



HAL
open science

The Ideal View on Rackoff's Coverability Technique

Ranko Lazić, Sylvain Schmitz

► **To cite this version:**

Ranko Lazić, Sylvain Schmitz. The Ideal View on Rackoff's Coverability Technique. RP 2015 - 9th International Workshop on Reachability Problems, Sep 2015, Warsaw, Poland. pp.1–13, 10.1007/978-3-319-24537-9_8 . hal-01176755v2

HAL Id: hal-01176755

<https://inria.hal.science/hal-01176755v2>

Submitted on 2 Feb 2016 (v2), last revised 21 Jun 2017 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Copyright

THE IDEAL VIEW ON RACKOFF'S COVERABILITY TECHNIQUE

RANKO LAZIĆ AND SYLVAIN SCHMITZ

ABSTRACT. Rackoff's small witness property for the coverability problem is the standard means to prove tight upper bounds in vector addition systems (VAS) and many extensions. We show how to derive the same bounds directly on the computations of the VAS instantiation of the generic backward coverability algorithm. This relies on a dual view of the algorithm using ideal decompositions of downwards-closed sets, which exhibits a key structural invariant in the VAS case. This reasoning generalises uniformly to several VAS extensions: we derive optimal upper bounds for coverability in branching and alternating VAS, matching the previously known results from the literature.

1. INTRODUCTION

Checking safety properties in infinite transition systems can often be reduced to *coverability* checks. The coverability problem asks, given a transition system and two configurations x and y and a quasi-ordering \leq over configurations, whether x might *cover* y , i.e. reach some configuration $y' \geq y$ in finitely many steps. The problem is decidable for the large class of (effective) *well-structured transition systems* (WSTS) where \leq is a *well-quasi-ordering* (wqo) compatible with the transition relation [1, 13]. The algorithm to that end is a generic *backward coverability* procedure, which computes successively the sets of configurations that can cover y in at most $0, 1, 2, \dots$ steps. Those sets are upwards-closed and since \leq is a wqo they can be represented through their finitely many minimal elements.

Nevertheless, the naive complexity upper bounds one can extract directly from the termination argument of the backward coverability algorithm—which also relies on \leq being a wqo—are sometimes very far from the optimal ones. A striking illustration is provided by vector addition systems (VAS): the complexity bounds offered e.g. by [11] are in ACKERMANN, whereas coverability in VAS has long been known to be EXPSPACE-complete thanks to a lower bound by Lipton [20] and an upper bound by Rackoff [22].

Rackoff's Technique is essentially combinatorial in nature: he shows by induction on the dimension of the VAS that, if x can reach one such $y' \geq y$, then there exists a small (doubly-exponential) run in the VAS witnessing this fact. A non-deterministic algorithm can then simply look for such a

An extended abstract of this work was published in the Proceedings of the 9th International Workshop on Reachability Problems, Lecture Notes in Computer Science vol. 9328 pp. 76–88 (M. Bojańczyk, S. Lasota, and I. Potapov, Eds.), Springer, 2015. Work funded in part by the Leverhulme Trust Visiting Professorship VP1-2014-041 and the EPSRC grant EP/M011801/1.

witness using only exponential space. The same general technique has since been extended to prove tight complexity upper bounds for coverability in numerous extensions of VASs [10, 6, 8, 17, 16]. It is however less clear how to adapt the technique for more general systems, where for instance the notion of dimension is absent or more involved.

Remarkably, Bozzelli and Ganty [7] showed that Rackoff’s small witness property can be applied to the backward coverability algorithm for VAS to obtain a 2EXPTIME upper bound.¹ However, their proof uses Rackoff’s analysis as a black box, and does not work directly with the structures manipulated by the backward coverability algorithm. As such, it is again unclear how this result could be translated to further classes of well-structured transition systems.

Contributions. In this paper, we revisit the backward coverability algorithm for VAS, and extract directly a 2EXPTIME upper bound for its running time. We take for this in Section 3 a dual view on the backward coverability algorithm, by considering successively the sets of configurations that do *not* cover y in 0, 1, 2, \dots or fewer steps. Such sets are downwards-closed, and enjoy a (usually effective) canonical representation as finite unions of *ideals* [5, 12, 15]. We show in Section 4 that, in the case of VAS, this dual view exhibits an additional structural property of ω -*monotonicity*, which allows to derive the desired doubly-exponential bound.

Our purpose is above all pedagogical, as we hope to see this type of reasoning applied more broadly where the simple proof argument of Rackoff fails. As illustrations of the versatility of the framework, we consider in Section 5 and Section 6 the top-down and bottom-up coverability problems in *alternating branching* VAS. In each case, we provide an instance of the generic backward algorithm that solves the problem, and show that its running time matches the known optimal complexities [10, 8, 17]. In contrast to the upper bound proofs in the literature, which employ ad-hoc adaptations of Rackoff’s technique, the proofs we obtain are uniform extensions from the basic case of VAS, so they proceed by analysing the downwards-closed sets that the backward algorithms produce. For top-down coverability, the ideals involved are the same as for VAS (vector ideals), whereas bottom-up coverability requires us to work one level higher (vector set ideals).

We start with some preliminaries on WSTS and ideals in Section 2.

2. PRELIMINARIES

We first recall the necessary background on well-quasi-orders, well-structured transition systems, and ideal decompositions, while illustrating systematically the definitions on VAS and reset VAS.

2.1. Well-Structured Transition Systems. A *well-quasi-order* (wqo) (X, \leq) is a set X equipped with a transitive reflexive relation \leq such that, along any infinite sequence x_0, x_1, \dots of elements from X , one can find two indices $i < j$ such that $x_i \leq x_j$. A finite or infinite sequence without such

¹In the same spirit, Majumdar and Wang [21] show that the ‘expand, enlarge, and check’ algorithm for bottom-up coverability in branching VASs runs in 2EXPTIME , using the combinatorial analysis of Demri et al. [10].

pair of indices is *bad*, and necessarily finite over a wqo. See for instance [26] for more background on wqos.

Example 2.1 (Dickson's Lemma). The set \mathbb{N}^d of d -dimensional vectors of natural numbers forms a wqo when endowed with the product ordering \sqsubseteq , defined by $\mathbf{u} \sqsubseteq \mathbf{v}$ if $\mathbf{u}(i) \leq \mathbf{v}(i)$ for all $1 \leq i \leq d$.

A *well-structured transition system* (WSTS) [1, 13] is a triple (X, \rightarrow, \leq) where X is a set of configurations, $\rightarrow \subseteq X \times X$ is a transition relation, and (X, \leq) is a wqo with the following *compatibility* condition: if $x \leq x'$ and $x \rightarrow y$, then there exists $y' \geq y$ with $x' \rightarrow y'$. In other words, \leq is a simulation relation on the transition system (X, \rightarrow) . We write as usual $\rightarrow^{\leq 0} \stackrel{\text{def}}{=} \{(x, x) \mid x \in X\}$ and $\rightarrow^{\leq k+1} \stackrel{\text{def}}{=} \rightarrow^{\leq k} \cup \{(x, y) \mid \exists z \in X. x \rightarrow z \rightarrow^{\leq k} y\}$ for the reachability relation in at most k steps, and $\rightarrow^* \stackrel{\text{def}}{=} \bigcup_k \rightarrow^{\leq k}$ for the reflexive transitive closure of \rightarrow .

Example 2.2 (VAS are WSTS). A d -dimensional *vector addition system* (VAS) is a finite set \mathbf{A} of vectors in \mathbb{Z}^d . It defines a WSTS $(\mathbb{N}^d, \rightarrow, \sqsubseteq)$ with configurations space \mathbb{N}^d and $\mathbf{u} \rightarrow \mathbf{u} + \mathbf{a}$ for all \mathbf{u} in \mathbb{N}^d and \mathbf{a} in \mathbf{A} such that $\mathbf{u} + \mathbf{a}$ is in \mathbb{N}^d .

For instance, the 2-dimensional VAS $\mathbf{A}_{\div 2} = \{(-2, 1)\}$ can be seen as weakly computing the halving function: from any configuration $(n, 0)$, it can reach $(n \bmod 2, \lfloor n/2 \rfloor)$ and all its reachable configurations (n', m) satisfy $m \leq n/2$.

Example 2.3 (Reset VAS are WSTS). A d -dimensional *reset VAS* is a finite subset \mathbf{A} of $\mathbb{Z}^d \times \mathbb{P}(\{1, \dots, d\})$. Given $R \subseteq \{1, \dots, d\}$ and a vector \mathbf{u} , we define the vector $R(\mathbf{u})$ by $R(\mathbf{u})(i) = 0$ if $i \in R$, and $R(\mathbf{u})(i) = \mathbf{u}(i)$ otherwise. A reset VAS defines a WSTS $(\mathbb{N}^d, \rightarrow, \sqsubseteq)$ where $\mathbf{u} \rightarrow R(\mathbf{u} + \mathbf{a})$ if there exists (\mathbf{a}, R) in \mathbf{A} such that $\mathbf{u} + \mathbf{a}$ is in \mathbb{N}^d .

For instance, the 5-dimensional reset VAS

$$\mathbf{A}_{\log} = \left\{ \begin{array}{l} (0, 0, -2, 1, 0, \emptyset), (0, 0, 1, -1, 0, \emptyset), \\ (-1, 1, -2, 1, 0, \{3\}), (1, -1, 1, -1, 1, \{4\}) \end{array} \right\}$$

is a weak computer for the logarithm function: from any configuration of the form $(1, 0, 2^n, 0, 0)$, it can reach $(1, 0, 1, 0, n)$, and all its reachable configurations of the form $(1, 0, n', m, \ell)$ satisfy $\ell \leq n$.

2.2. Ideal Decompositions. The *downward-closure* of a subset $S \subseteq X$ over a wqo (X, \leq) is $\downarrow X \stackrel{\text{def}}{=} \{x \in X \mid \exists s \in S. x \leq s\}$. A subset $D \subseteq X$ is *downwards-closed* if $\downarrow D = D$. We write $\downarrow x$ for the downward-closure of the singleton set $\{x\}$. Upward-closures and upwards-closed sets $U = \uparrow U$ are defined similarly.

Well-quasi-orders can also be characterised by the *descending chain condition*: a quasi-order (X, \leq) is a wqo if and only if every descending sequence $D_0 \supsetneq D_1 \supsetneq D_2 \supsetneq \dots$ of downwards-closed subsets $D_i \subseteq X$ is finite.

