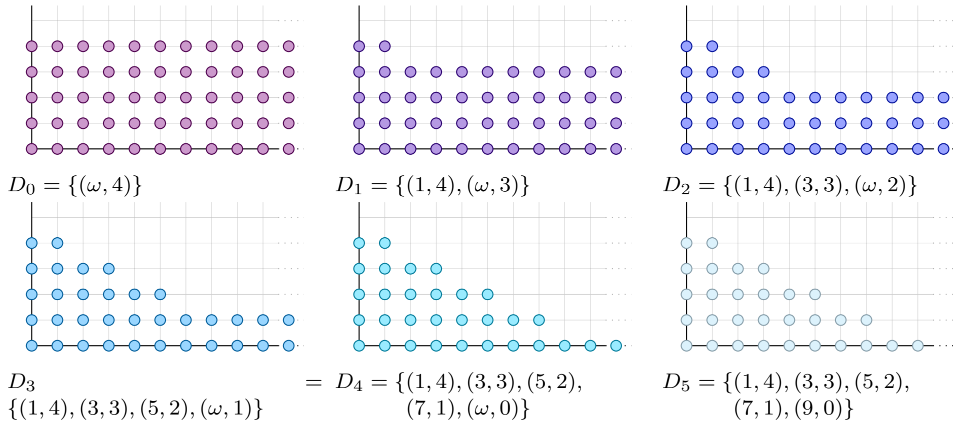

FIGURE 1. A descending chain $D_0 \supsetneq D_1 \supsetneq \cdots \supsetneq D_5$ over $\mathbb{N}^2$ [32, Ex. 3.6].

finite. Furthermore, any upwards-closed set U over a wqo has a *finite basis* B such that $U = {\uparrow}B$ [29, 41]; without loss of generality, we can take the elements of B to be minimal and mutually incomparable in U.

A well-studied wqo is $(\mathbb{N}^d, \sqsubseteq)$ the set of d-dimensional vectors of natural numbers along with the component-wise (aka product) ordering [17]; see Figure 1 for an illustration of a descending chain over $\mathbb{N}^2$, which happens to be produced by the backward coverability algorithm for a vector addition system [32, Ex. 3.6].

Order Ideals. An *order ideal* of X is a downwards-closed subset $I \subseteq X$, which is *directed*: it is non-empty, and if x, x' are two elements of I, then there exists y in I with $x \leq y$ and $x' \leq y$. Alternatively, order ideals are characterised as the *irreducible* non-empty downwards-closed sets of X: an order ideal is a non-empty downwards-closed set I with the property that, if $I \subseteq D_1 \cup D_2$ for two downwards-closed sets D_1 and D_2, then $I \subseteq D_1$ or $I \subseteq D_2$.

Over a wqo $(X, \leq)$, any downwards-closed set $D \subseteq X$ has a canonical decomposition as a finite union of order ideals $D = I_1 \cup \cdots \cup I_n$, where the I_j's are mutually incomparable for inclusion [11, 25]. We write $I \in D$ if I is an order ideal appearing in the canonical decomposition of D, i.e., if it is a maximal order ideal included in D. Then $D \subseteq D'$ if and only if, for all $I \in D$, there exists $I' \in D'$ such that $I \subseteq I'$.

Effective Representations over $\mathbb{N}^d$. Over the wqo $(\mathbb{N}^d, \sqsubseteq)$, the order ideals are exactly the sets of the form ${\downarrow}\boldsymbol{v} \cap \mathbb{N}^d$ where $\boldsymbol{v}$ ranges over $\mathbb{N}^d_\omega \stackrel{\text{def}}{=} (\mathbb{N} \uplus \{\omega\})^d$, where ω is a new top element [25]. From here on, we will abuse notations and identify an order ideal I of $\mathbb{N}^d$ with the vector $\boldsymbol{v}$ in $\mathbb{N}^d_\omega$ such that $I = {\downarrow}\boldsymbol{v} \cap \mathbb{N}^d$. See for instance the decompositions in Figure 1.

Let us introduce some notations for the sets of *infinite* and *finite components* of I, namely

$$\omega(I) \stackrel{\text{def}}{=} \{1 \leq i \leq d \mid I(i) = \omega\}, \qquad \text{fin}(I) \stackrel{\text{def}}{=} \{1 \leq i \leq d \mid I(i) < \omega\}, \tag{2}$$

along with its *dimension* and *finite dimension*, respectively defined as

$$\dim I \stackrel{\text{def}}{=} |\omega(I)|, \qquad \text{fdim}\, I \stackrel{\text{def}}{=} |\text{fin}(I)|. \tag{3}$$

Note that $\mathrm{fin}(I) = \{1,\dots,d\} \setminus \omega(I)$ and $\mathrm{fdim}\, I = d - \dim I$. For instance, the order ideal $I = (\omega, 4)$ in the decomposition of D_0 in Figure 1 satisfies $\omega(I) = \{1\}$ and $\dim I = 1$.

The order ideals of $\mathbb{N}^d$, when represented as vectors in $\mathbb{N}^d_\omega$, are rather easy to manipulate [25] — and thus so are the downwards-closed subsets of $\mathbb{N}^d$ when represented as finite sets of vectors in $\mathbb{N}^d_\omega$. For instance,

- $I \subseteq I'$ (as subsets of $\mathbb{N}^d$) if and only if $I \sqsubseteq I'$ (as vectors in $\mathbb{N}^d_\omega$) — which incidentally entails $\omega(I) \subseteq \omega(I')$ and therefore $\dim I \leq \dim I'$; also note that, if $I \subseteq I'$ and $\dim I = \dim I'$, then $\omega(I) = \omega(I')$;
- the intersection of two order ideals is again an order ideal, represented by the vector $I \wedge I'$ defined by $(I \wedge I')(i) \stackrel{\text{def}}{=} \min(I(i), I'(i))$ for all $1 \leq i \leq d$;
- the complement of an order ideal I is the upwards-closed set $\bigcup_{i \in \mathrm{fin}(I)} \uparrow\big((I(i)+1) \cdot \boldsymbol{e}_i\big)$, where $\boldsymbol{e}_i$ denotes the unit vector with "1" in coordinate i and "0" everywhere else.

Proper Ideals and Monotonicity. If $D \supsetneq D'$, then there must be an order ideal $I \in D$ such that $I \notin D'$. Coming back to a descending chain $D_0 \supsetneq D_1 \supsetneq \cdots \supsetneq D_\ell$, we then say that an order ideal I is *proper* at step k, for $0 \leq k < \ell$, if $I \in D_k$ but $I \notin D_{k+1}$; at each step $0 \leq k < \ell$, there must be at least one proper order ideal. In Figure 1, $(\omega, 4)$ is proper at step 0, and more generally $(\omega, 4-k)$ is the only proper order ideal at step $0 \leq k < 5$.

It turns out that the descending chains arising from some algorithmic procedures, including the backward coverability algorithm for VAS, enjoy additional relationships between their proper order ideals. Over $(\mathbb{N}^d, \sqsubseteq)$, we say that a descending chain $D_0 \supsetneq D_1 \supsetneq \cdots$ is

- *strongly monotone* [36, 7] if, whenever an ideal I_{k+1} is proper at some step $k+1$, then there exists I_k proper at step k such that $\dim I_{k+1} \leq \dim I_k$, and
- in particular *ω-monotone* [32] if, whenever an ideal I_{k+1} is proper at some step $k+1$, then there exists I_k proper at step k such that $\omega(I_{k+1}) \subseteq \omega(I_k)$.

The descending chain depicted in Figure 1 is ω-monotone—and thus strongly monotone— with $\omega((\omega, 4-(k+1))) \subseteq \omega((\omega, 4-k))$ for all $4 > k \geq 0$.

Controlled Sequences. While guaranteed to be finite, descending chains over a wqo can have arbitrary length. Nevertheless, their length can be bounded under additional assumptions. We define the *size* of a downwards-closed subset of $\mathbb{N}^d$ and of an order ideal of $\mathbb{N}^d$ as

$$\|D\| \stackrel{\text{def}}{=} \max_{I \in D} \|I\|\,, \qquad\qquad \|I\| \stackrel{\text{def}}{=} \max_{i \in \mathrm{fin}(I)} I(i)\,. \tag{4}$$

In Figure 1, $\|D_0\| = \|D_1\| = \|D_2\| = 4$, $\|D_3\| = 5$, $\|D_4\| = 7$, and $\|D_5\| = 9$.

Given a *control* function $g\colon \mathbb{N} \to \mathbb{N}$, which will always be monotone (i.e., $\forall x \leq y. g(x) \leq g(y)$) and expansive (i.e., $\forall x. x \leq g(x)$) along with an *initial size* $n_0 \in \mathbb{N}$, we say that a descending chain $D_0 \supsetneq D_1 \supsetneq \cdots$ over $\mathbb{N}^d$ is (asymptotically) (g, n_0)-*controlled* if, for all $k \geq 0$,

$$\|D_k\| \leq g^k(n_0) \tag{5}$$

where $g^k(n_0)$ is the kth iterate of g applied to n_0 [39]. In particular, $\|D_0\| \leq n_0$ initially. In Figure 1, the descending chain is $(g, 4)$-controlled for $g(x) \stackrel{\text{def}}{=} x + 1$.

3. Main Result

In this section, we establish a new bound on the length of controlled strongly monotone descending sequences. This relies on a generalisation of the notion of *thinness* from Künnemann et al. [30, Def. 3.6] (see § 3.1), before we can apply thinness in the setting of strongly monotone descending chains and prove our main result in § 3.2.

3.1. **Thinness.**

Fix a control function g, an initial size n_0, and a dimension $d \geq 0$. Define inductively the bounds on sizes $(N_i)_{0\leq i\leq d}$ and lengths $(L_i)_{0\leq i\leq d}$ as follows

$$N_0 \stackrel{\text{def}}{=} n_0 \;, \qquad N_{i+1} \stackrel{\text{def}}{=} g^{L_i+1}(n_0) \;, \tag{6}$$

$$L_0 \stackrel{\text{def}}{=} 0 \;, \qquad L_{i+1} \stackrel{\text{def}}{=} L_i + \prod_{1\leq j\leq i+1} (d-j+1)(N_j+1) \;. \tag{7}$$

Beware the abuse of notation, as the bounds above depend on (g, n_0) and d, but those will always be clear from the context.

Remark 3.1 (Monotonicity of $(N_i)_{0\leq i\leq d}$ and $(L_i)_{0\leq i\leq d}$)**.** By definition, for all $0 \leq i < j \leq d$, $0 \leq L_i < L_j$, and because g is assumed monotone expansive, $n_0 \leq N_i \leq N_j$. □

The following definition generalises [30, Def. 3.6] to handle order ideals and an arbitrary control function and initial size.

Definition 3.2 (Thin order ideal)**.** Let (g, n_0) be a control function and initial size and $d > 0$ a dimension. An order ideal I of $\mathbb{N}^d$ is *thin* if there exists a bijection $\sigma\colon \text{fin}(I) \to \{1, \ldots, \text{fdim}\, I\}$ such that, for all $i \in \text{fin}(I)$, $I(i) \leq N_{\sigma(i)}$.

Observe that that, if I' is thin, $I \subseteq I'$, and $\dim I = \dim I'$, then I is thin.