An *ideal* of X is a non-empty downwards-closed subset $I \subseteq X$, which is *directed*: if x, x' are two elements of I , then there exists y in I with $x \leq y$ and $x' \leq y$. Over a wqo (X, \leq) , any downwards-closed set $D \subseteq X$ has a unique *decomposition* as a finite union of ideals $D = I_1 \cup \dots \cup I_n$, where the I_j 's are mutually incomparable for inclusion [5, 12]. Alternatively, ideals are

characterised as *irreducible* downwards-closed sets: an 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 $I \subseteq D_2$.

Example 2.4 (Vector Ideals). Over $(\mathbb{N}^d, \sqsubseteq)$, observe that $\downarrow \mathbf{u}$ is an ideal for every \mathbf{u} in \mathbb{N}^d . Those are however not the only ideals, e.g. $I \stackrel{\text{def}}{=} \{(0, n, 0) \mid n \in \mathbb{N}\}$ is also an ideal of \mathbb{N}^3 . Write $\mathbb{N}_\omega \stackrel{\text{def}}{=} \mathbb{N} \uplus \{\omega\}$ where ω is a new top element; the product ordering \sqsubseteq extends naturally to \mathbb{N}_ω^d . Then the ideals of $(\mathbb{N}^d, \sqsubseteq)$ are exactly the downward-closures $\downarrow \mathbf{u}$ inside \mathbb{N}^d of vectors \mathbf{u} from \mathbb{N}_ω^d . For the previous example in \mathbb{N}^3 , $\downarrow(0, \omega, 0) = I$.

Although ideals provide finite representations for manipulating downwards-closed sets, some additional effectiveness assumptions are necessary to employ them in algorithms. In this paper, we will say that a wqo (X, \leq) has *effective* ideal representations [see 12, 15, for more stringent requisites] if every ideal can be represented, and there are algorithms on those representations:

- (CI): to check $I \subseteq I'$ for two ideals I and I' ,
- (II): to compute the ideal decomposition of $I \cap I'$ for two ideals I and I' ,
- (CU'): to compute the ideal decomposition of the residual $X \setminus \uparrow x = \{x' \in X \mid x \not\leq x'\}$ for any x in X .

Example 2.5 (Effective Representations of Vector Ideals). We shall use vectors in \mathbb{N}_ω^d as representations. For (CI), given two vectors \mathbf{u} and \mathbf{v} in \mathbb{N}_ω^d , $\downarrow \mathbf{u} \subseteq \downarrow \mathbf{v}$ if and only if $\mathbf{u} \sqsubseteq \mathbf{v}$. Furthermore, for (II), $\downarrow \mathbf{u} \cap \downarrow \mathbf{v} = \downarrow \mathbf{w}$ where $\mathbf{w}(i) \stackrel{\text{def}}{=} \min_{\leq}(\mathbf{u}(i), \mathbf{v}(i))$ for all $1 \leq i \leq d$. Finally, for (CU'), if \mathbf{u} is in \mathbb{N}^d , then $\mathbb{N}^d \setminus \uparrow \mathbf{u} = \bigcup_{1 \leq j \leq d \mid \mathbf{u}(j) > 0} \downarrow \mathbf{u}_{/j}$ where $\mathbf{u}_{/j}(i) = \omega$ if $i \neq j$ and $\mathbf{u}_{/j}(j) \stackrel{\text{def}}{=} \mathbf{u}(j) - 1$ otherwise.

Crucially for the applicability of our approach, effective ideal representations exist for most wqos of interest [12, 15].

3. BACKWARD COVERABILITY

Let us recall in this section the generic backward coverability algorithm for well-structured transition systems [1, 13]—the first published instance we know of this algorithm was incidentally for the case of reset VASS [3].

We take a dual view on this algorithm, by considering downwards-closed sets represented through their ideal decompositions, instead of the usual view using upwards-closed sets represented through their minimal elements. We present the generic algorithm, but will illustrate all the notions using the case of VAS and reset VAS in Section 3.2, and derive naive upper bounds for both in Section 3.3—which will turn out to be optimal for reset VAS.

3.1. Generic Algorithm. Consider a WSTS (X, \rightarrow, \leq) and a target configuration y from X to be covered. Define $D_* \stackrel{\text{def}}{=} \{x \in X \mid \forall y' \geq y. x \not\rightarrow^* y'\}$ as the set of configurations that do not cover y . The purpose of the backward coverability algorithm is to compute D_* ; solving a coverability instance with source configuration x_0 then amounts to checking whether x_0 belongs to D_* .

The idea of the algorithm is to compute successively for every k the set D_k of configurations that do *not* cover y in k steps or fewer:

$$D_* = \bigcap_k D_k, \quad D_k \stackrel{\text{def}}{=} \{x \in X \mid \forall y' \geq y. x \not\rightarrow^{\leq k} y'\}. \quad (1)$$

These over-approximations D_k can be computed inductively on k by

$$D_0 = X/y, \quad D_{k+1} = D_k \cap \text{Pre}_\forall(D_k), \quad (2)$$

where for any set $S \subseteq X$,

$$\text{Pre}_\forall(S) \stackrel{\text{def}}{=} \{x \in X \mid \forall z \in X. (x \rightarrow z \implies z \in S)\}. \quad (3)$$

The algorithm terminates as soon as $D_k \subseteq D_{k+1}$ (and thus $D_{k+j} = D_k = D_*$ for all j). This is guaranteed to arise eventually by the descending chain condition, since otherwise we would have an infinite descending chain of downwards-closed sets $D_0 \supseteq D_1 \supseteq D_2 \supseteq \dots$.

3.1.1. Correctness. The correctness of the algorithm hinges on the following claim:

Claim 3.1 (Correctness). Equations (1) and (2) define the same D_k .

Proof. By induction on k . For the base case, $x \not\rightarrow^{\leq 0} y'$ for all $y' \geq y$, if and only if $x \not\geq y$, i.e. if and only if x is in X/y . For the induction step and for all $y' \geq y$, $x \not\rightarrow^{\leq k+1} y'$ if and only if $x \not\rightarrow^{\leq k} y'$ and there does not exist any z with $x \rightarrow z$ and $z \rightarrow^{\leq k} y'$. The former is equivalent to x belonging to D_k by induction hypothesis. The latter occurs if and only if for all z in X , if $x \rightarrow z$ then $z \not\rightarrow^{\leq k} y'$, i.e. if and only if x belongs to $\text{Pre}_\forall(D_k)$ by induction hypothesis. \square

3.1.2. Effective Ideal Representations. The algorithm as presented above relies on the effectiveness of Eq. (2). We are going to use effective representations of the ideal decompositions of the D_k to this end. Let us first check that we are indeed dealing with downwards-closed sets:

Claim 3.2 (Downward-closure). For all k , D_k is downwards-closed.

Proof. By induction on k . For the base case, $D_0 = X/y$ is downwards-closed. For the induction step, first observe that, if D is downwards-closed, then $\text{Pre}_\forall(D)$ is also downwards-closed. Indeed, let $x \leq x'$ for some x' in $\text{Pre}_\forall(D)$. Consider any z such that $x \rightarrow z$. Then by WSTS compatibility, there exists $z' \geq z$ such that $x' \rightarrow z'$. Since x' belongs to $\text{Pre}_\forall(D)$, z' belongs to D . Because D is downwards-closed, z also belongs to D . This shows x in $\text{Pre}_\forall(D)$ as desired. We conclude by noting that downwards-closed sets are closed under intersection, hence $D_{k+1} = D_k \cap \text{Pre}_\forall(D_k)$ is downwards-closed by applying the induction hypothesis to D_k . \square

The only additional effectiveness assumption we make is that:

(Pre): the ideal decomposition of $\text{Pre}_\forall(D)$ is computable for all downwards-closed D .

This is sufficient to compute the ideal decompositions of all the D_k . Indeed, initially D_0 is computed using (CU'). Later, $\text{Pre}_\forall(D_k)$ is computable by (Pre), and its intersection with D_k is also computable by (II). Finally, recall that, by ideal irreducibility, $I_1 \cup \dots \cup I_n \subseteq J_1 \cup \dots \cup J_m$ for ideals I_1, \dots, I_n and downwards-closed J_1, \dots, J_m if and only if for all $1 \leq i \leq n$ there exists $1 \leq j \leq m$ such that $I_i \subseteq J_j$. Therefore, the termination check $D_k \subseteq D_{k+1}$ is effective by (CI).

Remark 3.3 (Existential Predecessors). The classical presentation of the backward coverability algorithm works with upwards-closed sets U , which are represented through their finitely many minimal elements in $\min U$. The assumption (Pre) is therefore replaced by the computability of a set of minimal elements for

$$\text{Pre}_\exists(U) \stackrel{\text{def}}{=} \{x \in X \mid \exists z \in U. x \rightarrow z\} \quad (4)$$

for an upwards-closed U .

It turns out that, under the usual effectiveness assumptions on (X, \leq_X) plus the ideal effectiveness assumptions, the two assumptions on the computability of Pre_\forall and Pre_\exists are equivalent. Indeed, we have the duality

$$\text{Pre}_\forall(D) = X \setminus (\text{Pre}_\exists(X \setminus D)), \quad (5)$$

where an ideal decomposition of $X \setminus U$ for an upwards-closed $U = \bigcup_{x \in \min U} \uparrow x$ can be computed since $X \setminus U = \bigcap_{x \in \min U} X \setminus \uparrow x$, while the minimal elements of $X \setminus D$ can be computed using the *Generalised Valk-Jantzen Lemma* [14].²


Hence, assuming effective ideal representations, our view of the backward coverability algorithm using downwards-closed sets is effective in the same cases as the usual upwards-closed version of the algorithm.

3.2. Coverability for VAS and Reset VAS. In order to instantiate the backward coverability algorithm for VAS and reset VAS, we need to prove that they satisfy the various effectiveness assumptions. Example 2.5 dealt with (CI), (II), and (CU'). By Remark 3.3, (Pre) is also effective since existential predecessors are well-known to be computable [1, 13].

We nevertheless provide here a direct computation of universal predecessors. Thus, given a downwards-closed $D = \downarrow \mathbf{u}_1 \cup \dots \cup \downarrow \mathbf{u}_m$ for some $\mathbf{u}_1, \dots, \mathbf{u}_m$ in \mathbb{N}_ω^d , we want to compute a finite set of vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$ from \mathbb{N}_ω^d such that $\text{Pre}_\forall(D) = \downarrow \mathbf{v}_1 \cup \dots \cup \downarrow \mathbf{v}_n$. Using (CI) we can then select the maximal such \mathbf{v}_j to obtain incomparable ideals.