Remark 3.3 (Number of thin order ideals)**.** There cannot be more than $\binom{d}{i} \cdot i! \cdot \prod_{1\leq j\leq i}(N_j + 1) = \prod_{1\leq j\leq i}(d - j + 1)(N_j + 1)$ distinct thin order ideals of finite dimension i. As will become apparent in the proofs, this is what motivates the definition in (7).

Furthermore, if we let $\text{Idl}^{\textsf{thin}}(\mathbb{N}^d)$ denote the set of thin order ideals of $\mathbb{N}^d$, there is only one thin order ideal of finite dimension 0—namely $(\omega, \ldots, \omega)$—, and

$$\begin{aligned} |\text{Idl}^{\textsf{thin}}(\mathbb{N}^d)| &\leq 1 + \sum_{1\leq i\leq d} \prod_{1\leq j\leq i} (d-j+1)(N_j+1) \\ &= 1 + \sum_{1\leq i\leq d} (L_i - L_{i-1}) \\ &= 1 + L_d - L_0 = 1 + L_d \;. \end{aligned}$$

□

3.2. **Thinness Lemma.**

The crux of our result is the following lemma.

Lemma 3.4 (Thinness)**.** *Consider a (g, n_0)-controlled strongly monotone descending chain $D_0 \supsetneq D_1 \supsetneq \cdots$ of downwards-closed subsets of $\mathbb{N}^d$. If I_ℓ is a proper order ideal at some step ℓ, then I_ℓ is thin and $\ell \leq L_{\text{fdim}\, I_\ell}$.*

The proof of Lemma 3.4 proceeds by induction over the finite dimension $\text{fdim}\, I_\ell = d - \dim I_\ell$. For the base case where I_ℓ has full dimension $\dim I_\ell = d$, then $I_\ell = (\omega, \ldots, \omega)$ is thin and $D_\ell = \mathbb{N}^d$ is the full space, which can only occur at step

$\ell = 0 = L_0$. For the induction step, we first establish thinness with the following claim; note that, as just argued, an order ideal of dimension d is necessarily thin. We then follow with the bound on ℓ to complete the proof of Lemma 3.4.

Claim 3.5. *Let $0 \leq d' < d$ and assume that Lemma 3.4 holds for all proper order ideals I' of dimension $\dim I' > d'$. If I is any (not necessarily proper) order ideal of dimension $\dim I = d'$ appearing as a maximal ideal in the descending chain $D_0 \supsetneq D_1 \supsetneq \cdots$, then I is thin.*

Proof of Claim 3.5. Let k be a step where I appears in the descending chain $D_0 \supsetneq D_1 \supsetneq \cdots$, i.e., $I \in D_k$, and let us write $I_k \stackrel{\text{def}}{=} I$. If $k > 0$, since $D_k \subseteq D_{k-1}$, there exists an order ideal $I_{k-1} \in D_{k-1}$ such that $I_k \subseteq I_{k-1}$. If $k = 0$, or by repeating this argument if $k > 0$, we obtain a chain of order ideals (with decreasing indices)

$$I_k \subseteq I_{k-1} \subseteq \cdots \subseteq I_0 \tag{8}$$

where $I_m \in D_m$ for all $k \geq m \geq 0$. Every order ideal in that chain must have dimension at least $\dim I_k = d'$ since they all contain I_k. Two cases arise.

(1) If every order ideal in the chain (8) has dimension $\dim I_k$, then because the descending chain $D_0 \supsetneq D_1 \supsetneq \cdots$ is (g, n_0)-controlled, we have $\|I_0\| \leq n_0 = N_0$ and we know by Remark 3.1 that I_0 is thin. Since $I_k \subseteq I_0$ and $\dim I_k = \dim I_0$, I_k is also thin.

(2) Otherwise there exists a first index K along the chain (8) where the dimension increases, i.e., such that $\dim I_k < \dim I_K$ and $\dim I_m = \dim I_k$ for all $k \geq m > K$. Then I_K is proper, as otherwise D_{K+1} would contain two distinct but comparable order ideals in its canonical decomposition, namely I_K and I_{K+1}: indeed, $I_{K+1} \subseteq I_K$ and $\dim I_{K+1} = \dim I_k < \dim I_K$ imply $I_{K+1} \subsetneq I_K$. By assumption, Lemma 3.4 can be applied to I_K of dimension $\dim I_K > \dim I_k = d'$, thus I_K is thin and $K \leq L_{\mathrm{fdim}\, I_K}$.

Let us now show that I_{K+1} is thin, which will also yield that I_k is thin since $I_k \subseteq I_{K+1}$ and $\dim I_k = \dim I_{K+1}$.

Since $\dim I_{K+1} < \dim I_K$, we let $f \stackrel{\text{def}}{=} \dim I_K - \dim I_{K+1} = \mathrm{fdim}\, I_{K+1} - \mathrm{fdim}\, I_K > 0$. As furthermore $I_{K+1} \subseteq I_K$, $\omega(I_{K+1}) \subsetneq \omega(I_K)$ and we let $\{i_1, \ldots, i_f\} \stackrel{\text{def}}{=} \omega(I_K) \setminus \omega(I_{K+1}) = \mathrm{fin}(I_{K+1}) \setminus \mathrm{fin}(I_K)$.

Since I_K is thin, there exists a bijection $\sigma \colon \mathrm{fin}(I_K) \to \{1, \ldots, \mathrm{fdim}\,(I_K)\}$ such that $I_K(i) \leq N_{\sigma(i)}$ for all $i \in \mathrm{fin}(I_K)$. We extend σ to a bijection $\sigma' \colon \mathrm{fin}(I_K) \uplus \{i_1, \ldots, i_f\} \to \{1, \ldots, \mathrm{fdim}\, I_K + f\}$: we let $\sigma'(i) \stackrel{\text{def}}{=} \sigma(i)$ for all $i \in \mathrm{fin}(I_K)$, and $\sigma'(i_j) \stackrel{\text{def}}{=} \mathrm{fdim}\, I_K + j$ for all $1 \leq j \leq f$. Let us check that σ' witnesses the thinness of I_{K+1}.

- Because $I_{K+1} \subseteq I_K$, for all those $i \in \mathrm{fin}(I_K)$, $I_{K+1}(i) \leq I_K(i) \leq N_{\sigma(i)} = N_{\sigma'(i)}$.
- Since $K + 1 \leq L_{\mathrm{fdim}\, I_K} + 1$ and since the descending chain $D_0 \supsetneq D_1 \supsetneq \cdots$ is (g, n_0)-controlled, we have a bound of $g^{L_{\mathrm{fdim}\, I_K}+1}(n_0) = N_{\mathrm{fdim}\, I_K+1}$ on all the finite components of I_{K+1}, and in particular $I_{K+1}(i_j) \leq N_{\mathrm{fdim}\, I_K+1}$ for all $1 \leq j \leq f$. By Remark 3.1, we conclude that $I_{K+1}(i_j) \leq N_{\mathrm{fdim}\, I_K+j} = N_{\sigma'(i_j)}$ for all $1 \leq j \leq f$. [3.5]

Proof of Lemma 3.4. We have already argued for the base case, so let us turn to the inductive step where $\dim I_\ell < d$. If $\ell > 0$ and since our descending chain is strongly monotone, we can find an order ideal $I_{\ell-1}$ proper at step $\ell - 1$ such that $\dim I_\ell \leq \dim I_{\ell-1}$. Both if $\ell = 0$ or by repeating this argument, we obtain a sequence of order ideals (with decreasing indices)

$$I_\ell, I_{\ell-1}, \ldots, I_0 \tag{9}$$

where, for each $\ell > k \geq 0$, I_k is proper at step k, and $\dim I_{k+1} \leq \dim I_k$.

Let us decompose our sequence (9) by identifying the first step L where $\dim I_{L+1} < \dim I_L$; let $L \stackrel{\text{def}}{=} -1$ if this never occurs. After this step, for all $L \geq k \geq 0$, $\dim I_k > \dim I_\ell$. Within the initial segment, for $\ell \geq k > L$, the dimension $\dim I_k$ remains constant equal to $\dim I_\ell$, and the induction hypothesis allows to apply Claim 3.5 and infer that every order ideal I_k in this initial segment, and in particular I_ℓ among them, is thin.

It remains to provide a bound on ℓ. The $\ell - L$ order ideals in the initial segment are thin, and distinct since they are proper, hence by Remark 3.3,

$$\ell \leq L + \prod_{1 \leq i \leq \operatorname{fdim} I_\ell} (d - i + 1)(N_i + 1) \ . \tag{10}$$

If $L \geq 0$: we can apply the induction hypothesis to the proper order ideal I_L of finite dimension $\operatorname{fdim} I_L < \operatorname{fdim} I_\ell$ along with Remark 3.1 to yield $L \leq L_{\operatorname{fdim} I_L} \leq L_{\operatorname{fdim} I_\ell - 1}$ and therefore

$$\ell \leq L_{\operatorname{fdim} I_\ell - 1} + \prod_{1 \leq i \leq \operatorname{fdim} I_\ell} (d - i + 1)(N_i + 1) = L_{\operatorname{fdim} I_\ell} \ . \tag{11}$$

If $L = -1$: then (11) also holds since $L_{\operatorname{fdim} I_\ell - 1} \geq 0 > L$ in (10). □

We deduce a general combinatorial statement on the length of controlled strongly monotone descending chains, that generalises and refines [32, Thm. 4.4] thanks to thinness.

Theorem 3.6 (Length function for strongly monotone descending chains)**.** *Consider a (g, n_0)-controlled strongly monotone descending chain $D_0 \supsetneq \cdots \supsetneq D_\ell$ of downwards-closed subsets of $\mathbb{N}^d$. Then $\ell \leq L_d + 1$.*

Proof. In such a descending chain, either $\ell = 0 \leq L_d + 1$, or $\ell > 0$ and there must be an order ideal I proper at step $\ell - 1$, and I has finite dimension at most d. By Lemma 3.4 and Remark 3.1, $\ell - 1 \leq L_{\operatorname{fdim} I} \leq L_d$ in that case. □

3.3. **Thin Order Ideals and Filters.** Let us conclude this section with some consequences of Lemma 3.4 and Claim 3.5. Whereas thinness was posited *a priori* in the proof of Künnemann et al. [30, Thm. 3.3] and then shown to indeed allow a suitable decomposition of minimal covering executions and to eventually prove their result, here in the descending chain setting it is an inherent property of all the order ideals appearing in the chain, thereby providing a "natural" explanation for thinness.

Corollary 3.7. *Consider a (g, n_0)-controlled strongly monotone descending chain $D_0 \supsetneq D_1 \supsetneq \cdots$ of downwards-closed subsets of $\mathbb{N}^d$. Then every order ideal appearing in the chain is thin.*

Corollary 3.7 also entails a form of thinness of the minimal configurations in the complement of the downwards-closed sets D_k. Recall that such a complement is the upward-closure of a finite basis $B_k \stackrel{\text{def}}{=} \min_{\sqsubseteq} \mathbb{N}^d \setminus D_k$. Each element $\boldsymbol{v} \in B_k$ is a vector defining a so-called *(principal) order filter* $\uparrow\boldsymbol{v}$ of $\mathbb{N}^d$. Let us call a vector

$\boldsymbol{v} \in \mathbb{N}^d$ *nearly thin* if there exists a permutation $\sigma\colon \{1,\dots,d\} \to \{1,\dots,d\}$ such that, for all $1 \le i \le d$, $\boldsymbol{v}(i) \le N_{\sigma(i)} + 1$. We can relate thin order ideals with nearly thin order filters, which by Corollary 3.7 applies to every vector $\boldsymbol{v} \in \bigcup_k B_k$.