3.2.1. Universal Predecessors in VAS. Thanks to (II) and the fact that \mathbf{A} is finite (VAS are finitely branching), we start by reducing our computation to that of predecessors along a specific action \mathbf{a} from \mathbf{A} : $\text{Pre}_\forall(D) =$

²This last result is quite general, but it might be simpler in practice to implement this computation directly; for instance, in the case of $X = \mathbb{N}^d$, $\mathbb{N}^d \setminus \downarrow \mathbf{u}$ for $\mathbf{u} \in \mathbb{N}_\omega^d$ is decomposed as $\bigcup_{1 \leq j \leq d \mid \mathbf{u}(j) < \omega} \uparrow \mathbf{u}_{\setminus j}$ where $\mathbf{u}_{\setminus j}(i) = 0$ if $i \neq j$ and $\mathbf{u}_{\setminus j}(j) = \mathbf{u}(j) + 1$ otherwise, and $\uparrow \mathbf{u} \cap \uparrow \mathbf{v} = \uparrow \mathbf{w}$ where $\mathbf{w}(i) = \max_{\leq}(\mathbf{u}(i), \mathbf{v}(i))$ for all $1 \leq i \leq d$.

FIGURE 1. The successive D_k for $\mathbf{A}_{\div 2}$ with target $\mathbf{t} = (0, 5)$.

$\bigcap_{\mathbf{a} \in \mathbf{A}} \text{Pre}_{\forall}^{\mathbf{a}}(D)$ where

$$\text{Pre}_{\forall}^{\mathbf{a}}(D) \stackrel{\text{def}}{=} \{\mathbf{v} \in \mathbb{N}^d \mid \mathbf{v} + \mathbf{a} \in \mathbb{N}^d \implies \mathbf{v} + \mathbf{a} \in D\} \quad (6)$$

$$= \{\mathbf{v} \in \mathbb{N}^d \mid \mathbf{v} + \mathbf{a} \notin \mathbb{N}^d\} \cup \{\mathbf{v} \in \mathbb{N}^d \mid \mathbf{v} + \mathbf{a} \in D\} \quad (7)$$

$$= \mathbb{N}^d \setminus \uparrow\theta(\mathbf{a}) \cup \{\mathbf{v} \in \mathbb{N}^d \mid \mathbf{v} + \mathbf{a} \in D\}, \quad (8)$$

where $\theta(\mathbf{a}) \stackrel{\text{def}}{=} \min_{\square} \{\mathbf{v} \in \mathbb{N}^d \mid \mathbf{v} + \mathbf{a} \in \mathbb{N}^d\}$ is called the *threshold* of \mathbf{a} and can be computed for all $1 \leq i \leq d$ by

$$\theta(\mathbf{a})(i) = \begin{cases} 0 & \text{if } \mathbf{a}(i) \geq 0 \\ -\mathbf{a}(i) & \text{otherwise.} \end{cases} \quad (9)$$

Thus by (CU') it only remains to compute a representation for the decomposition of $\{\mathbf{v} \in \mathbb{N}^d \mid \mathbf{v} + \mathbf{a} \in D\} = \bigcup_{1 \leq j \leq m} \{\mathbf{v} \in \mathbb{N}^d \mid \mathbf{v} + \mathbf{a} \in \downarrow \mathbf{u}_j\}$. For each ideal $\downarrow \mathbf{u}_j$ in the decomposition of D , $\{\mathbf{v} \in \mathbb{N}^d \mid \mathbf{v} + \mathbf{a} \in \downarrow \mathbf{u}_j\}$ is either the empty set if $\mathbf{u}_j \not\geq \theta(-\mathbf{a})$, or $\downarrow(\mathbf{u}_j - \mathbf{a})$ otherwise, where addition is extended with $\omega + z = \omega$ for all z in \mathbb{Z} .

Example 3.4. Recall the VAS $\mathbf{A}_{\div 2} = \{(-2, 1)\}$ from Example 2.2, and consider the target configuration $\mathbf{t} = (0, 5)$. This yields $D_0 = \downarrow(\omega, 4)$, and the backward coverability algorithm computes the set of all configurations from which $\mathbf{A}_{\div 2}$ cannot compute at least 5 in its second component; see Figure 1.

3.2.2. *Universal Predecessors in Reset VAS.* The same reasoning holds for reset VAS as for VAS:

$$\text{Pre}_{\forall}(D) = \bigcap_{(\mathbf{a}, R) \in \mathbf{A}} \left(\mathbb{N}^d \setminus \uparrow\theta(\mathbf{a}) \cup \bigcup_{1 \leq j \leq m} \{\mathbf{v} \in \mathbb{N}^d \mid R(\mathbf{v} + \mathbf{a}) \in \downarrow \mathbf{u}_j\} \right), \quad (10)$$

(where $R(\mathbf{v} + \mathbf{a})$ is defined as in Example 2.4). In order to compute a representation for this last set, given a vector \mathbf{v} in \mathbb{N}_{ω}^d and $R \subseteq \{1, \dots, d\}$,

define $\bar{\mathbf{v}}^R$ as the R -closure of \mathbf{v} , which replaces the components in R by ω 's:

$$\bar{\mathbf{v}}^R(i) \stackrel{\text{def}}{=} \begin{cases} \omega & \text{if } i \in R \\ \mathbf{v}(i) & \text{otherwise.} \end{cases} \quad (11)$$

Then $\{\mathbf{v} \in \mathbb{N}^d \mid R(\mathbf{v} + \mathbf{a}) \in \downarrow \mathbf{u}_j\}$ is either the empty set if $\bar{\mathbf{u}}_j^R \not\preceq \theta(-\mathbf{a})$, or $\downarrow(\bar{\mathbf{u}}_j^R - \mathbf{a})$ otherwise.

Example 3.5. Recall the reset VAS \mathbf{A}_{\log} from Example 2.3, in which the first two vector components are used to encode two control states. Setting

$$D_0 = \downarrow(1, 0, \omega, \omega, 1) \cup \downarrow(0, 1, \omega, \omega, 0),$$

the backward coverability algorithm computes as follows the set of all configurations from which \mathbf{A}_{\log} cannot compute in its last component either at least 2 in state $(1, 0)$ or at least 1 in state $(0, 1)$.

$$\begin{aligned} D_1 &= \downarrow(0, 0, \omega, \omega, 1) \cup \downarrow(1, 0, 1, \omega, 1) \cup \downarrow(1, 0, \omega, \omega, 0) \cup \downarrow(0, 1, \omega, \omega, 0), \\ D_2 &= \downarrow(0, 0, \omega, \omega, 1) \cup \downarrow(1, 0, 1, 0, 1) \cup \downarrow(1, 0, 0, \omega, 1) \cup \downarrow(1, 0, \omega, \omega, 0) \\ &\quad \cup \downarrow(0, 1, \omega, 0, 0) \cup \downarrow(0, 1, 0, \omega, 0), \\ D_3 &= \downarrow(0, 0, \omega, \omega, 1) \cup \downarrow(1, 0, 1, 0, 1) \cup \downarrow(1, 0, 0, 1, 1) \cup \downarrow(1, 0, \omega, \omega, 0) \\ &\quad \cup \downarrow(0, 1, 2, 0, 0) \cup \downarrow(0, 1, 0, 1, 0), \\ D_4 &= \downarrow(0, 0, \omega, \omega, 1) \cup \downarrow(1, 0, 1, 0, 1) \cup \downarrow(1, 0, 0, 1, 1) \cup \downarrow(1, 0, 1, \omega, 0) \\ &\quad \cup \downarrow(1, 0, \omega, 0, 0) \cup \downarrow(0, 1, 2, 0, 0) \cup \downarrow(0, 1, 0, 1, 0), \\ D_5 &= \downarrow(0, 0, \omega, \omega, 1) \cup \downarrow(1, 0, 1, 0, 1) \cup \downarrow(1, 0, 0, 1, 1) \cup \downarrow(1, 0, 0, \omega, 0) \\ &\quad \cup \downarrow(1, 0, 1, 1, 0) \cup \downarrow(1, 0, 3, 0, 0) \cup \downarrow(0, 1, 2, 0, 0) \cup \downarrow(0, 1, 0, 1, 0), \\ D_* = D_6 &= \downarrow(0, 0, \omega, \omega, 1) \cup \downarrow(1, 0, 1, 0, 1) \cup \downarrow(1, 0, 0, 1, 1) \cup \downarrow(1, 0, 0, 2, 0) \\ &\quad \cup \downarrow(1, 0, 1, 1, 0) \cup \downarrow(1, 0, 3, 0, 0) \cup \downarrow(0, 1, 2, 0, 0) \cup \downarrow(0, 1, 0, 1, 0). \end{aligned}$$

3.3. Ackermann Upper Bounds. Let us finally show how to bound the running time of the backward coverability algorithm on VAS and reset VAS. The main ingredient to that end is a combinatorial statement on the length of *controlled* descending chains of downwards-closed sets.

3.3.1. Controlled Descending Chains. Consider some set X with a norm $\|\cdot\|: X \rightarrow \mathbb{N}$. Given a monotone *control* function $g: \mathbb{N} \rightarrow \mathbb{N}$ and an *initial norm* $n \in \mathbb{N}$, we say that a sequence x_0, x_1, \dots of elements from X is (g, n) -*controlled* if $\|x_i\| \leq g^i(n)$ the i th iterate of g applied to n . In particular, $\|x_0\| \leq n$ initially.

This notion can be applied to the descending chain $D_0 \supseteq D_1 \supseteq \dots$ constructed by the backward coverability algorithm for a d -dimensional VAS or reset VAS \mathbf{A} and target vector $\mathbf{t} \in \mathbb{N}^d$. We define for this $\|\cdot\|$ as the infinity norm on elements and finite subsets of $\mathbb{Z}_\omega^d \stackrel{\text{def}}{=} (\mathbb{Z} \uplus \{\omega\})^d$, i.e. the maximum absolute value of any occurring integer. For instance, $\|\{(1, \omega, 5), (0, \omega, \omega)\}\| = 5$, and in Example 2.2 $\|\mathbf{A}_{\dot{2}}\| = 2$. When considering a downwards-closed set D with decomposition $\downarrow \mathbf{u}_1 \cup \dots \cup \downarrow \mathbf{u}_m$, we define $\|D\| \stackrel{\text{def}}{=} \|\{\mathbf{u}_1, \dots, \mathbf{u}_m\}\|$. Hence what is controlled in a descending chain $D_0 \supseteq D_1 \supseteq \dots$ is actually its ideal representation.

Claim 3.6 (Control for VAS and Reset VAS). The descending chain $D_0 \supseteq D_1 \supseteq \dots$ is (g, n) -controlled for $g(x) \stackrel{\text{def}}{=} x + \|\mathbf{A}\|$ and $n \stackrel{\text{def}}{=} \|\mathbf{t}\|$.