Proposition 3.8. *If every order ideal in the canonical decomposition of a downwards-closed set $D \subseteq \mathbb{N}^d$ is thin, then each $\boldsymbol{v} \in \min_{\sqsubseteq} \mathbb{N}^d \setminus D$ is nearly thin.*

Proof. Consider the canonical decomposition $D = I_1 \cup \cdots \cup I_m$ of D. Then $U \stackrel{\text{def}}{=} \mathbb{N}^d \setminus D = (\mathbb{N}^d \setminus I_1) \cap \cdots \cap (\mathbb{N}^d \setminus I_m)$. In turn, for each $1 \le j \le m$, $\mathbb{N}^d \setminus I_j = \bigcup_{i \in \operatorname{fin}(I_j)} \uparrow\big((I_j(i)+1)\cdot \boldsymbol{e}_i\big)$ where $\boldsymbol{e}_i$ denotes the unit vector such that $\boldsymbol{e}_i(i) \stackrel{\text{def}}{=} 1$ and $\boldsymbol{e}_i(j) \stackrel{\text{def}}{=} 0$ for all $j \ne i$. Distributing intersections over unions, we obtain that

$$U = \bigcup_{(i_1,\dots,i_m)\in\operatorname{fin}(I_1)\times\cdots\times\operatorname{fin}(I_m)} \ \bigcap_{1\le j\le m} \uparrow\big((I_j(i_j)+1)\cdot \boldsymbol{e}_{i_j}\big)\ . \tag{$*$}$$

For two order filters $\uparrow\boldsymbol{v}$ and $\uparrow\boldsymbol{v}'$, $(\uparrow\boldsymbol{v}) \cap (\uparrow\boldsymbol{v}') = \uparrow(\boldsymbol{v} \vee \boldsymbol{v}')$ where $\boldsymbol{v} \vee \boldsymbol{v}'$ denotes the component-wise maximum of $\boldsymbol{v}$ and $\boldsymbol{v}'$. Therefore, by $(*)$, any $\boldsymbol{v} \in \min_{\sqsubseteq} U$ is of the form

$$\boldsymbol{v}_{i_1,\dots,i_m} \stackrel{\text{def}}{=} \bigvee_{1\le j\le m} \big((I_j(i_j)+1)\cdot \boldsymbol{e}_{i_j}\big) \tag{$\dagger$}$$

for some $(i_1,\dots,i_m) \in \operatorname{fin}(I_1)\times\cdots\times\operatorname{fin}(I_m)$. Note that not all the vectors $\boldsymbol{v}_{i_1,\dots,i_m}$ defined by $(\dagger)$ are necessarily minimal in U, but that

$$\min_{\sqsubseteq} U = \min_{\sqsubseteq}\{\boldsymbol{v}_{i_1,\dots,i_m} \mid (i_1,\dots,i_m) \in \operatorname{fin}(I_1)\times\cdots\times\operatorname{fin}(I_m)\}\ . \tag{$\ddagger$}$$

Assume by contradiction that there exists some minimal vector $\boldsymbol{v} \in \min_{\sqsubseteq} U$ that is not nearly thin. Without loss of generality, $\boldsymbol{v}(1) \le \boldsymbol{v}(2) \le \cdots \le \boldsymbol{v}(d)$, as otherwise we could apply a suitable permutation of $\{1,\dots,d\}$ on the components of each ideal $I_j \in D$. Then, because $N_i \le N_{i'}$ for all $i < i'$ by Remark 3.1, $\boldsymbol{v}$ not being nearly thin entails that there exists an index $k \in \{1,\dots,d\}$ such that $\boldsymbol{v}(k) > N_k + 1$ but $\boldsymbol{v}(i) \le N_i + 1$ for all $i < k$.

By $(\ddagger)$, there exists $(i_1,\dots,i_m) \in \operatorname{fin}(I_1)\times\cdots\times\operatorname{fin}(I_m)$ such that $\boldsymbol{v} = \boldsymbol{v}_{i_1,\dots,i_m}$. We are going to show that there exists $(i'_1,\dots,i'_m) \in \operatorname{fin}(I_1)\times\cdots\times\operatorname{fin}(I_m)$ such that $\boldsymbol{v}_{i'_1,\dots,i'_m} \sqsubsetneq \boldsymbol{v}$, which by $(\ddagger)$ contradicts the minimality of $\boldsymbol{v}$.

Looking more closely at the individual components of $\boldsymbol{v}$ in $(\dagger)$, define for all $1 \le i \le d$ the set $S_i \stackrel{\text{def}}{=} \{1 \le j \le m \mid i_j = i\}$ of indices $j \in \{1,\dots,m\}$ such that the value of $\boldsymbol{v}(i)$ "stems" from I_j. Then

$$\boldsymbol{v}_{i_1,\dots,i_m}(i) = \begin{cases} 0 & \text{if } S_i = \emptyset \\ \max\{I_j(i)+1 \mid j \in S_i\} & \text{otherwise.} \end{cases} \tag{$**$}$$

In particular, for the kth component, $S_k \ne \emptyset$ and we let $V_k \stackrel{\text{def}}{=} \{j \in S_k \mid I_j(k) > N_k\}$ denote the indices j of the ideals $I_j \in D$ responsible for the violation of near thinness.

Example 3.9. Let us illustrate the previous notations. Let $d \stackrel{\text{def}}{=} 4$ and assume for the sake of simplicity that

$$N_1 \stackrel{\text{def}}{=} 2\ , \qquad N_2 \stackrel{\text{def}}{=} 4\ , \qquad N_3 \stackrel{\text{def}}{=} 6\ , \qquad N_4 \stackrel{\text{def}}{=} 8\ .$$

Consider $D \overset{\text{def}}{=} \{I_1, I_2, I_3, I_4, I_5\}$ with

$$I_1 \overset{\text{def}}{=} (1,4,6,7)\ , \qquad I_2 \overset{\text{def}}{=} (2,6,4,8)\ , \qquad I_3 \overset{\text{def}}{=} (3,1,7,6)\ ,$$
$$I_4 \overset{\text{def}}{=} (3,1,7,6)\ , \qquad I_5 \overset{\text{def}}{=} (4,5,3,0)\ .$$

Then $\boldsymbol{v} \overset{\text{def}}{=} (2,7,7,7) = \boldsymbol{v}_{3,2,4,1,2}$ is not nearly thin with $k = 2$, stem sets

$$S_1 = \{4\}\ , \qquad S_2 = \{2,5\}\ , \qquad S_3 = \{1\}\ , \qquad S_4 = \{3\}\ ,$$

and $V_2 = \{2,5\}$, and indeed $\boldsymbol{v}(2)$ stems from the ideals I_2 and I_5, which are such that $I_2(2) = 6 > N_2$ and $I_5(2) = 5 > N_2$. □

For all $1 \le j \le m$, because I_j is thin, there exists a bijection $\sigma_j\colon \operatorname{fin}(I_j) \to \{1, \dots, \operatorname{fdim} I_j\}$ such that, for all $i \in \operatorname{fin}(I_j)$, $I_j(i) \le N_{\sigma_j(i)}$. Without loss of generality, we can assume that for all $i, i' \in \operatorname{fin}(I_j)$, $I_j(i) \le I_j(i')$ whenever $\sigma_j(i) < \sigma_j(i')$.

Example 3.9 (continuing from p. 9)**.** Here are suitable bijections witnessing thinness:

$$\sigma_1 = (1\,2\,3\,4)\ , \qquad \sigma_2 = (1\,3\,2\,4)\ , \qquad \sigma_3 = (2\,1\,4\,3)\ ,$$
$$\sigma_4 = (2\,1\,4\,3)\ , \qquad \sigma_5 = (3\,4\,2\,1)\ .$$
□

For every $j \in V_k$, $\sigma_j^{-1}(\{1, \dots, k\}) \setminus \{1, \dots, k-1\}$ is non empty. Therefore it contains an element $i'_j \ge k$ such that $I_j(i'_j) \le N_k$. For every $1 \le j \le m$ such that $j \notin V_k$, let $i'_j \overset{\text{def}}{=} i_j$.

Let us check that $\boldsymbol{v}_{i'_1,\dots,i'_m} \subsetneq \boldsymbol{v}$, which will allow to conclude. Define $S'_i \overset{\text{def}}{=} \{1 \le j \le m \mid i'_j = i\}$ for each $1 \le i \le d$; then equation (∗∗) holds mutatis mutandis for $\boldsymbol{v}_{i'_1,\dots,i'_m}$ and

for $i < k$**:** $S'_i = S_i$ hence $\boldsymbol{v}_{i'_1,\dots,i'_m}(i) = \boldsymbol{v}(i)$;

for $i = k$**:** $S'_k = S_k \setminus V_k = \{j \in S_k \mid I_j(k) \le N_k\}$ hence $\boldsymbol{v}_{i'_1,\dots,i'_m}(k) \le N_k + 1 < \boldsymbol{v}(k)$ by definition of k;

for $i > k$**:** $S'_i = S_i \cup \{j \in V_k \mid i'_j = i\}$ hence $\boldsymbol{v}_{i'_1,\dots,i'_m}(i) = \max\{I_j(i) + 1 \mid j \in S'_i\} = \max(\max\{I_j(i) + 1 \mid j \in S_i\}, \max\{I_j(i) + 1 \mid j \in V_k \text{ and } i'_j = i\})$.

- On the one hand, $\max\{I_j(i) + 1 \mid j \in S_i\} = \boldsymbol{v}(i)$.
- On the other hand, $I_j(i'_j) \le N_k$ for all $j \in V_k$ by definition of i'_j, hence $\max\{I_j(i) + 1 \mid j \in V_k \text{ and } i'_j = i\} \le N_k + 1 < \boldsymbol{v}(k)$ by definition of k.

As $\boldsymbol{v}(k) \le \boldsymbol{v}(i)$ by assumption since $i > k$, we conclude $\boldsymbol{v}_{i'_1,\dots,i'_m}(i) = \boldsymbol{v}(i)$.

Example 3.9 (continuing from p. 10)**.** We have $\sigma_2^{-1}(\{1,2\}) = \{1,3\}$ and $\sigma_5^{-1}(\{1,2\}) = \{3,4\}$, hence we can pick $i'_2 \overset{\text{def}}{=} 3$ and $i'_5 \overset{\text{def}}{=} 4$. This defines $\boldsymbol{v}_{3,3,4,1,4}$ with stem sets

$$S'_1 = \{4\}\ , \qquad S'_2 = \emptyset\ , \qquad S'_3 = \{1,2\}\ , \qquad S'_4 = \{3,5\}\ .$$

Then $\boldsymbol{v}_{3,3,4,1,4} = (2,0,7,7) \subsetneq \boldsymbol{v}$ as desired. □

4. Applications

We describe two applications of Theorem 3.6 in this section. The first application in § 4.2 is to the coverability problem in vector addition systems, and relies on the analysis of the backward coverability algorithm done in [32]. Thus we can indeed recover the improved upper bound of Künnemann et al. [30] for the coverability problem in the more general setting of descending chains, and show that the backward coverability algorithm achieves this $n^{2^{O(d)}}$ upper bound (see Corollary 4.5).