Proof. The fact that $\|D_0\| \leq \|\mathbf{t}\|$ follows from analysing (CU') in Example 2.5. Regarding the control function g , observe that taking unions and intersections of ideals using (II) cannot increase the norm. Hence it suffices to show that $\|\text{Pre}_\forall(D)\| \leq \|D\| + \|\mathbf{A}\|$ for all $D = \downarrow \mathbf{u}_1 \cup \dots \cup \downarrow \mathbf{u}_m$. Note that for reset VAS, $\|\overline{\mathbf{u}_j^R} - \mathbf{a}\| \leq \|\mathbf{u}_j - \mathbf{a}\|$. Hence for both VAS and reset VAS, $\|\text{Pre}_\forall(D)\| \leq \max_{\mathbf{a}} \max_{1 \leq j \leq m} (\|\mathbb{N}^d \setminus \uparrow \theta(\mathbf{a})\|, \|\mathbf{u}_j - \mathbf{a}\|)$. We conclude by observing that $\|\mathbb{N}^d \setminus \uparrow \theta(\mathbf{a})\| \leq \|\mathbf{a}\| \leq \|\mathbf{A}\|$ by (CU') and $\|\mathbf{u}_j - \mathbf{a}\| \leq \|\mathbf{u}_j\| + \|\mathbf{a}\| \leq \|D\| + \|\mathbf{A}\|$. \square

3.3.2. Proper Ideals & Upper Bound. Consider a computation $D_0 \supseteq D_1 \supseteq \dots \supseteq D_\ell = D_{\ell+1}$ of the backward coverability algorithm for a VAS or a reset VAS. At each step $0 \leq k \leq \ell$, the cost of computing D_{k+1} from D_k and of checking for termination is polynomial in $\|A\|$ and $\|D_k\|$. The difficulty is to evaluate how large ℓ can be.

The idea here is that, at every step $0 \leq k < \ell$, there is at least one *proper* ideal I_k : an ideal appearing in the representation of D_k but not in that of D_{k+1} ; then $I_k \subseteq D_k$ but $I_k \not\subseteq D_{k+1}$. Note that for all $0 \leq j < k < \ell$, $I_j \not\subseteq I_k$, since the contrary would entail $I_j \subseteq I_k \subseteq D_k \subseteq D_{j+1}$. Hence the sequence $(I_k)_{0 \leq k < \ell}$ of proper ideals is a *bad* sequence.

In the case of VAS and reset VAS, those proper ideals I_k are of the form $\downarrow \mathbf{v}_k$ for some representant \mathbf{v}_k in \mathbb{N}_ω^d , and the sequence $\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_{\ell-1}$ is also controlled by (g, n) according to Claim 3.6. Using the combinatorial results from [26, Corollary 2.25 and Theorem 2.34] on such bad sequences, we obtain the following *length function theorem* for descending chains of downwards-closed subsets of \mathbb{N}^d ; as the point of this paper is rather to avoid such high complexities, the reader will find all the necessary details in A.

Theorem 3.7 (Length Function Theorem for Descending Chains). *Let $n > 0$. Any (g, n) -controlled descending chain $D_0 \supseteq D_1 \supseteq \dots$ of downwards-closed subsets of \mathbb{N}^d is of length at most $h_{\omega^{d+1}}(n \cdot d!)$, where $h(x) \stackrel{\text{def}}{=} d \cdot g(x)$.*

Here h_α for an ordinal α and base function h denotes the α th Cichoń function [26]. Each of the ℓ steps of computation can furthermore be performed in time polynomial in $g^\ell(n)$.

Since g is primitive-recursive according to Claim 3.6, the overall complexity for an instance of size n is bounded by $\text{ackermann}(p(n))$ for some primitive-recursive function p , which lies within the complexity class ACKERMANN [25, Corollary 4.3]. Such an upper bound is overly pessimistic for VAS, but is actually tight for reset VAS: coverability for reset VAS is indeed complete for ACKERMANN [27, 26].

4. COMPLEXITY FOR VAS

We know from Bozzelli and Ganty's 2EXPTIME upper bound [7] for the backward coverability algorithm that the ACKERMANN upper bound from the previous section is far from tight in the case of VAS. We show in this section that the descending chains $D_0 \supseteq D_1 \supseteq \dots$ computed by the backward coverability algorithm for VAS enjoy a structural invariant, which we dub

ω -monotonicity, and which is absent from the chains computed for reset VAS. In turn, we show in Theorem 4.4, that controlled decreasing chains that are ω -monotone are much shorter, allowing us to derive the desired 2EXPTIME bound in Corollary 4.6.

4.1. Transitions Between Proper Ideals. The proof of ω -monotonicity in the case of VAS can be shown directly, but reflects a more general *proper transition sequence* property of the generic backward coverability algorithm, which we are going to show in the general setting.

4.1.1. *Transitions over Ideals.* Let us first lift the transition relation \rightarrow of a WSTS (X, \rightarrow, \leq) to work over ideals. Define for any ideal I of X

$$\text{Post}_{\exists}(I) \stackrel{\text{def}}{=} \{z \in X \mid \exists x \in I . x \rightarrow z\} . \quad (12)$$

Then $\downarrow\text{Post}_{\exists}(I)$ is downwards-closed with a unique decomposition into maximal ideals. We follow Blondin et al. [4] and write ' $I \rightarrow J$ ' if J is an ideal from the decomposition of $\downarrow\text{Post}_{\exists}(I)$.

Example 4.1 (Transitions over Vector Ideals). In the case of a VAS \mathbf{A} , observe that, if \mathbf{v} is a vector from \mathbb{N}_{ω}^d , then $\text{Post}_{\exists}(\downarrow\mathbf{v}) = \bigcup_{\mathbf{a} \in \mathbf{A}} \downarrow(\mathbf{v} + \mathbf{a})$. Each such $\downarrow(\mathbf{v} + \mathbf{a})$ is already an ideal. In the case of a reset VAS \mathbf{A} , we have similarly $\text{Post}_{\exists}(\downarrow\mathbf{v}) = \bigcup_{(\mathbf{a}, R) \in \mathbf{A}} \downarrow R(\mathbf{v} + \mathbf{a})$.

4.1.2. *Proper Transition Sequences.* We can now state the result that motivates this subsection:

Claim 4.2 (Proper Transition Sequence). If I_{k+1} is a proper ideal of D_{k+1} , then there exist an ideal J and a proper ideal I_k of D_k such that $I_{k+1} \rightarrow J \subseteq I_k$.

Proof. Consider first a computation $D_0 \supseteq D_1 \supseteq \dots \supseteq D_{\ell} = D_{\ell+1}$ of the generic backward coverability algorithm, where each of the D_k is represented as a finite union of ideals. At each refinement step $D_k \supseteq D_{k+1}$ of the algorithm, some of the ideals from the decomposition of D_k —namely the *proper* ones—might be removed, while others might remain untouched in the decomposition of D_{k+1} . Thus, an ideal is proper in D_k if and only if it intersects the set of elements *excluded* between steps k and $k+1$: by basic set operations, first observe that (2) is equivalent to

$$D_{k+1} = D_k \setminus \{x \in D_k \mid \exists z \notin D_k . x \rightarrow z\} . \quad (13)$$

Moreover, noting $D_{-1} \stackrel{\text{def}}{=} X$, z in (13) must belong to D_{k-1} , or x would have already been excluded before step k . We have therefore $D_{k+1} = D_k \setminus E_k$ where

$$E_{-1} \stackrel{\text{def}}{=} \{x \in X \mid x \geq y\} , \quad E_k \stackrel{\text{def}}{=} \{x \in D_k \mid \exists z \in E_{k-1} . x \rightarrow z\} . \quad (14)$$

Consider now a proper ideal I_{k+1} of D_{k+1} : this means $I_{k+1} \cap E_{k+1} \neq \emptyset$. This implies in turn $\downarrow\text{Post}_{\exists}(I_{k+1}) \cap E_k \neq \emptyset$ by (14), thus there exists J such that $I_{k+1} \rightarrow J$ and $J \cap E_k \neq \emptyset$.

Since $I_{k+1} \subseteq D_{k+1} \subseteq \text{Pre}_{\forall}(D_k)$ by (2), we also know that $\text{Post}_{\exists}(I_{k+1}) \subseteq D_k$. By ideal irreducibility, it means that $J \subseteq I_k$ for some ideal I_k from the decomposition of D_k . Observe finally that $I_k \cap E_k \supseteq J \cap E_k \neq \emptyset$, i.e. that I_k is proper. \square

4.2. ω -Monotonicity. For \mathbf{u} in \mathbb{N}_ω^d , its ω -set is the subset $\omega(\mathbf{u})$ of $\{1, \dots, d\}$ such that $\mathbf{u}(i) = \omega$ if and only if $i \in \omega(\mathbf{u})$. We say that a descending chain $D_0 \supseteq D_1 \supseteq \dots \supseteq D_\ell$ of downwards-closed subsets of \mathbb{N}^d is ω -monotone if for all $0 \leq k < \ell - 1$ and all proper ideals $\downarrow \mathbf{v}_{k+1}$ in the decomposition of D_{k+1} , there exists a proper ideal $\downarrow \mathbf{v}_k$ in the decomposition of D_k such that $\omega(\mathbf{v}_{k+1}) \subseteq \omega(\mathbf{v}_k)$. Note that, in VAS and reset VAS alike, because $D_k \supseteq D_{k+1}$ and by ideal irreducibility, for any \mathbf{v}_{k+1} from the decomposition of D_{k+1} there exists some \mathbf{v}_k from that of D_k with $\downarrow \mathbf{v}_{k+1} \subseteq \downarrow \mathbf{v}_k$ and thus $\omega(\mathbf{v}_{k+1}) \subseteq \omega(\mathbf{v}_k)$, but that \mathbf{v}_k may not be proper (then $\mathbf{v}_{k+1} = \mathbf{v}_k$).

Claim 4.3 (VAS Descending Chains are ω -Monotone). The descending chains computed by the backward coverability algorithm for VAS are ω -monotone.

Proof. Let $D_0 \supseteq D_1 \supseteq \dots \supseteq D_\ell$ be the descending chain computed for a VAS \mathbf{A} . Suppose $0 \leq k < \ell - 1$ and $\downarrow \mathbf{v}_{k+1}$ is a proper ideal in the decomposition of D_{k+1} . By Claim 4.2, there exists a proper ideal $\downarrow \mathbf{v}_k$ in the decomposition of D_k such that $\mathbf{v}_{k+1} + \mathbf{a} \sqsubseteq \mathbf{v}_k$. We conclude that $\omega(\mathbf{v}_{k+1}) \subseteq \omega(\mathbf{v}_k)$. \square

As we can see with Example 3.5 however, the descending chains computed for reset VAS are in general *not* ω -monotone: $(1, 0, \omega, \omega, 0)$ is proper in D_3 and has a proper transition to $(0, 1, 0, \omega, 0)$ in D_2 using $(-1, 1, -2, 1, 0, \{3\})$ from \mathbf{A}_{\log} , but no ideal with $\{3, 4\}$ as ω -set is proper in D_2 .