The second application in § 4.3 focuses on the coverability problem in invertible affine nets, a class introduced by Benedikt et al. [7], who analysed the complexity of the problem through a reduction to zeroness in invertible polynomial automata. We give a direct analysis of the complexity of the backward coverability algorithm, which follows the same lines as in the VAS case, and allows to improve on the 2EXPSPACE upper bound shown in [7] for the problem, by showing that it is actually EXPSPACE-complete (see Corollary 4.13). This application additionally illustrates the usefulness of considering strongly monotone descending chains rather than the ω-monotone ones, as the descending chains constructed by the backward algorithm for invertible affine nets are in general not ω-monotone.

As both applications take place in the framework of well-structured transition systems [1, 22], we start with a quick refresher on this framework, the backward coverability algorithm, and its dual view using downwards-closed sets [32] in the upcoming § 4.1.

4.1. Coverability in Well-Structured Transition Systems.

Well-structured transition systems (WSTS) form an abstract family of computational models where the set of configurations is equipped with a well-quasi-ordering "compatible" with the computation steps. This wqo ensures the termination of generic algorithms checking some important behavioural properties like coverability and termination. While the idea can be traced back to the 1980's [21], this framework has been especially popularised through two landmark surveys [1, 22] that emphasised its wide applicability, and new WSTS models keep being invented in multiple areas to this day.

4.1.1. *Well-Structured Transition Systems.* A *well-structured transition system* (WSTS) [1, 22] is a triple $(X, \to, \leq)$ where X is a set of configurations, $\to \subseteq X \times X$ is a transition relation, and $(X, \leq)$ is a wqo with the following *compatibility* condition: if $x \leq x'$ and $x \to y$, then there exists $y' \geq y$ with $x' \to y'$.

The coverability problem below corresponds to the verification of safety properties, i.e., to checking that no bad configuration can ever be reached from a given initial configuration $s \in X$. Here we are given an error configuration $t \in X$, and we assume that any configuration larger than t is also an error.

Problem (Coverability in well-structured transition systems)**.**

input: a well-structured transition system $(X, \to, \leq)$ and two configurations s and t in X

question: does s *cover* t, i.e., does there exist $t' \in X$ such that $s \to^* t' \geq t$?

4.1.2. *The Backward Coverability Algorithm.* The first published version of this algorithm seems to date back to [3], where it was used to show the decidability of coverability in vector addition systems extended with reset capabilities, before it was rediscovered and generalised to well-structured transition systems [1].

The Algorithm. Given an instance of the coverability problem, the *backward coverability algorithm* [3, 1, 22] computes (a finite basis for) the upwards-closed set

$$U_* \stackrel{\text{def}}{=} \{x \in X \mid \exists t' \geq t \,.\, x \to^* t'\} \tag{12}$$

of all the configurations that cover t, and then checks whether $s \in U_*$.

The set U_* itself is computed by letting

$$U_0 \stackrel{\text{def}}{=} \uparrow t \,, \qquad\qquad U_{k+1} \stackrel{\text{def}}{=} U_k \cup \mathrm{Pre}_\exists(U_k) \,, \tag{13}$$

where, for a set $S \subseteq X$,

$$\mathrm{Pre}_\exists(S) \stackrel{\text{def}}{=} \{x \in X \mid \exists y \in S \,.\, x \to y\}.$$

Then $U_k = \{x \in X \mid \exists t' \geq t \,.\, x \to^{\leq k} t'\}$ is the set of configurations that can cover t in at most k steps. Equation (13) defines a chain $U_0 \subseteq U_1 \subseteq \cdots$ of upwards-closed subsets of X. Furthermore, if $U_\ell = U_{\ell+1}$ at some step, then we have reached stabilisation: $U_\ell = U_{\ell+k} = U_*$ for all k. Thus we focus in this algorithm on ascending chains $U_0 \subsetneq U_1 \subsetneq \cdots$, which are finite thanks to the ascending chain condition of the wqo $(X, \leq)$. In order to turn (13) into an actual algorithm, one needs to make some effectiveness assumptions on $(X, \to, \leq)$, typically that $\leq$ is decidable and a finite basis for $\mathrm{Pre}_\exists(\uparrow x)$ can be computed for all $x \in X$ [22, Prop. 3.5].

A Dual View. Lazić and Schmitz [32] take a dual view of the algorithm and define from (13) a descending chain $D_0 \supsetneq D_1 \supsetneq \cdots$ of the same length where

$$D_k \stackrel{\text{def}}{=} X \setminus U_k \tag{14}$$

for each k; this stops with $D_* = X \setminus U_*$ the set of configurations that do *not* cover t. The entire computation in (13) can be recast in this dual view, by setting

$$D_0 \stackrel{\text{def}}{=} X \setminus \uparrow t \,, \qquad\qquad D_{k+1} \stackrel{\text{def}}{=} D_k \cap \mathrm{Pre}_\forall(D_k) \,, \tag{15}$$

where, for a set $S \subseteq X$,

$$\mathrm{Pre}_\forall(S) \stackrel{\text{def}}{=} \{x \in X \mid \forall y \in X \,.\, (x \to y \implies y \in S)\} = X \setminus (\mathrm{Pre}_\exists(X \setminus S)).$$

Under some effectiveness assumptions, in particular for manipulating ideal representations over X, this can be turned into an actual algorithm [32, Sec. 3.1].

4.2. **Coverability in Vector Addition Systems.** Vector addition systems are a well-established model for simple concurrent processes [27] equivalent to Petri nets, with far-reaching connections to many topics in theoretical computer science. In particular, their coverability problem, which essentially captures safety checking, has been thoroughly investigated from both a theoretical [27, 34, 37, 13, 32, 30] and a more practical [19, 8, 24, 10] standpoint.

4.2.1. *Vector Addition Systems.* A d-dimensional *vector addition system* (VAS) [27] is a finite set $\boldsymbol{A}$ of vectors in $\mathbb{Z}^d$. It defines a well-structured transition system $(\mathbb{N}^d, \to_{\boldsymbol{A}}, \sqsubseteq)$ with $\mathbb{N}^d$ as set of configurations and transitions $\boldsymbol{u} \to_{\boldsymbol{A}} \boldsymbol{u} + \boldsymbol{a}$ for all $\boldsymbol{u}$ in $\mathbb{N}^d$ and $\boldsymbol{a}$ in $\boldsymbol{A}$ such that $\boldsymbol{u} + \boldsymbol{a}$ is in $\mathbb{N}^d$. We work with a unary encoding, and let $\|\boldsymbol{u}\| \stackrel{\text{def}}{=} \max_{1 \leq i \leq d} |\boldsymbol{u}(i)|$ and $\|\boldsymbol{A}\| \stackrel{\text{def}}{=} \max_{\boldsymbol{a} \in \boldsymbol{A}} \|\boldsymbol{a}\|$ for all $\boldsymbol{u} \in \mathbb{Z}^d$ and $\boldsymbol{A} \subseteq \mathbb{Z}^d$ finite.

The coverability problem in vector addition systems was first shown decidable in 1969 by Karp and Miller [27], before being proven EXPSPACE-complete when d is part of the input by Lipton [34] and Rackoff [37]. Note that the problem parameterised by d is trivial for $d = 1$ (a target $\boldsymbol{t}$ is coverable if and only if $\boldsymbol{s} \geq \boldsymbol{t}$ or there exists $\boldsymbol{a} \in \boldsymbol{A}$ with $\boldsymbol{a} > 0$), hence we will assume $d \geq 2$.

4.2.2. *Complexity Upper Bounds.* The dual backward coverability algorithm of §4.1.2 is straightforward to instantiate in the case of a vector addition system. Figure 1 displays the computed descending chain for the 2-dimensional VAS $\boldsymbol{A}_{\div 2} \stackrel{\text{def}}{=} \{(-2, 1)\}$ and target configuration $\boldsymbol{t} \stackrel{\text{def}}{=} (0, 5)$ [32, Ex. 3.6].

Fact 4.1 ([32, claims 3.9 and 4.3])**.** *The descending chain $D_0 \supsetneq D_1 \supsetneq \cdots$ defined by equations (13–15) for a d-dimensional VAS $\boldsymbol{A}$ and a target vector $\boldsymbol{t}$ is (g, n_0)-controlled for $g(x) \stackrel{\text{def}}{=} x + \|\boldsymbol{A}\|$ and $n_0 \stackrel{\text{def}}{=} \|\boldsymbol{t}\|$, and is ω-monotone.*

The length of the descending chain defined by equations (13–15) is the main source of complexity for the whole backward coverability algorithm, and we can apply our own Theorem 3.6 instead of [32, Thm. 4.4] in order to prove the following bound on this length, where the combinatorics are somewhat similar to those of [30, Lem. 3.5].

Theorem 4.2. *The backward coverability algorithm terminates after at most $n^{2^{O(d)}}$ iterations on a d-dimensional VAS encoded in unary.*

Proof. Let n be the size of the input to the coverability problem; we assume in the following that $n, d \geq 2$. By Fact 4.1 and due to the unary encoding, the descending chain $D_0 \supsetneq D_1 \supsetneq \cdots \supsetneq D_\ell = D_*$ is (g, n_0)-controlled for $g(x) \stackrel{\text{def}}{=} x + n$ and $n_0 \stackrel{\text{def}}{=} n$, and is ω-monotone and thus strongly monotone. By Theorem 3.6, $\ell \leq L_d + 1$. Let us bound this value.

Claim 4.3. *Let $g(x) \stackrel{\text{def}}{=} x + n$ and $n_0 \stackrel{\text{def}}{=} n$. Then, for all $i \leq d$,*

$$N_{i+1} = n \cdot (L_i + 2) \;, \qquad\qquad L_i + 4 \leq n^{3^i \cdot (\lg d + 1)} \;.$$

Proof of Claim 4.3. In the case of N_{i+1}, by the definition of N_{i+1} in (6), $N_{i+1} = g^{L_i+1}(n_0) = n + (L_i + 1) \cdot n = n \cdot (L_i + 2)$ as desired.