4.3. Upper Bound. We are now in position to state a refinement of Theorem 3.7 for ω -monotone controlled descending chains. For a control function $g: \mathbb{N} \rightarrow \mathbb{N}$, define the function $\tilde{g}: \mathbb{N}^2 \rightarrow \mathbb{N}$ by induction on its first argument:

$$\tilde{g}(0, n) \stackrel{\text{def}}{=} 1, \quad \tilde{g}(m+1, n) \stackrel{\text{def}}{=} \tilde{g}(m, n) + (g^{\tilde{g}(m, n)}(n) + 1)^{m+1}. \quad (15)$$

Theorem 4.4 (Length Function Theorem for ω -Monotone Descending Chains). *Any (g, n) -controlled ω -monotone descending chain $D_0 \supseteq D_1 \supseteq \dots$ of downwards-closed subsets of \mathbb{N}^d is of length at most $\tilde{g}(d, n)$.*

Proof. Note that D_ℓ the last element of the chain has the distinction of not having any proper ideal, hence it suffices to bound the index k of the last set D_k with a proper ideal $\downarrow \mathbf{v}_k$, and add one to get a bound on ℓ . We are going to establish by induction on $d - |I|$ that if $\downarrow \mathbf{v}_k$ is a proper ideal from the decomposition of D_k and its ω -set is I , then $k < \tilde{g}(d - |I|, n)$, which by monotonicity of \tilde{g} in its first argument entails $k < \tilde{g}(d, n)$ as desired.

For the base case, $|I| = d$ implies that \mathbf{v}_k is the vector with ω 's in every coordinate, which can only occur in D_0 . The inductive step is handled by the following claim, when setting $k < \tilde{g}(d - |I| - 1, n)$ by induction hypothesis for the last index with a proper ideal whose ω -set strictly includes I :

Claim 4.5. Let I and $k < k'$ be such that:

- (i) for all $j \in \{k+1, \dots, k'-1\}$, the decomposition of D_j does not contain a proper ideal whose ω -set strictly includes I ;
- (ii) the decomposition of $D_{k'}$ contains a proper ideal whose ω -set is I .

Then we have $k' - k \leq (\|D_{k+1}\| + 1)^{(d-|I|)}$.

For a proof, from assumption (ii), by applying the ω -monotonicity for $j = k' - 1, k' - 2, \dots, k + 1$ and due to assumption (i), there exists a proper ideal $\downarrow \mathbf{v}_j$ in the decomposition of D_j and such that $\omega(\mathbf{v}_j) = I$ for all

$j \in \{k+1, \dots, k'\}$. Since they are proper, those $k' - k$ vectors are mutually distinct.

Consider any such \mathbf{v}_j . Since $D_{k+1} \supseteq D_j$, by ideal irreducibility there exists a vector \mathbf{u}_j in the decomposition of D_{k+1} such that $\mathbf{v}_j \sqsubseteq \mathbf{u}_j$. We have that $\omega(\mathbf{u}_j) = I$, since otherwise \mathbf{u}_j would be proper at $D_{j'}$ for some $j' \in \{k+1, \dots, j-1\}$, which would contradict assumption (i). Hence $\|\mathbf{v}_j\| \leq \|\mathbf{u}_j\| \leq \|D_{k+1}\|$.

To conclude, note that there can be at most $(\|D_{k+1}\| + 1)^{(d-|I|)}$ mutually distinct vectors in \mathbb{N}_ω^d with I as ω -set and norm bounded by $\|D_{k+1}\|$. \square

Finally, putting together Claim 3.6 (control for VAS), Claim 4.3 (ω -monotonicity), and Theorem 4.4 (lengths of controlled ω -monotone descending chains), we obtain that the backward coverability algorithm for VAS runs in 2EXPTIME, and in pseudo-polynomial time if the dimension d is fixed.

Corollary 4.6. *For any d -dimensional VAS \mathbf{A} and target vector \mathbf{t} , the backward coverability algorithm terminates after at most $((\|\mathbf{A}\| + 1)(\|\mathbf{t}\| + 2))^{(d+1)!}$ steps.*

Proof. Let $h(m, n) = \tilde{g}(m, n)(\|\mathbf{A}\| + 1)(n + 2)$ where $g(x) = x + \|\mathbf{A}\|$. We have $h(m + 1, n) \leq (h(m, n))^{m+2}$, so $\tilde{g}(m, n) \leq h(m, n) \leq ((\|\mathbf{A}\| + 1)(n + 2))^{(m+1)!}$. \square

As mentioned in the introduction, this 2EXPTIME upper bound is not tight for the coverability problem for VAS—which is EXPSpace-complete [20, 22]. However, the bound is tight for the backward coverability algorithm: as shown by Bozzelli and Ganty [7] using Lipton’s construction, there is a uniform family of VAS $(\mathbf{A}_n)_{n \in \mathbb{N}}$ each of size $O(n)$ on which the algorithm requires 2^{2^n} iterations.

5. TOP-DOWN TREE COVERABILITY

We turn to demonstrating how easily our new proof of the doubly-exponential bound for the backward coverability algorithm on VAS can be extended to derive (the known) optimal bounds for top-down alternating branching VAS: TOWER in general [17] and 2EXPTIME with alternation only [8]. While this complexity analysis does not yield any new results, we believe it brings some additional insights into our dual view of the backward coverability algorithm. Also, unlike with VAS, the upper bounds we obtain are tight.

5.1. Alternating Branching VAS. Recall that an *alternating branching vector addition system (ABVAS)* of dimension $d \in \mathbb{N}$ consists of: a finite set of *unary* rules $\mathbf{A} \subseteq \mathbb{Z}^d$, a finite set of *fork* rules $\mathbf{B}_\wedge \subseteq \mathbb{Z}^d \times \mathbb{Z}^d$, and a finite set of *split* rules $\mathbf{B}_+ \subseteq \mathbb{Z}^d$. A computation is a tree whose nodes are labelled by vectors in \mathbb{N}^d , and such that every non-leaf node ν is obtained by applying some rule (we write $\mathbf{v}(\cdot)$ for node labels):

- either ν has one child ν' and $\mathbf{v}(\nu) + \mathbf{a} = \mathbf{v}(\nu')$ for some unary rule \mathbf{a} from \mathbf{A} ,
- or ν has two children ν_1, ν_2 and $\mathbf{v}(\nu) + \mathbf{b}_1 = \mathbf{v}(\nu_1)$ and $\mathbf{v}(\nu) + \mathbf{b}_2 = \mathbf{v}(\nu_2)$ for some fork rule $(\mathbf{b}_1, \mathbf{b}_2)$ from \mathbf{B}_\wedge ,

- or ν has two children ν_1, ν_2 and $\mathbf{v}(\nu) + \mathbf{b} = \mathbf{v}(\nu_1) + \mathbf{v}(\nu_2)$ for some split rule \mathbf{b} from \mathbf{B}_+ .

Thus a VAS is an ABVAS with empty \mathbf{B}_\wedge and \mathbf{B}_+ ; an ABVAS with $\mathbf{B}_\wedge = \emptyset$ is called a *branching* VAS (BVAS), and one with $\mathbf{B}_+ = \emptyset$ is called an *alternating* VAS (AVAS). Such systems appear in a variety of contexts in relation with fragments of propositional linear logic, simulation relations, computational linguistics, etc. [9, 24, 8, 17].

The *top-down coverability* problem (also called the *leaf coverability* problem), given a d -dimensional top-down ABVAS $(\mathbf{A}, \mathbf{B}_\wedge, \mathbf{B}_+)$, and root and target vectors $\mathbf{r}, \mathbf{t} \in \mathbb{N}^d$, asks whether there exists a computation whose root label is \mathbf{r} and whose every leaf label is $\sqsupseteq \mathbf{t}$. We remark that the problem is equivalent to the *join* variant in which the fork rules are applied by taking pointwise maxima (i.e. $\mathbf{v}(\nu) = \max\{\mathbf{v}(\nu_1) - \mathbf{b}_1, \mathbf{v}(\nu_2) - \mathbf{b}_2\}$), and also to the reachability problem for *lossy* ABVAS [17].

5.2. Backward Coverability. The theory of well-structured transition systems can be lifted in a very natural way to top-down tree computations, although we are not aware of any literature on the subject. One can nevertheless find enlightening parallels with *monotonic games* [2]. A top-down coverability instance for an ABVAS defines indeed a two-players game over an infinite arena with positions in \mathbb{N}^d . Starting from \mathbf{r} , Player 1 attempts to cover \mathbf{t} , while Player 2 attempts to prevent it. At each round, in a configuration $\mathbf{v} \in \mathbb{N}^d$,

- (1) Player 1 chooses an applicable rule \mathbf{a} from \mathbf{A} with $\mathbf{v} + \mathbf{a} \sqsupseteq \mathbf{0}$, or $(\mathbf{b}_1, \mathbf{b}_2)$ from \mathbf{B}_\wedge with $\mathbf{v} + \mathbf{b}_1 \sqsupseteq \mathbf{0}$ and $\mathbf{v} + \mathbf{b}_2 \sqsupseteq \mathbf{0}$, or \mathbf{b} from \mathbf{B}_+ with $\mathbf{v} + \mathbf{b} \sqsupseteq \mathbf{0}$ and in that case further chooses two vectors $\mathbf{v}_1, \mathbf{v}_2$ from \mathbb{N}^d such that $\mathbf{v} + \mathbf{b} = \mathbf{v}_1 + \mathbf{v}_2$.
- (2) Player 2 then chooses a branch, i.e. $\mathbf{v}' = \mathbf{v} + \mathbf{a}$ in case of a unary rule \mathbf{a} from \mathbf{A} , or $\mathbf{v}' = \mathbf{v} + \mathbf{b}_i$ for $i \in \{1, 2\}$ in case of a fork rule $(\mathbf{b}_1, \mathbf{b}_2)$ in \mathbf{B}_\wedge , or $\mathbf{v}' = \mathbf{v}_i + \mathbf{b}$ for $i \in \{1, 2\}$ in case of a split rule \mathbf{b} from \mathbf{B}_+ , and the game proceeds to the next round from \mathbf{v}' .

Player 1 wins the game if a configuration $\mathbf{v} \sqsupseteq \mathbf{t}$ is eventually reached. One can see that ABVAS computation trees represent strategies for Player 1, and such a tree has all its leaves $\sqsupseteq \mathbf{t}$ if and only if it is a winning strategy.

5.2.1. Winning Regions. Observe that, using this game-theoretic viewpoint, the *winning region* for Player 1 is an upwards-closed set of configurations, and conversely—the game is determined—the winning region for Player 2 is downwards-closed. In our dual view of the backward coverability algorithm using downwards-closed sets, we actually compute the winning region D_* for Player 2 by successive refinements of the winning region D_k for the game limited to at most k rounds. A configuration \mathbf{v} then belongs to D_{k+1} , if for all choices of rules by Player 1, there exists a branch with \mathbf{v}' in the winning region D_k of Player 2. Using this insight, the generic algorithm from Section 3 can be adapted to top-down coverability by replacing the Pre_\forall operator with

$$\text{Pre}_{\forall\exists}(S) \stackrel{\text{def}}{=} \text{Pre}_\forall^{\mathbf{A}}(S) \cap \text{Pre}_{\forall\exists}^{\mathbf{B}_\wedge}(S) \cap \text{Pre}_{\forall\exists}^{\mathbf{B}_+}(S) \quad (16)$$

where

$$\text{Pre}_{\forall}^{\mathbf{A}}(S) \stackrel{\text{def}}{=} \{v \in \mathbb{N}^d \mid \forall \mathbf{a} \in \mathbf{A}. v + \mathbf{a} \in \mathbb{N}^d \Rightarrow v + \mathbf{a} \in S\}, \quad (17)$$

$$\text{Pre}_{\forall\exists}^{\mathbf{B}_\wedge}(S) \stackrel{\text{def}}{=} \left\{ v \in \mathbb{N}^d \mid \begin{array}{l} \forall (\mathbf{b}_1, \mathbf{b}_2) \in \mathbf{B}_\wedge. v + \mathbf{b}_1 \in \mathbb{N}^d \wedge v + \mathbf{b}_2 \in \mathbb{N}^d \\ \Rightarrow v + \mathbf{b}_1 \in S \vee v + \mathbf{b}_2 \in S \end{array} \right\}, \quad (18)$$

$$\text{Pre}_{\forall\exists}^{\mathbf{B}_+}(S) \stackrel{\text{def}}{=} \left\{ v \in \mathbb{N}^d \mid \begin{array}{l} \forall \mathbf{b} \in \mathbf{B}_+. \forall \mathbf{v}_1, \mathbf{v}_2 \in \mathbb{N}^d. v + \mathbf{b} = \mathbf{v}_1 + \mathbf{v}_2 \\ \Rightarrow \mathbf{v}_1 \in S \vee \mathbf{v}_2 \in S \end{array} \right\}. \quad (19)$$

5.2.2. Downward-Closure and Correctness. It is straightforward to check that $\text{Pre}_{\forall\exists}$ preserves the property of downward closure as in Claim 3.2. Given an initial downwards-closed $D_0 \subseteq \mathbb{N}^d$, as in Section 3, we write D_* for the last set in the longest (necessarily finite) descending chain of downwards-closed sets $D_0 \supseteq D_1 \supseteq \dots$ defined by $D_{k+1} = D_k \cap \text{Pre}_{\forall\exists}(D_k)$. From the definition of $\text{Pre}_{\forall\exists}$, we have the required correctness property, matching Claim 3.1:

Claim 5.1 (Correctness). For all k , D_k consists of all vectors $v \in \mathbb{N}^d$ such that all computation trees of $(\mathbf{A}, \mathbf{B}_\wedge, \mathbf{B}_+)$ whose root label is v and whose height is at most k have some leaf label in D_0 .

Thus, the set D_* consists of all vectors $v \in \mathbb{N}^d$ such that all computations of $(\mathbf{A}, \mathbf{B}_\wedge, \mathbf{B}_+)$ whose root label is v have some leaf label in D_0 .

5.2.3. Effectiveness and Control. Similarly to the cases of VAS and reset VAS in Sections 3.2 and 3.3, we shall give an explicit computation for predecessors, allowing us to provide bounds on the norms of ideals in the course of the backward coverability algorithm. The counterpart to Claim 3.6 is then:

Claim 5.2 (Control for ABVAS).

- (a) The ideal decomposition of $\text{Pre}_{\forall\exists}(D)$ is computable for all downwards-closed $D \subseteq \mathbb{N}^d$.
- (b) The descending chain $D_0 \supseteq D_1 \supseteq \dots$ is (g, n) -controlled for

$$g(x) \stackrel{\text{def}}{=} \max\{x + \|\mathbf{A}\|, x + \|\mathbf{B}_\wedge\|, 2x + \|\mathbf{B}_+\| + 1\}$$

and $n \stackrel{\text{def}}{=} \|\mathbf{t}\|$. For AVAS, i.e. when \mathbf{B}_+ is empty, the term $2x + \|\mathbf{B}_+\| + 1$ disappears.

Proof. Let us examine the three cases in (16); by (II) the norm of the ideals in $\text{Pre}_{\forall\exists}(D)$ will be bounded by the maximal norm over all three cases. For (17),

$$\text{Pre}_{\forall}^{\mathbf{A}}(D) = \bigcap_{\mathbf{a} \in \mathbf{A}} \text{Pre}_{\forall}^{\mathbf{a}}(D) \quad (20)$$

as defined in (6), entailing $\|\text{Pre}_{\forall}^{\mathbf{A}}(D)\| \leq \|D\| + \|\mathbf{A}\|$. For (18),

$$\text{Pre}_{\forall\exists}^{\mathbf{B}_\wedge}(D) = \bigcap_{(\mathbf{b}_1, \mathbf{b}_2) \in \mathbf{B}_\wedge} (\text{Pre}_{\forall}^{\mathbf{b}_1}(D) \cup \text{Pre}_{\forall}^{\mathbf{b}_2}(D)) \quad (21)$$

and thus $\|\text{Pre}_{\forall\exists}^{\mathbf{B}_\wedge}(D)\| \leq \|D\| + \|\mathbf{B}_\wedge\|$. Finally, regarding (19), we are here in a situation where Remark 3.3 and footnote 2 lead to a much simpler analysis:

$$\text{Pre}_{\forall\exists}^{\mathbf{B}_+}(D) = \mathbb{N}^d \setminus (\text{Pre}_{\exists\forall}^{\mathbf{B}_+}(\mathbb{N}^d \setminus D)) \quad (22)$$

where $\|\mathbb{N}^d \setminus D\| \leq \|D\| + 1$ when considering the norms of the minimal elements of $\mathbb{N}^d \setminus D$, and

$$\text{Pre}_{\exists\forall}^{\mathbf{B}_+}(\mathbb{N}^d \setminus D) = \bigcup_{\mathbf{b} \in \mathbf{B}_+} \{\mathbf{v}_1 + \mathbf{v}_2 - \mathbf{b} \in \mathbb{N}^d \mid \mathbf{v}_1, \mathbf{v}_2 \in \mathbb{N}^d \setminus D\}. \quad (23)$$

This last expression yields $\|\text{Pre}_{\exists\forall}^{\mathbf{B}_+}(\mathbb{N}^d \setminus D)\| \leq 2(\|D\| + 1) + \|\mathbf{B}_+\|$, and finally $\|\text{Pre}_{\forall\exists}^{\mathbf{B}_+}(D)\| \leq 2\|D\| + \|\mathbf{B}_+\| + 1$ as desired. \square

5.3. Complexity Upper Bounds.

5.3.1. ω -Monotonicity. The property that allows us to deduce that the backward coverability algorithm is optimal for top-down ABVAS, and also when restricted to AVAS, is again ω -monotonicity of the downwards-closed sets that it computes. We extend for this claims 4.2 and 4.3.

Claim 5.3. The descending chains computed by the backward coverability algorithm for top-down ABVAS are ω -monotone.

Proof. Let $D_0 \supseteq D_1 \supseteq \dots \supseteq D_\ell$ be the descending chain computed for a top-down ABVAS $(\mathbf{A}, \mathbf{B}_\wedge, \mathbf{B}_+)$. Suppose $0 \leq k < \ell - 1$ and \mathbf{v}_{k+1} is a proper vector in the decomposition of D_{k+1} . Then $\downarrow \mathbf{v}_{k+1} \not\subseteq \text{Pre}_{\forall\exists}(D_{k+1})$, so there is a case for each of the three types of rule.

Split: Here some $\mathbf{b} \in \mathbf{B}_+$ and $\mathbf{v}', \mathbf{v}'' \in \mathbb{N}_\omega^d$ are such that $\mathbf{v}_{k+1} + \mathbf{b} = \mathbf{v}' + \mathbf{v}''$ and $\downarrow \mathbf{v}', \downarrow \mathbf{v}'' \not\subseteq D_{k+1}$. Without loss of generality, we may assume that the ω -sets of \mathbf{v}, \mathbf{v}' and \mathbf{v}'' are the same. Since $\downarrow \mathbf{v}_{k+1} \subseteq \text{Pre}_{\forall\exists}(D_k)$, we have that either $\downarrow \mathbf{v}' \subseteq D_k$ or $\downarrow \mathbf{v}'' \subseteq D_k$, say the former. Let \mathbf{v}_k be any vector in the decomposition of D_k such that $\mathbf{v}' \sqsubseteq \mathbf{v}_k$. We conclude that \mathbf{v}_k is proper and that $\omega(\mathbf{v}_{k+1}) \subseteq \omega(\mathbf{v}_k)$.

Fork: This case is similar but easier as we may assume $\mathbf{v}_{k+1} + \mathbf{b}_1 = \mathbf{v}'$ and $\mathbf{v}_{k+1} + \mathbf{b}_2 = \mathbf{v}''$.

Unary: This case is as for VAS, cf. the proof of Claim 4.3. \square

5.3.2. Upper Bounds. We are now equipped to establish, by applying the length function theorem for ω -monotone descending chains (Theorem 4.4) that the backward coverability algorithm for top-down ABVAS runs in TOWER in general and 2EXPTIME with alternation only. Since the ideal decomposition of each D_{k+1} is computable in time polynomial in the bound on the norm of D_k , it suffices to bound the number of iterations of the main loop.