Regarding L_i, we proceed by induction over i. For the base case $i = 0$, $L_0 + 4 = 4 \leq n^{3^0 \cdot (\lg d + 1)}$ since we assumed $n, d \geq 2$. For the induction step, by the definition of L_{i+1} in (7)

$$\begin{aligned} L_{i+1} + 4 &= L_i + 4 + \prod_{0 \leq j \leq i} (d - j)(N_{j+1} + 1) \\ &\leq L_i + 4 + \prod_{0 \leq j \leq i} (d - j) \cdot n \cdot (L_j + 3) \\ &\leq 2 \cdot (dn)^{i+1} \cdot \prod_{0 \leq j \leq i} (L_j + 3) \;. \end{aligned}$$

Here, since $n \geq 2$,

$$2 \cdot (dn)^{i+1} \leq n^{(i+1)(\lg d + 1) + 1}$$

and by induction hypothesis for $j \leq i$

$$\prod_{0 \leq j \leq i} (L_j + 3) \leq n^{\sum_{0 \leq j \leq i} 3^j (\lg d + 1)} \;.$$

Thus, it only remains to see that, since $i > 0$,

$$\begin{aligned} 3^{i+1} \cdot (\lg d + 1) &= (1 + 2 \cdot \sum_{0 \leq j \leq i} 3^j) \cdot (\lg d + 1) \\ &\geq (1 + 3^0 + 3^i) \cdot (\lg d + 1) + \sum_{0 \leq j \leq i} 3^j \cdot (\lg d + 1) \\ &\geq (i + 1) \cdot (\lg d + 1) + 1 + \sum_{0 \leq j \leq i} 3^j \cdot (\lg d + 1) \;. \end{aligned} \qquad [4.3]$$

Thus $L_d + 1 \leq n^{3^d \cdot (\lg d + 1)}$ by Claim 4.3, which is in $n^{2^{O(d)}}$. □

Remark 4.4 (Branching or alternating vector addition systems)**.** The improved upper bound parameterised by the dimension d in Theorem 4.2 also applies to some extensions of vector addition systems, for which Lazić and Schmitz [32] have shown that the backward coverability algorithm was constructing an ω-monotone descending chain controlled as in Fact 4.1, namely

- in [32, claims 6.7 and 6.8] for bottom-up coverability in branching vector addition systems (BVAS)—which is 2EXP-complete [16]—, and
- in [32, claims 5.4 and 5.5] for top-down coverability in alternating vector addition systems (AVAS)—which is 2EXP-complete as well [15]. □

Recall that U_ℓ is the set of configurations that can cover the target $\boldsymbol{t}$ in at most ℓ steps, hence Theorem 4.2 provides an alternative proof for [30, Thm. 3.3]: if there exists a covering execution, then there is one of length in $n^{2^{O(d)}}$, from which an algorithm in $n^{2^{O(d)}}$ follows by [30, Thm. 3.2]. Regarding the optimality of Theorem 4.2, recall that Lipton [34] shows an $n^{2^{\Omega(d)}}$ lower bound on the length of a minimal covering execution, which translates into the same lower bound on the number ℓ of iterations of the backward coverability algorithm [13, Cor. 2]. Finally, this also yields an improved upper bound on the complexity of the (original) backward coverability algorithm. Here, we can rely on the analysis performed by Bozzelli and Ganty [13, Sec. 3] and simply replace Rackoff's $n^{2^{O(d\lg d)}}$ bound on the length of minimal covering executions by the bound from Theorem 4.2.

Corollary 4.5. *The backward coverability algorithm runs in time $n^{2^{O(d)}}$ on d-dimensional VAS encoded in unary.*

Proof. Let n be the size of the input to the coverability problem and $U_0 \subsetneq U_1 \subsetneq \cdots \subsetneq U_\ell = U_*$ be the ascending chain constructed by the backward coverability according to (13). By Theorem 4.2, ℓ is in $n^{2^{O(d)}}$.

Let $B_k \stackrel{\text{def}}{=} \min_\sqsubseteq U_k$ be the minimal basis at each step k. The algorithm computes B_{k+1} from B_k as per (13) by computing $\min_\sqsubseteq \mathrm{Pre}_\exists(\uparrow\boldsymbol{v})$ for each $\boldsymbol{v} \in B_k$, adding the elements of B_k, and removing any non-minimal vector. Thus each step can be performed in time polynomial in n, d, and the number of vectors in B_k. Here, Bozzelli and Ganty's analysis in [13, Sec. 3] shows that $\|\boldsymbol{v}'\| \leq g(\|\boldsymbol{v}\|)$ for all $\boldsymbol{v}' \in \min_\sqsubseteq \mathrm{Pre}_\exists(\uparrow\boldsymbol{v})$, yielding a bound of $|B_k| \leq (g^k(n)+1)^d \leq ((\ell+1)\cdot n+1)^d$, which is still in $n^{2^{O(d)}}$.

We can do slightly better. By Corollary 3.7, all the ideals in the canonical decomposition of $D_k \stackrel{\text{def}}{=} \mathbb{N}^d \setminus U_k$ are thin, and in turn Proposition 3.8 shows that all the vectors in B_k are nearly thin. Accordingly, let us denote by $\mathrm{Fil}^{\mathsf{thin}+1}(\mathbb{N}^d)$ the set of order filters $\uparrow\boldsymbol{v}$ such that $\boldsymbol{v}$ is nearly thin. Then $|B_k| \leq |\mathrm{Fil}^{\mathsf{thin}+1}(\mathbb{N}^d)|$, and the latter is in $n^{2^{O(d)}}$:

$$
\begin{aligned}
|\mathrm{Fil}^{\mathsf{thin}+1}(\mathbb{N}^d)| &\leq d! \cdot \prod_{1\leq i\leq d}(N_i+2) \\
&\leq d!\cdot n^d \cdot \prod_{0\leq i\leq d-1}(L_i+4) && \text{(by Claim 4.3 on } N_i\text{)}\\
&\leq n^{2d+\sum_{0\leq i\leq d-1} 3^i\cdot(\lg d+1)} && \text{(because } d\leq n \text{ and by Claim 4.3 on } L_i\text{)}\\
&\leq n^{3^d\cdot(\lg d+1)}\,. && (16)
\end{aligned}
$$

Therefore, the overall complexity of the backward coverability algorithm is polynomial in ℓ, $\max_{0\leq k\leq \ell}|B_k|$, n, and d, which is in $n^{2^{O(d)}}$.

Observe that the dual version of the backward coverability algorithm enjoys the same upper bound: at each step k, the algorithm computes D_{k+1} from D_k as per (15); this computation of D_{k+1} can be performed in time polynomial in n, d, and the number of ideals in the canonical decomposition of D_k [32, Sec. 3.2.1]. By Corollary 3.7 and Remark 3.3, $|D_k| \leq |\mathrm{Idl}^{\mathsf{thin}}(\mathbb{N}^d)| \leq 1 + L_d$, hence the overall complexity of the dual algorithm is polynomial in L_d, n, and d, which is still in $n^{2^{O(d)}}$. □

The bounds in $n^{2^{O(d)}}$ for $\|\boldsymbol{v}\| \leq N_d + 1$ for all $\boldsymbol{v} \in \min_{\sqsubseteq} U_k$ and for $|\min_{\sqsubseteq} U_k| \leq |\mathrm{Fil}^{\mathsf{thin}+1}(\mathbb{N}^d)|$ in the previous proof also improve on the corresponding bounds in [44, Thm. 9] and [13, Thm. 2]. Recall that Künnemann et al. [30, Thm. 4.2] show that, assuming the exponential time hypothesis, there does not exist a deterministic $n^{o(2^d)}$ time algorithm deciding coverability in unary encoded d-dimensional VAS, hence the backward coverability algorithm is conditionally optimal.

4.3. **Coverability in Affine Nets.**

Affine nets [23], also known as affine vector addition systems, are a broad generalisation of VAS and Petri nets encompassing multiple extended VAS operations designed for greater modelling power.

4.3.1. *Affine Nets.*

A d-dimensional (well-structured) *affine net* [23] is a finite set $\mathcal{N}$ of triples $(\boldsymbol{a}, A, \boldsymbol{b}) \in \mathbb{N}^d \times \mathbb{N}^{d\times d} \times \mathbb{N}^d$. It defines a well-structured transition system $(\mathbb{N}^d, \rightarrow_{\mathcal{N}}, \sqsubseteq)$ with $\mathbb{N}^d$ as set of configurations and transitions $\boldsymbol{u} \rightarrow_{\mathcal{N}} A \cdot (\boldsymbol{u} - \boldsymbol{a}) + \boldsymbol{b}$ for all $\boldsymbol{u}$ in $\mathbb{N}^d$ and $(\boldsymbol{a}, A, \boldsymbol{b})$ in $\mathcal{N}$ such that $\boldsymbol{u} - \boldsymbol{a}$ is in $\mathbb{N}^d$. This model encompasses notably

- VAS and Petri nets when (each such) A is the identity matrix I_d,
- *reset nets* [2, 3] when A is component-wise smaller or equal to I_d,
- *transfer nets* [14] when the sum of values in every column of A is one,
- *post self-modifying nets* [43]—also known as *strongly increasing affine nets* [23, 12]—when A is component-wise larger or equal to I_d, and
- *invertible* affine nets [7] when A is invertible over the rationals, i.e., $A \in \mathsf{GL}_d(\mathbb{Q})$.

As in the case of VAS, we will work with a unary encoding, and we let $\|\mathcal{N}\| \stackrel{\text{def}}{=} \max\{\|\boldsymbol{a}\| \mid (\boldsymbol{a}, A, \boldsymbol{b}) \in \mathcal{N}\}$; note that the entries from $\boldsymbol{b}$ and A are not taken into account.

Example 4.6. Consider the affine nets

$$\mathcal{N}_1 \stackrel{\text{def}}{=} \left\{\begin{bmatrix} 2 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix}\right\}$$

$$\mathcal{N}_2 \stackrel{\text{def}}{=} \left\{\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \end{bmatrix}\right\}$$

$$\mathcal{N}_3 \stackrel{\text{def}}{=} \left\{\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & 1 \\ 2 & 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \end{bmatrix}\right\}.$$

Then $\mathcal{N}_1$ defines the same WSTS as the 2-dimensional VAS $\boldsymbol{A}_{\div 2} = \{(-2, 1)\}$. Focusing on the effects of their transition matrices, $\mathcal{N}_2$ performs a transfer from its second component into its first component, while $\mathcal{N}_3$ sums the values of its first two components into the first one, and puts the double of its first component into its second one. ☐

The coverability problem for reset VAS was first shown decidable in 1978 by Arnold and Latteux [3] using the backward coverability algorithm, and the same algorithm applies to all affine nets [18, 23]. Its complexity is considerable: their coverability problem has already an Ackermannian complexity in the reset or transfer cases [42, 20, 40]. In the strongly increasing case, Bonnet, Finkel, and Praveen [12, Lem. 11 and Thm. 13] show how to adapt Rackoff's original argument to derive an upper bound in $n^{2^{O(d \lg d)}}$ on the length of minimal coverability witnesses, with an EXPSPACE upper bound for the problem when d is part of the input, while in the invertible case, Benedikt et al. [7, Thm. 6] show a 2EXPSPACE upper bound.

Control. Before we turn to the case of invertible affine nets, let us show that the descending chains defined by the backward coverability algorithm for affine nets are controlled, with a control very similar to the VAS case (c.f. Fact 4.1).

Proposition 4.7. *The descending chain $D_0 \supsetneq D_1 \supsetneq \cdots$ defined by equations (13–15) for a d-dimensional affine net $\mathcal{N}$ and a target vector $\boldsymbol{t}$ is (g, n_0)-controlled for $g(x) \stackrel{\text{def}}{=} x + \|\mathcal{N}\|$ and $n_0 \stackrel{\text{def}}{=} \|\boldsymbol{t}\|$.*

Proof. Rather than handling $\mathrm{Pre}_\forall$ computations directly, we use the fact that $\mathrm{Pre}_\forall(S) = \mathbb{N}^d \setminus (\mathrm{Pre}_\exists(\mathbb{N}^d \setminus S))$ for all $S \subseteq \mathbb{N}^d$ and the following statement on $\mathrm{Pre}_\exists$ computations.