Recall Equation 15, and let $h(m, n) \stackrel{\text{def}}{=} 32(n + L)(\tilde{g}^\dagger(m, n))^2$ where

$$L \stackrel{\text{def}}{=} \max\{\|\mathbf{A}\|, \|\mathbf{B}_\wedge\|, \|\mathbf{B}_+\| + 1\} \quad (24)$$

$$g^\dagger(x) \stackrel{\text{def}}{=} 2x + L. \quad (25)$$

We have

$$\tilde{g}^\dagger(m+1, n) \leq (2^{\tilde{g}^\dagger(m, n)}(n+L+1))^{m+1} \leq 2^{(\tilde{g}^\dagger(m, n)+1)(n+L+1)(m+1)} \leq 2^{h(m, n)/4}$$

and hence

$$h(m+1, n) \leq 32(n+L)2^{h(m, n)/2} \leq 2^{5+n+L+h(m, n)/2} \leq 2^{h(m, n)},$$

so $h(d, n) \leq \underbrace{2^{\dots^2}}_d^{32(n+L)}$.

Let also $h'(m, n) = \tilde{g}^\ddagger(m, n)(L' + 1)(n + 2)$ where $L' = \max\{\|\mathbf{A}\|, \|\mathbf{B}_\wedge\|\}$ and $g^\ddagger(x) = x + L'$. We have $h'(m + 1, n) \leq (h'(m, n))^{m+2}$, so $h'(d, n) \leq ((L' + 1)(n + 2))^{(d+1)!}$.

Theorem 4.4 gives us:

Corollary 5.4. *For any d -dimensional top-down ABVAS $(\mathbf{A}, \mathbf{B}_\wedge, \mathbf{B}_+)$ and initial downwards-closed $D_0 \subseteq \mathbb{N}^d$, the backward coverability algorithm terminates after at most $\underbrace{2^{\dots^2}}_d^{32(\|D_0\|+L)}$ steps, where $L = \max\{\|\mathbf{A}\|, \|\mathbf{B}_\wedge\|, \|\mathbf{B}_+\| + 1\}$. For AVAS, i.e. when \mathbf{B}_+ is empty, the algorithm terminates after at most $((L' + 1)(\|D_0\| + 2))^{(d+1)!}$ steps, where $L' = \max\{\|\mathbf{A}\|, \|\mathbf{B}_\wedge\|\}$.*

These upper bounds are optimal: top-down coverability is indeed TOWER-hard for ABVAS [17] and 2EXPTIME-hard for AVAS [8].

6. BOTTOM-UP TREE COVERABILITY

As the last case study in this paper of the ideal view of Rackoff's technique, we consider the coverability problem for bottom-up coverability in *meet* ABVAS. We discover a backward coverability algorithm that, instead of configurations that are vectors of natural numbers as in Sections 4 and 5, works with downwards-closed sets of such vectors. The ideals involved are therefore one level higher than the vector ideals, and we recall some of their properties from the literature. We follow the pattern from the preceding sections and derive the optimal complexity bounds: ACKERMANN in general [17] and 2EXPTIME for BVAS [10].

6.1. Vector Set Ideals. We first recall some properties of the downwards-closed powerset of \mathbb{N}^d . We write $\mathbb{D}(X)$ for the set of all subsets of X that are downwards-closed with respect to a quasi-ordering \leq . The subset relation is a quasi-ordering on $\mathbb{D}(X)$. Because $(\mathbb{N}^d, \sqsubseteq)$ is an ω^2 -wqo (it is even a *better-quasi-order*), $(\mathbb{D}(\mathbb{N}^d), \subseteq)$ is a wqo.

Fact 6.1 ([12, 15]). *For any $d \in \mathbb{N}$, $(\mathbb{D}(\mathbb{N}^d), \subseteq)$ is a wqo with effective ideal representations.*

Indeed, the ideals of $(\mathbb{D}(\mathbb{N}^d), \subseteq)$ are all $\downarrow D = \{D' \in \mathbb{D}(\mathbb{N}^d) \mid D' \subseteq D\}$ for downwards-closed $D \subseteq \mathbb{N}^d$, and can thus be represented as finite sets of vectors from \mathbb{N}_ω^d .

6.2. Meet ABVAS are defined as ABVAS (cf. Section 5), except that applications of fork rules take component-wise minima. Thus, for every non-leaf node ν in a computation:

- either ν has one child ν' and $\mathbf{v}(\nu) + \mathbf{a} = \mathbf{v}(\nu')$ for some unary rule \mathbf{a} from \mathbf{A} ,
- or ν has two children ν_1, ν_2 and $\mathbf{v}(\nu) = \min\{\mathbf{v}(\nu_1) - \mathbf{b}_1, \mathbf{v}(\nu_2) - \mathbf{b}_2\}$ for some fork rule $(\mathbf{b}_1, \mathbf{b}_2)$ from \mathbf{B}_\wedge ,

- or ν has two children ν_1, ν_2 and $\mathbf{v}(\nu) + \mathbf{b} = \mathbf{v}(\nu_1) + \mathbf{v}(\nu_2)$ for some split rule \mathbf{b} from \mathbf{B}_+ .

The *bottom-up coverability* problem (also known as the *root coverability* problem), given a d -dimensional meet ABVAS $(\mathbf{A}, \mathbf{B}_\wedge, \mathbf{B}_+)$, and leaf and target vectors $\ell, \mathbf{t} \in \mathbb{N}^d$, asks whether there exists a computation whose every leaf label is ℓ and whose root label is $\sqsupseteq \mathbf{t}$. We remark that the problem is equivalent to the reachability problem for *gainy* ABVAS [17] but that the variant of bottom-up coverability in which fork rules are applied with exact alternation is undecidable [19].

To solve the coverability problem, we first define an operator on sets of vectors:

$$\begin{aligned} \text{Post}(S) &\stackrel{\text{def}}{=} \{\mathbf{v} - \mathbf{a} \in \mathbb{N}^d \mid \mathbf{v} \in S \wedge \mathbf{a} \in \mathbf{A}\} \\ &\cup \{\min\{\mathbf{v}_1 - \mathbf{b}_1, \mathbf{v}_2 - \mathbf{b}_2\} \in \mathbb{N}^d \mid \mathbf{v}_1, \mathbf{v}_2 \in S \wedge (\mathbf{b}_1, \mathbf{b}_2) \in \mathbf{B}_\wedge\} \\ &\cup \{\mathbf{v}_1 + \mathbf{v}_2 - \mathbf{b} \in \mathbb{N}^d \mid \mathbf{v}_1, \mathbf{v}_2 \in S \wedge \mathbf{b} \in \mathbf{B}_+\}, \end{aligned}$$

and then instantiate the generic backward algorithm from Section 3 by the following operator on sets of downward-closed sets of vectors:

$$\mathcal{PRE}(\mathcal{S}) \stackrel{\text{def}}{=} \{D \in \mathbb{D}(\mathbb{N}^d) \mid D \cup \downarrow \text{Post}(D) \in \mathcal{S}\}. \quad (26)$$

6.3. Downward-Closure and Correctness. It is straightforward to check that \mathcal{PRE} preserves the property of downward closure. Given an initial downwards-closed $\mathcal{D}_0 \subseteq \mathbb{D}(\mathbb{N}^d)$, as in Section 3, we write \mathcal{D}_* for the last set in the longest (necessarily finite) descending chain of downwards-closed sets $\mathcal{D}_0 \supseteq \mathcal{D}_1 \supseteq \dots$ defined by $\mathcal{D}_{k+1} = \mathcal{D}_k \cap \mathcal{PRE}(\mathcal{D}_k)$. Observing that $\mathcal{PRE}(\mathcal{D}) \subseteq \mathcal{D}$ whenever \mathcal{D} is downwards-closed, we actually have $\mathcal{D}_{k+1} = \mathcal{PRE}(\mathcal{D}_k)$. From the definitions of Post and \mathcal{PRE} , we have the required correctness property:

Claim 6.2 (Correctness).

- Any \mathcal{D}_k consists of all downwards-closed sets of vectors D for which there exists $D' \in \mathcal{D}_0$ such that all computations of $(\mathbf{A}, \mathbf{B}_\wedge, \mathbf{B}_+)$ whose leaves are in D and whose height is at most k have their root in D' .
- The set \mathcal{D}_* consists of all downwards-closed sets of vectors D for which there exists $D' \in \mathcal{D}_0$ such that all computations of $(\mathbf{A}, \mathbf{B}_\wedge, \mathbf{B}_+)$ whose leaves are in D have their root in D' .

6.4. Effectiveness and Control. Let the norm of a downwards-closed subset of $\mathbb{D}(\mathbb{N}^d)$ whose ideal decomposition is $\downarrow D_1 \cup \dots \cup \downarrow D_n$ be the maximum of the norms of D_1, \dots, D_n .

Similarly to the cases of VAS and reset VAS in Sections 3.2 and 3.3, we obtain:

Claim 6.3 (Control for meet ABVAS).

- The ideal decomposition of $\mathcal{PRE}(\mathcal{D})$ is computable for all downwards-closed $\mathcal{D} \subseteq \mathbb{D}(\mathbb{N}^d)$.
- The descending chain $\mathcal{D}_0 \supseteq \mathcal{D}_1 \supseteq \dots$ is (g, n) -controlled for

$$g(x) \stackrel{\text{def}}{=} x + \max\{\|\mathbf{A}\|, \|\mathbf{B}_\wedge\|, \|\mathbf{B}_+\|\}$$

and $n \stackrel{\text{def}}{=} \|\mathbf{t}\|$.

6.5. ω -Monotonicity. The distinction between bottom-up ABVAS and BVAS appears here. Inside the descending chain of downwards-closed subsets of $\mathbb{D}(\mathbb{N}^d)$ computed by the backward coverability algorithm, we can always find a descending chain of downwards-closed subsets of \mathbb{N}^d , whose length is bounded by the Ackermann function. However, for bottom-up BVAS, the latter descending chain is ω -monotone, enabling us to infer that its length is then at most doubly exponential.

Claim 6.4. Suppose $\mathcal{D}_0 \supseteq \cdots \supseteq \mathcal{D}_\ell$ is the descending chain computed by the backward coverability algorithm for a d -dimensional bottom-up ABVAS $(\mathbf{A}, \mathbf{B}_\wedge, \mathbf{B}_+)$ and an initial downwards-closed $\mathcal{D}_0 \subseteq \mathbb{D}(\mathbb{N}^d)$.

- (a) There exist proper downwards-closed D_k in the ideal decomposition of \mathcal{D}_k for all $k \in \{0, \dots, \ell - 1\}$ such that

$$D_{k+1} \cup \downarrow \text{Post}(D_{k+1}) \subseteq D_k$$

for all $k \in \{0, \dots, \ell - 2\}$.

- (b) For BVAS, i.e. when \mathbf{B}_\wedge is empty, any descending chain $D_0, \dots, D_{\ell-1}$ as in (a) is ω -monotone.