Claim 4.8. *If $\boldsymbol{u}' \in \min_\sqsubseteq \mathrm{Pre}_\exists(\uparrow\boldsymbol{u})$, then $\|\boldsymbol{u}'\| \leq \|\boldsymbol{u}\| + \|\mathcal{N}\|$.*

Proof of Claim 4.8. In such a situation, there exists a triple $(\boldsymbol{a}, A, \boldsymbol{b}) \in \mathcal{N}$ such that $\boldsymbol{u}' \sqsupseteq \boldsymbol{a}$ and $A \cdot (\boldsymbol{u}' - \boldsymbol{a}) \sqsupseteq \boldsymbol{u} - \boldsymbol{b}$. Let $\boldsymbol{y}$ be defined by $\boldsymbol{y}(i) \stackrel{\text{def}}{=} \max(\boldsymbol{u}(i), \boldsymbol{b}(i)) - \boldsymbol{b}(i)$ for all $1 \leq i \leq d$, thus of size $\|\boldsymbol{y}\| \leq \|\boldsymbol{u}\|$. Then $\boldsymbol{u}' = \boldsymbol{x} + \boldsymbol{a}$ where $\boldsymbol{x}$ is a $\sqsubseteq$-minimal solution of the system of inequalities $A\boldsymbol{x} \sqsupseteq \boldsymbol{y}$.

We are going to show that if $\boldsymbol{x}$ is an $\sqsubseteq$-minimal solution, then $\|\boldsymbol{x}\| \leq \|\boldsymbol{y}\|$. This will yield the result, as then $\|\boldsymbol{u}'\| \leq \|\boldsymbol{y}\| + \|\boldsymbol{a}\| \leq \|\boldsymbol{u}\| + \|\mathcal{N}\|$. Assume by contradiction that $\boldsymbol{x}$ is a $\sqsubseteq$-minimal solution with $\boldsymbol{x}(j) > \|\boldsymbol{y}\|$ for some $1 \leq j \leq d$. Consider $\boldsymbol{x}'$ defined by $\boldsymbol{x}'(j) \stackrel{\text{def}}{=} \|\boldsymbol{y}\|$ and $\boldsymbol{x}'(i) \stackrel{\text{def}}{=} \boldsymbol{x}(i)$ for all $i \neq j$; note that $\boldsymbol{x}' \sqsubsetneq \boldsymbol{x}$. Let us show that $\boldsymbol{x}'$ is also a solution, i.e., that $A\boldsymbol{x}' \sqsupseteq \boldsymbol{y}$: for all $1 \leq i \leq d$,

- if $A(i,j) > 0$ then $\sum_{1 \leq k \leq d} A(i,k) \cdot \boldsymbol{x}'(k) \geq \boldsymbol{x}'(j) \geq \|\boldsymbol{y}\| \geq \boldsymbol{y}(i)$, and
- otherwise $\sum_{1 \leq k \leq d} A(i,k) \cdot \boldsymbol{x}'(k) = \sum_{1 \leq k \leq d} A(i,k) \cdot \boldsymbol{x}(k) \geq \boldsymbol{y}(i)$ since $\boldsymbol{x}$ is a solution.

Thus $\boldsymbol{x}'$ is a solution, contradicting the $\sqsubseteq$-minimality of $\boldsymbol{x}$. [4.8]

Now, since $D_0 = \mathbb{N}^d \setminus \uparrow\boldsymbol{t}$, $\|D_0\| \leq \|\boldsymbol{t}\| - 1$ by [32, Lem. 3.8]. Regarding the control function, $D_{k+1} = D_k \cap \mathrm{Pre}_\forall(D_k)$ is such that $\|D_{k+1}\| \leq \max(\|D_k\|, \|\mathrm{Pre}_\forall(D_k)\|)$ also by [32, Lem. 3.8]. In turn, regarding $\mathrm{Pre}_\forall(D_k) = \mathbb{N}^d \setminus \mathrm{Pre}_\exists(U_k)$, the minimal elements $\boldsymbol{u}$ of $U_k = \mathbb{N}^d \setminus D_k$ have size $\|\boldsymbol{u}\| \leq \|D_k\| + 1$ still by [32, Lem. 3.8], thus the minimal elements $\boldsymbol{u}'$ of $\mathrm{Pre}_\exists(U_k)$ have size $\|\boldsymbol{u}'\| \leq \|D_k\| + 1 + \|\mathcal{N}\|$ by Claim 4.8, hence $\|\mathrm{Pre}_\forall(D_k)\| \leq \|D_k\| + \|\mathcal{N}\|$ by a last application of [32, Lem. 3.8]. ☐

4.3.2. *Invertible Affine Nets.* The restriction to invertible affine nets [7] is somehow orthogonal to the usual restrictions to reset/transfer/post self-modifying/... nets. For instance, in Example 4.6, the identity matrix in $\mathcal{N}_1$ is clearly invertible, and the transfer matrix in $\mathcal{N}_2$ is not. More generally, reset nets are never invertible (when they perform resets), and transfer nets are invertible exactly when their matrices are permutation matrices. Nevertheless, some more involved affine nets are invertible, like $\mathcal{N}_3$ in Example 4.6, whose matrix is invertible with inverse $\begin{bmatrix} 0 & 1/2 \\ 1 & -1/2 \end{bmatrix}$.

Strong Monotonicity. When dealing with a descending sequence of downwards-closed sets produced by the dual backward coverability algorithm for WSTS, a key observation made in [32] allows to sometimes derive monotonocity. For this, in a WSTS $(X, \to, \leq)$, define $\mathrm{Post}_\exists(S) \stackrel{\text{def}}{=} \{y \in X \mid \exists x \in S \,.\, x \to y\}$. Following [9], for two order ideals I and I', write $I \rightsquigarrow I'$ if I' appears in the canonical decomposition of $\downarrow\mathrm{Post}_\exists(I)$.

Fact 4.9 ([32, Claim 4.2])**.** *Let $D_0 \supsetneq D_1 \supsetneq \cdots$ be a descending chain of downwards-closed sets defined by equations (13–15). If I_{k+1} is a ideal proper at step $k+1$, then there exists an order ideal I and an order ideal I_k proper at step k such that $I_{k+1} \rightsquigarrow I \subseteq I_k$.*

In the case of affine nets, and identifying order ideals I with vectors in $\mathbb{N}^d_\omega$ with $\omega + n = \omega - n = \omega \cdot n = \omega$ for all n in $\mathbb{N}$, $\downarrow\mathrm{Post}_\exists(I) = \downarrow\{A \cdot (I - \boldsymbol{a}) + \boldsymbol{b} \mid (\boldsymbol{a}, A, \boldsymbol{b}) \in \mathcal{N}, I \sqsupseteq \boldsymbol{a}\}$.

Proposition 4.10. *The descending chain $D_0 \supsetneq D_1 \supsetneq \cdots$ defined by equations (13–15) for a d-dimensional invertible affine net $\mathcal{N}$ and a target vector $\boldsymbol{t}$ is strongly monotone.*

Proof. Let I_{k+1} be proper at step $k+1$. By Fact 4.9, there exists an order ideal I and an order ideal I_k proper at step k such that $I_{k+1} \rightsquigarrow_{\mathcal{N}} I \subseteq I_k$. Let us show that $\dim I_{k+1} \leq \dim I$; as $\dim I \leq \dim I_k$ because $I \subseteq I_k$, this will yield the result.

Since $I_{k+1} \rightsquigarrow_{\mathcal{N}} I$, there exists $(\boldsymbol{a}, A, \boldsymbol{b})$ in $\mathcal{N}$ such that $I - \boldsymbol{b} = A \cdot (I_{k+1} - \boldsymbol{a})$. For this to hold, note that for all $i \in \mathrm{fin}(I)$, the ith row of A must be such that $A(i,j) = 0$ for all $j \in \omega(I_{k+1})$. As A is invertible, those $(\mathrm{fdim}\, I)$-many rows must be linearly independent. As just argued, the jth column for each of these rows is made of zeroes whenever $j \in \omega(I_{k+1})$. Thus the remaining $(\mathrm{fdim}\, I_{k+1})$-many columns must make those $\mathrm{fdim}\, I$ rows linearly independent, hence necessarily $\mathrm{fdim}\, I_{k+1} \geq \mathrm{fdim}\, I$, i.e., $\dim I_{k+1} \leq \dim I$. ☐

Observe that the proof of Proposition 4.10 does not work for the transfer net $\mathcal{N}_2$ of Example 4.6: $\begin{bmatrix} \omega \\ \omega \end{bmatrix} \rightsquigarrow_{\mathcal{N}_2} \begin{bmatrix} \omega \\ 0 \end{bmatrix}$; this is exactly the kind of non-monotone behaviour invertibility was designed to prevent. Also observe that $\begin{bmatrix} 2 \\ \omega \end{bmatrix} \rightsquigarrow_{\mathcal{N}_3} \begin{bmatrix} \omega \\ 4 \end{bmatrix}$ in the invertible affine net $\mathcal{N}_3$, which is not an ω-monotone behaviour: this illustrates the usefulness of capturing strongly monotone descending chains, as [32, Thm. 4.4 and Cor. 4.6] do not apply.

Complexity Upper Bounds. We are now equipped to analyse the complexity of the backward coverability algorithm in invertible affine nets. Regarding the length ℓ of the chain constructed by the algorithm, by propositions 4.7 and 4.10 we are in the same situation as in Theorem 4.2 and we can simply repeat the arguments from its proof.

Theorem 4.11. *The backward coverability algorithm terminates after at most $n^{2^{O(d)}}$ iterations on d-dimensional invertible affine nets encoded in unary when $d \geq 2$.*

We deduce two corollaries from Theorem 4.11: one pertaining to the complexity of the backward coverability algorithm in dimension d, which mirrors Corollary 4.5, and one for the coverability problem when d is part of the input. Let us start with the backward coverability algorithm.

Corollary 4.12. *The backward coverability algorithm runs in time $n^{2^{O(d)}}$ on d-dimensional invertible affine nets encoded in unary when $d \geq 2$.*

Proof. Theorem 4.11 shows that the length ℓ of the ascending chain $U_0 \subsetneq U_1 \subsetneq \cdots \subsetneq U_\ell = U_*$ constructed by the backward coverability algorithm is at most $L_d + 1$, which is in $n^{2^{O(d)}}$.

Let $B_k \stackrel{\text{def}}{=} \min_\sqsubseteq U_k$ denote the minimal basis at step k. In order to compute B_{k+1} as per (13), thanks to Claim 4.8, we could essentially argue as in the proof of Corollary 4.5, with the caveat that computing bluntly $\min_\sqsubseteq \mathrm{Pre}_\exists(\uparrow\boldsymbol{v})$ for each $\boldsymbol{v} \in B_k$ is dangerously similar to a linear integer programming question and will incur an additional cost.

Alternatively, recall from equation (16) that $\mathrm{Fil}^{\mathsf{thin}+1}(\mathbb{N}^d)$, the set of order filters $\uparrow\boldsymbol{v}$ such that $\boldsymbol{v}$ is nearly thin, has at most $n^{2^{O(d)}}$ elements, and that $|B_k| \leq |\mathrm{Fil}^{\mathsf{thin}+1}(\mathbb{N}^d)|$ by Corollary 3.7 and Proposition 3.8. Thus in order to compute B_{k+1} one can enumerate the nearly thin vectors $\boldsymbol{v'} \in \mathrm{Fil}^{\mathsf{thin}+1}(\mathbb{N}^d)$ and check for each $(\boldsymbol{a}, A, \boldsymbol{b}) \in \mathcal{N}$ such that $\boldsymbol{v'} \sqsupseteq \boldsymbol{a}$ whether there exists $\boldsymbol{v} \in B_k$ such that $A \cdot (\boldsymbol{v'} - \boldsymbol{a}) + \boldsymbol{b} \sqsupseteq \boldsymbol{v}$. Each such check can be performed in time polynomial in $\|\boldsymbol{v'}\| \leq N_d + 1 = n \cdot (L_{d-1} + 2) + 1$, n, d, and $|B_k| \leq |\mathrm{Fil}^{\mathsf{thin}+1}(\mathbb{N}^d)|$. Thus the entire computation can be carried out in $n^{2^{O(d)}}$.

The same upper bound holds for the dual version of the backward coverability algorithm. At each step k, by Corollary 3.7, in order to compute D_{k+1} one can enumerate the thin order ideals $I \in \mathrm{Idl}^{\mathsf{thin}}(\mathbb{N}^d)$ and check for each such I whether $I \subseteq D_k$ and $I \subseteq \mathrm{Pre}_\forall(D_k)$, before removing the non-maximal ones. Note that $I \subseteq \mathrm{Pre}_\forall(D_k)$ if and only if $\mathrm{Post}_\exists(I) \subseteq D_k$, if and only if $\{A \cdot (I - \boldsymbol{a}) + \boldsymbol{b} \mid (\boldsymbol{a}, A, \boldsymbol{b}) \in \mathcal{N}, I \sqsupseteq \boldsymbol{a}\} \subseteq D_k$, which can be checked in time polynomial in $\|I\| \leq N_d = n \cdot (L_{d-1} + 2)$, n, d, and $|D_k| \leq |\mathrm{Idl}^{\mathsf{thin}}(\mathbb{N}^d)| \leq L_d + 1$ by Remark 3.3. The entire computation can be performed in time polynomial in L_d, n, and d, and this remains in $n^{2^{O(d)}}$. ∎

As VAS are a particular case of invertible affine nets, the upper bounds in Corollary 4.12 are optimal assuming the exponential time hypothesis by [30, Thm. 4.2].

Our last result concerns the complexity of coverability in invertible affine nets when d is part of the input. Note that the arguments leading to an algorithm working in space $O(d \lg(n \cdot \ell))$ in the VAS case [30, Thm. 3.2]—which are essentially the same as those used to derive a 2EXPSPACE upper bound for invertible affine nets in [7, Thm. 6]—do not work here, as the configurations along an execution of an affine net can grow exponentially with ℓ.

Corollary 4.13. *The coverability problem for invertible affine nets is* EXPSPACE*-complete.*

Proof. The hardness for EXPSPACE follows from the hardness of the coverability problem for VAS [34].

Regarding the upper bound, consider the execution of the classical backward coverability algorithm as defined in equation (13) on an invertible affine net $\mathcal{N}$

with target configuration $\boldsymbol{t}$: this is an ascending chain $U_0 \subsetneq U_1 \subsetneq \cdots \subsetneq U_\ell$ where $U_\ell = U_{\ell+1} = U_*$. The following characterisation of coverability actually holds more generally in WSTS.

Claim 4.14. *In an affine net $\mathcal{N}$, $\boldsymbol{s}$ covers $\boldsymbol{t}$ if and only if there exists $\ell' \leq \ell$ and a sequence of configurations $\boldsymbol{t}_0, \ldots, \boldsymbol{t}_{\ell'}$, called a* coverability pseudo-witness*, satisfying*

$$\boldsymbol{t}_0 \stackrel{\text{def}}{=} \boldsymbol{t}\ , \qquad \boldsymbol{t}_{k+1} \in \min_{\sqsubseteq} \mathrm{Pre}_\exists(\uparrow \boldsymbol{t}_k)\ , \qquad \boldsymbol{t}_{\ell'} \sqsubseteq \boldsymbol{s}\ . \tag{17}$$

Proof of Claim 4.14. If a coverability pseudo-witness exists, then we claim that for all $\ell' \geq k \geq 0$ there exists $\boldsymbol{s}_k \sqsupseteq \boldsymbol{t}_k$ such that $\boldsymbol{s} = \boldsymbol{s}_{\ell'} \to_{\mathcal{N}} \boldsymbol{s}_{\ell'-1} \to_{\mathcal{N}} \cdots \to_{\mathcal{N}} \boldsymbol{s}_k$, and thus in particular $\boldsymbol{s} \to^*_{\mathcal{N}} \boldsymbol{s}_0 \geq \boldsymbol{t}_0$ for $k = 0$. We can check this by induction over k. For the base case $k = \ell'$, define $\boldsymbol{s}_{\ell'} \stackrel{\text{def}}{=} \boldsymbol{s}$. For the induction step k, since $\boldsymbol{t}_{k+1} \in \mathrm{Pre}_\exists(\uparrow \boldsymbol{t}_k)$ there exists $\boldsymbol{s}'_k \sqsupseteq \boldsymbol{t}_k$ such that $\boldsymbol{t}_{k+1} \to_{\mathcal{N}} \boldsymbol{s}'_k$; by WSTS compatibility and since $\boldsymbol{s}_{k+1} \sqsupseteq \boldsymbol{t}_{k+1}$, there exists $\boldsymbol{s}_k \sqsupseteq \boldsymbol{s}'_k$ such that $\boldsymbol{s}_{k+1} \to_{\mathcal{N}} \boldsymbol{s}_k$.

Conversely, assume that $\boldsymbol{s}$ covers $\boldsymbol{t}$ in $\mathcal{N}$. Then $\boldsymbol{s} \in U_\ell$, and let $\ell' \leq \ell$ be the least index such that $\boldsymbol{s} \in U_{\ell'}$. Then either $\ell' = 0$, i.e., $\boldsymbol{s} \sqsupseteq \boldsymbol{t} = \boldsymbol{t}_0$ and we are done, or $\ell' > 0$. Because $\boldsymbol{s} \in U_{\ell'}$ there must be some $\boldsymbol{t}_{\ell'} \in \min_{\sqsubseteq} U_{\ell'}$ with $\boldsymbol{s} \sqsupseteq \boldsymbol{t}_{\ell'}$, and $\boldsymbol{t}_{\ell'} \notin U_{\ell'-1}$ as otherwise $\boldsymbol{s}$ would be in $U_{\ell'-1}$, contradicting the minimality of ℓ'. In general, if we have found a sequence $(\boldsymbol{t}_j)_{\ell' \geq j \geq k > 0}$ satisfying (17) until rank $k+1$ included and know that $\boldsymbol{t}_k \in (\min_{\sqsubseteq} U_k) \setminus U_{k-1}$, then either $k = 1$ and $\boldsymbol{t}_1 \in \min_{\sqsubseteq} \mathrm{Pre}_\exists(\uparrow \boldsymbol{t}_0)$ by definition of U_0 and U_1 in (13), or $k > 1$ and because $\boldsymbol{t}_k \notin U_{k-1}$, there exists $\boldsymbol{t}_{k-1} \in \min_{\sqsubseteq} U_{k-1}$ such that $\boldsymbol{t}_k \in \min_{\sqsubseteq} \mathrm{Pre}_\exists(\uparrow \boldsymbol{t}_{k-1})$, and $\boldsymbol{t}_{k-1} \notin U_{k-2}$ as otherwise we would have $\boldsymbol{t}_k$ in U_{k-1}. Repeating this process yields a coverability pseudo-witness. [4.14]

By Claim 4.14, a non-deterministic algorithm for coverability can guess and check the existence of a coverability pseudo-witness. By Theorem 4.11, such a pseudo-witness has a length $\ell' \leq \ell$ in $n^{2^{O(d)}}$. Furthermore, by Claim 4.8 the components in each $\boldsymbol{t}_k$ in such a pseudo-witness are bounded by $\|\boldsymbol{t}\| + \|\mathcal{N}\| \cdot k \leq (\ell+1) \cdot n$, which is still in $n^{2^{O(d)}}$. Thus exponential space suffices. Note that this also holds when we assume the invertible affine net to be encoded in binary, by substituting 2^n for n in the bound $n^{2^{O(d)}}$. $\square$

Remark 4.15 (Strictly increasing affine nets)**.** Strictly increasing affine nets [43, 23, 12] are intuitively the affine nets devoid of any form of reset or transfer; in Example 4.6, only $\mathcal{N}_1$ is strictly increasing. All the results we have proven for invertible affine nets in this section—namely in Theorem 4.11 and corollaries 4.12 and 4.13—also hold for strictly increasing affine nets, because the descending chains of downwards-closed sets they generate when running the backward coverability algorithm are ω-monotone.

Claim 4.16. *The descending chain $D_0 \supsetneq D_1 \supsetneq \cdots$ defined by equations (13–15) for a d-dimensional strictly increasing affine net $\mathcal{N}$ and a target vector $\boldsymbol{t}$ is ω-monotone.*

Proof of Claim 4.16. Let I_{k+1} be proper at step $k+1$. By Fact 4.9, there exists an order ideal I and an order ideal I_k proper at step k such that $I_{k+1} \rightsquigarrow_{\mathcal{N}} I \subseteq I_k$. Let us show that $\omega(I_{k+1}) \subseteq \omega(I)$; as $\omega(I) \subseteq \omega(I_k)$ because $I \subseteq I_k$, this will yield the result.

Since $I_{k+1} \rightsquigarrow_{\mathcal{N}} I$, there exists $(\boldsymbol{a}, A, \boldsymbol{b})$ in $\mathcal{N}$ such that $I_{k+1} \sqsupseteq \boldsymbol{a}$ and $I = A \cdot (I_{k+1} - \boldsymbol{a}) + \boldsymbol{b}$. Because $\mathcal{N}$ is strictly increasing, $A = I_d + A'$ for some matrix

$A' \in \mathbb{N}^{d\times d}$, hence $I = I_{k+1} - \boldsymbol{a} + A' \cdot (I_{k+1} - \boldsymbol{a}) + \boldsymbol{b}$. Thus $I \sqsupseteq (I_{k+1} - \boldsymbol{a})$ and therefore $\omega(I) \supseteq \omega(I_{k+1})$. [4.16]

An EXPSPACE upper bound was already shown by Bonnet et al. [12] for the coverability problem, but the $n^{2^{O(d)}}$ bound for the problem parameterised by d is an improvement over the $n^{2^{O(d \lg d)}}$ bounds of [12, Lem. 11 and Thm. 13], and the bounds for the backward coverability algorithm are new. □

References

[1] P. A. Abdulla, K. Čerāns, B. Jonsson, and Y.-K. Tsay. Algorithmic analysis of programs with well quasi-ordered domains. *Inform. and Comput.*, 160(1–2):109–127, 2000. doi:10.1006/inco.1999.2843.

[2] T. Araki and T. Kasami. Some decision problems related to the reachability problem for Petri nets. *Theor. Comput. Sci.*, 3(1):85–104, 1976. doi:10.1016/0304-3975(76)90067-0.

[3] A. Arnold and M. Latteux. Récursivité et cônes rationnels fermés par intersection. *CALCOLO*, 15(4):381–394, 1978. doi:10.1007/BF02576519.

[4] A. R. Balasubramanian. Complexity of controlled bad sequences over finite sets of $\mathbb{N}^d$. In *Proc. LICS 2020*, pages 130–140. ACM, 2020. doi:10.1145/3373718.3394753.

[5] A. R. Balasubramanian. Complexity of coverability in depth-bounded processes. In *Proc. CONCUR 2022*, volume 243 of *Leibniz Int. Proc. Inf.*, pages 17:1–17:19. LZI, 2022. doi:10.4230/LIPIcs.CONCUR.2022.17.

[6] A. R. Balasubramanian, T. Lang, and R. Ramanayake. Decidability and complexity in weakening and contraction hypersequent substructural logics. In *Proc. LICS 2021*, pages 1–13. IEEE, 2021. doi:10.1109/LICS52264.2021.9470733.

[7] M. Benedikt, T. Duff, A. Sharad, and J. Worrell. Polynomial automata: Zeroness and applications. In *Proc. LICS 2017*, pages 1–12. IEEE, 2017. doi:10.1109/LICS.2017.8005101.

[8] M. Blondin, A. Finkel, C. Haase, and S. Haddad. Approaching the coverability problem continuously. In *Proc. TACAS 2016*, volume 9636 of *Lect. Notes in Comput. Sci.*, pages 480–496. Springer, 2016. doi:10.1007/978-3-662-49674-9_28.

[9] M. Blondin, A. Finkel, and P. McKenzie. Handling infinitely branching well-structured transition systems. *Inform. and Comput.*, 258:28–49, 2018. doi:10.1016/j.ic.2017.11.001.

[10] M. Blondin, C. Haase, and P. Offtermatt. Directed reachability for infinite-state systems. In *Proc. TACAS 2021*, volume 12652 of *Lect. Notes in Comput. Sci.*, pages 3–23. Springer, 2021. doi:10.1007/978-3-030-72013-1_1.

[11] R. Bonnet. On the cardinality of the set of initial intervals of a partially ordered set. In *Infinite and finite sets: to Paul Erdős on his 60th birthday, Vol. 1*, volume 10 of *Coll. Math. Soc. János Bolyai*, pages 189–198. North-Holland, 1975.

[12] R. Bonnet, A. Finkel, and M. Praveen. Extending the Rackoff technique to affine nets. In *Proc. FSTTCS 2012*, volume 18 of *Leibniz Int. Proc. Inf.*, pages 301–312. LZI, 2012. doi:10.4230/LIPIcs.FSTTCS.2012.301.

[13] L. Bozzelli and P. Ganty. Complexity analysis of the backward coverability algorithm for VASS. In *Proc. RP 2011*, volume 6945 of *Lect. Notes in Comput. Sci.*, pages 96–109. Springer, 2011. doi:10.1007/978-3-642-24288-5_10.

[14] G. Ciardo. Petri nets with marking-dependent arc cardinality: Properties and analysis. In *Proc. Petri Nets 1994*, volume 815 of *Lect. Notes in Comput. Sci.*, pages 179–198. Springer, 1994. doi:10.1007/3-540-58152-9_11.

[15] J.-B. Courtois and S. Schmitz. Alternating vector addition systems with states. In *Proc. MFCS 2014*, volume 8634 of *Lect. Notes in Comput. Sci.*, pages 220–231. Springer, 2014. doi:10.1007/978-3-662-44522-8_19.

[16] S. Demri, M. Jurdziński, O. Lachish, and R. Lazić. The covering and boundedness problems for branching vector addition systems. *J. Comput. Syst. Sci.*, 79(1):23–38, 2013. doi:10.1016/j.jcss.2012.04.002.

[17] L. E. Dickson. Finiteness of the odd perfect and primitive abundant numbers with n distinct prime factors. *Amer. J. Math.*, 35(4):413–422, 1913. doi:10.2307/2370405.

[18] C. Dufourd, A. Finkel, and Ph. Schnoebelen. Reset nets between decidability and undecidability. In *Proc. ICALP 1998*, volume 1443 of *Lect. Notes in Comput. Sci.*, pages 103–115. Springer, 1998. doi:10.1007/BFb0055044.

[19] J. Esparza, R. Ledesma-Garza, R. Majumdar, P. Meyer, and F. Niksic. An SMT-based approach to coverability analysis. In *Proc. CAV 2014*, volume 8559 of *Lect. Notes in Comput. Sci.*, pages 603–619. Springer, 2014. doi:10.1007/978-3-319-08867-9_40.

[20] D. Figueira, S. Figueira, S. Schmitz, and Ph. Schnoebelen. Ackermannian and primitive-recursive bounds with Dickson's Lemma. In *Proc. LICS 2011*, pages 269–278. IEEE, 2011. doi:10.1109/LICS.2011.39.

[21] A. Finkel. A generalization of the procedure of Karp and Miller to well structured transition systems. In *Proc. ICALP 1987*, volume 267 of *Lect. Notes in Comput. Sci.*, pages 499–508. Springer, 1987. doi:10.1007/3-540-18088-5_43.

[22] A. Finkel and Ph. Schnoebelen. Well-structured transition systems everywhere! *Theor. Comput. Sci.*, 256(1–2):63–92, 2001. doi:10.1016/S0304-3975(00)00102-X.

[23] A. Finkel, P. McKenzie, and C. Picaronny. A well-structured framework for analysing Petri net extensions. *Inform. and Comput.*, 195(1):1–29, 2004. doi:10.1016/j.ic.2004.01.005.

[24] T. Geffroy, J. Leroux, and G. Sutre. Occam's Razor applied to the Petri net coverability problem. *Theor. Comput. Sci.*, 750:38–52, 2018. doi:10.1016/j.tcs.2018.04.014.

[25] J. Goubault-Larrecq, S. Halfon, P. Karandikar, K. Narayan Kumar, and Ph. Schnoebelen. The ideal approach to computing closed subsets in well-quasi-orderings. In *Well-Quasi-Orders in Computation, Logic, Language and Reasoning*, volume 53 of *Trends in Logic*, pages 55–105. Springer, 2020. doi:10.1007/978-3-030-30229-0_3.

[26] L. Guillou, C. Mascle, and N. Waldburger. Parameterized broadcast networks with registers: from NP to the frontiers of decidability. In *Proc. FoSSaCS 2024*, volume 14575 of *Lect. Notes in Comput. Sci.*, pages 250–270, 2024. doi:10.1007/978-3-031-57231-9_12.

[27] R. M. Karp and R. E. Miller. Parallel program schemata. *J. Comput. Syst. Sci.*, 3(2):147–195, 1969. doi:10.1016/S0022-0000(69)80011-5.

[28] U. Koppenhagen and E. W. Mayr. Optimal algorithms for the coverability, the subword, the containment, and the equivalence problems for commutative semigroups. *Inform. and Comput.*, 158(2):98–124, 2000. doi:10.1006/inco.1999.2812.

[29] J. B. Kruskal. The theory of well-quasi-ordering: A frequently discovered concept. *J. Comb. Theory A*, 13(3):297–305, 1972. doi:10.1016/0097-3165(72)90063-5.

[30] M. Künnemann, F. Mazowiecki, L. Schütze, H. Sinclair-Banks, and K. Węgrzycki. Coverability in VASS revisited: Improving Rackoff's bound to obtain conditional optimality. In *Proc. ICALP 2023*, Leibniz Int. Proc. Inf., pages 131:1–131:20. LZI, 2023. doi:10.4230/LIPIcs.ICALP.2023.131.

[31] R. Lazić and S. Schmitz. The complexity of coverability in ν-Petri nets. In *Proc. LICS 2016*, pages 467–476. ACM, 2016. doi:10.1145/2933575.2933593.

[32] R. Lazić and S. Schmitz. The ideal view on Rackoff's coverability technique. *Inform. and Comput.*, 277:104582, 2021. doi:10.1016/j.ic.2020.104582.

[33] J. Leroux. Vector addition system reversible reachability problem. *Logic. Meth. in Comput. Sci.*, 9(1), 2013. doi:10.2168/LMCS-9(1:5)2013.

[34] R. J. Lipton. The reachability problem requires exponential space. Technical Report 62, Department of Computer Science, Yale University, 1976. URL http://www.cs.yale.edu/publications/techreports/tr63.pdf.

[35] E. W. Mayr and A. R. Meyer. The complexity of the word problems for commutative semigroups and polynomial ideals. *Adv. Math.*, 46(3):305–329, 1982. doi:10.1016/0001-8708(82)90048-2.

[36] D. Novikov and S. Yakovenko. Trajectories of polynomial vector fields and ascending chains of polynomial ideals. *Ann. Inst. Fourier*, 49(2):563–609, 1999. doi:10.5802/aif.1683.

[37] C. Rackoff. The covering and boundedness problems for vector addition systems. *Theor. Comput. Sci.*, 6(2):223–231, 1978. doi:10.1016/0304-3975(78)90036-1.

[38] L. E. Rosier and H.-C. Yen. A multiparameter analysis of the boundedness problem for vector addition systems. *J. Comput. Syst. Sci.*, 32(1):105–135, 1986. doi:10.1016/0022-0000(86)90006-1.

[39] S. Schmitz. *Algorithmic Complexity of Well-Quasi-Orders*. Habilitation thesis, École Normale Supérieure Paris-Saclay, 2017. URL http://tel.archives-ouvertes.fr/tel-01663266.

[40] S. Schmitz. The parametric complexity of lossy counter machines. In *Proc. ICALP 2019*, volume 132 of *Leibniz Int. Proc. Inf.*, pages 129:1–129:15. LZI, 2019. doi:10.4230/LIPIcs.ICALP.2019.129.

[41] S. Schmitz and Ph. Schnoebelen. Algorithmic aspects of WQO theory. Lecture notes, 2012. URL http://cel.archives-ouvertes.fr/cel-00727025.

[42] Ph. Schnoebelen. Revisiting Ackermann-hardness for lossy counter machines and reset Petri nets. In *Proc. MFCS 2010*, volume 6281 of *Lect. Notes in Comput. Sci.*, pages 616–628. Springer, 2010. doi:10.1007/978-3-642-15155-2_54.

[43] R. Valk. Self-modifying nets, a natural extension of Petri nets. In *Proc. ICALP 1978*, volume 62 of *Lect. Notes in Comput. Sci.*, pages 464–476. Springer, 1978. doi:10.1007/3-540-08860-1_35.

[44] H.-C. Yen and C.-L. Chen. On minimal elements of upward-closed sets. *Theor. Comput. Sci.*, 410(24):2442–2452, 2009. doi:10.1016/j.tcs.2009.02.036.