Proof. For (a), pick any proper $D_{\ell-1}$ in the decomposition of $\mathcal{D}_{\ell-1}$. Assuming that D_{k+1} has been picked, we have $\downarrow D_{k+1} \not\subseteq \mathcal{P}\mathcal{R}\mathcal{E}(\mathcal{D}_{k+1})$ and thus $D_{k+1} \not\subseteq \mathcal{P}\mathcal{R}\mathcal{E}(\mathcal{D}_{k+1})$, so $D_{k+1} \cup \downarrow \text{Post}(D_{k+1}) \in \mathcal{D}_k \setminus \mathcal{D}_{k+1}$. Pick any D_k in the decomposition of \mathcal{D}_k such that $D_{k+1} \cup \downarrow \text{Post}(D_{k+1}) \subseteq D_k$.

For (b), suppose $0 \leq j < \ell - 2$ and \mathbf{v}_{k+1} is a proper vector in the decomposition of D_{k+1} , i.e. $\downarrow \mathbf{v}_{k+1} \not\subseteq D_{k+2}$. Since $\downarrow D_{k+2}$ is in the decomposition of \mathcal{D}_{k+2} , we have that $D_{k+2} \cup \downarrow \mathbf{v}_{k+1} \notin \mathcal{D}_{k+2}$. But $D_{k+2} \cup \downarrow \mathbf{v}_{k+1} \subseteq D_{k+1}$, so $\text{Post}(D_{k+2} \cup \downarrow \mathbf{v}_{k+1}) \not\subseteq D_{k+1}$. Recalling that $\text{Post}(D_{k+2}) \subseteq D_{k+1}$, we have $\mathbf{w} \in \mathbb{N}_\omega^d$ and $\downarrow \mathbf{w} \not\subseteq D_{k+1}$ where:

- either $\mathbf{w} = \mathbf{v}_{k+1} - \mathbf{a}$ for some $\mathbf{a} \in \mathbf{A}$,
- or $\mathbf{w} = \mathbf{v}' + \mathbf{v}_{k+1} - \mathbf{b}$ for some \mathbf{v}' in the decomposition of D_{k+2} and $\mathbf{b} \in \mathbf{B}_+$.

In both cases, $\downarrow \mathbf{w} \subseteq \downarrow \text{Post}(D_{k+2} \cup \downarrow \mathbf{v}_{k+1}) \subseteq D_k$. Let \mathbf{v}_k be any vector in the decomposition of D_k such that $\mathbf{w} \sqsubseteq \mathbf{v}_k$. We conclude that \mathbf{v}_k is proper and that $\omega(\mathbf{v}_{k+1}) \subseteq \omega(\mathbf{v}_k)$. \square

6.6. Upper Bounds. Finally, by applying the length function theorem for descending chains (Theorem 3.7), or for ω -monotone descending chains (Theorem 4.4) in the case of BVAS, we establish the ACKERMANN and 2EXPTIME bounds. Again, the ideal decomposition of each \mathcal{D}_{k+1} is computable in time polynomial in $g^k(\|\mathcal{D}_0\|)$, so it suffices to bound the number of iterations of the main loop.

Corollary 6.5. *For any d -dimensional bottom-up ABVAS $(\mathbf{A}, \mathbf{B}_\wedge, \mathbf{B}_+)$ and initial downwards-closed $\mathcal{D}_0 \subseteq \mathbb{D}(\mathbb{N}^d)$, the backward coverability algorithm terminates after at most $\text{ackermann}(p(d, L, \|\mathcal{D}_0\|))$ steps, where p is a primitive-recursive function and $L = \max\{\|\mathbf{A}\|, \|\mathbf{B}_\wedge\|, \|\mathbf{B}_+\|\}$. For BVAS, i.e. when \mathbf{B}_\wedge is empty, the algorithm terminates after at most $((L + 1)(\|\mathcal{D}_0\| + 2))^{(d+1)!} + 1$ steps.*

These upper bounds are tight: ACKERMANN-hardness was shown originally by Urquhart [28] for reachability in a variant of gainy AVAS, and also holds for bottom-up coverability for meet AVAS [17]; 2EXPTIME-hardness was first shown by Demri et al. [10] for bottom-up coverability for BVAS.

7. CONCLUDING REMARKS

Rackoff's technique has successfully been employed to prove tight upper bounds for the coverability problem in VAS and extensions [10, 6, 8, 17, 16]. However, the technique does not readily generalise to more complex classes of well-structured transition systems, e.g. where configurations are not vectors of natural numbers.

We have shown that the same complexity bounds can be extracted in a principled way, by considering the ideal view of the backward coverability algorithm for VAS, and by noticing a structural invariant on its computations. Essentially the same arguments suffice to re-prove several recent upper bounds [10, 8, 17].

This paves the way for future investigations on coverability problems with large complexity gaps (where different structural invariants will need to be found). As an instance of such results, we have recently shown in [18] that the techniques presented in this paper yield new tight upper bounds for coverability in ν -Petri nets [see e.g. 23], improving on the previously known HYPERACKERMANN upper bound.

APPENDIX A. ACKERMANN UPPER BOUNDS

One way to obtain ACKERMANN upper bounds for the backward coverability algorithm on VAS and reset VAS would be to consider the dual ascending chain of *upwards-closed* sets employed in the usual description of the backward coverability algorithm. The resulting bounds would be similar (and the proof somewhat simpler), see [e.g. 26, Section 2.2.2]. Instead, we prove the bounds directly on descending chains, thanks to Theorem 3.7:

Theorem 3.7 (Length Function Theorem for Descending Chains). *Let $n > 0$. Any (g, n) -controlled descending chain $D_0 \supseteq D_1 \supseteq \dots$ of downwards-closed subsets of \mathbb{N}^d is of length at most $h_{\omega^{d+1}}(n \cdot d!)$, where $h(x) \stackrel{\text{def}}{=} d \cdot g(x)$.*

The main tool to this end is the following statement, which combines Corollary 2.25 and Theorem 2.34 from [26]:

Theorem A.1 (Length Function Theorem for Bad Sequences). *Let $n > 0$. Any (g, n) -controlled bad sequence over a polynomial normed wqo $(X, \leq, |\cdot|_X)$ with maximal order type $o(X) < \omega^{d+1}$ is of length at most $h_{o(X)}(n \cdot d)$, where $h(x) \stackrel{\text{def}}{=} d \cdot g(x)$.*

In Section 3.3, we have already sketched how to extract a (g, n) -controlled bad sequence $\mathbf{v}_0, \mathbf{v}_1, \dots$ of vectors from \mathbb{N}_{ω}^d out of a (g, n) -controlled descending chain $D_0 \supseteq D_1 \supseteq \dots$ of downwards-closed subsets of \mathbb{N}^d . What needs to be shown in order to apply Theorem A.1 to that bad sequence is that we can use a polynomial normed wqo (X, \leq) with $o(X) < \omega^{d+1}$ instead of \mathbb{N}_{ω}^d , and derive the $h_{\omega^{d+1}}(n \cdot d!)$ bound from it. This is a routine application of the results from [26], but we shall give a detailed account for the reader's sake.

A.1. Polynomial Normed WQOs. Let us denote by Γ_0 the empty wqo and by Γ_1 the singleton set $\{\bullet\}$ well-quasi-ordered with equality. A *polynomial* wqo is one that can be constructed from Γ_0 , Γ_1 , and \mathbb{N} through Cartesian products ‘ \times ’ and disjoint unions ‘ $+$ ’, using respectively the product and sum orderings [26, Section 2.1.2].

A *normed* quasi-order associates a *norm* function $|\cdot|_X: X \rightarrow \mathbb{N}$ to a quasi-order (X, \leq) , such that $X_{\leq n} \stackrel{\text{def}}{=} \{x \in X \mid |x|_X \leq n\}$ is finite for every n . For polynomial wqos, we use the zero norm on Γ_1 and the infinity norm on Cartesian products.

A.2. Reflections. A *normed order reflection* between two normed quasi-orders $(X, \leq_X, |\cdot|_X)$ and $(Y, \leq_Y, |\cdot|_Y)$ is a function $r: X \rightarrow Y$ such that, for all x and x' from X ,

- $r(x) \leq_Y r(x')$ implies $x \leq_X x'$, and
- $|r(x)|_Y \leq |x|_X$.

If $(Y, \leq_Y, |\cdot|_Y)$ is a normed wqo and there is a normed order reflection from $(X, \leq_X, |\cdot|_X)$ to it, then $(X, \leq_X, |\cdot|_X)$ is a normed wqo with (g, n) -controlled bad sequences of at most the same length [26, Proposition 2.16].

Observe that r defined by $r(\omega) \stackrel{\text{def}}{=} \bullet$ and $r(n) \stackrel{\text{def}}{=} n$ is a normed order reflection from $(\mathbb{N}_\omega, \leq, \|\cdot\|)$ into $(\mathbb{N} + \Gamma_1, \leq, |\cdot|)$. Hence $(\mathbb{N}_\omega^d, \sqsubseteq, \|\cdot\|)$ reflects into

$$((\mathbb{N} + \Gamma_1)^d, \sqsubseteq, |\cdot|) \quad (27)$$

by [26, Rem. 2.17]. This will be the polynomial normed wqo on which we will apply Theorem A.1.

A.3. Maximal Order Types. It remains to compute the maximal order type of $((\mathbb{N} + \Gamma_1)^d, \sqsubseteq, |\cdot|)$. This can be done algebraically [26, Section 2.4.1] using natural sums ‘ \oplus ’ and natural products ‘ \otimes ’ of maximal order types:

$$o(\mathbb{N} + \Gamma_1) = \omega \oplus 1 = \omega + 1, \quad (28)$$

$$o((\mathbb{N} + \Gamma_1)^d) = \underbrace{(\omega + 1) \otimes \cdots \otimes (\omega + 1)}_{d \text{ times}} = \sum_{d \geq i \geq 0} \omega^i \binom{d}{i}. \quad (29)$$

A.4. Cichoń Functions. Let us recall that, given a monotone expansive $h: \mathbb{N} \rightarrow \mathbb{N}$, and an ordinal α , the α th *Cichoń function* h_α is defined by induction on α by

$$h_0(x) \stackrel{\text{def}}{=} 0, \quad h_\alpha(x) \stackrel{\text{def}}{=} 1 + h_{P_x(\alpha)}(h(x)), \quad (30)$$

where $P_x(\alpha) < \alpha$ denotes the predecessor ordinal at x of α , defined for $0 < \alpha < \varepsilon_0$ by:

$$P_x(\alpha + 1) \stackrel{\text{def}}{=} \alpha, \quad P_x(\gamma + \omega^\beta) \stackrel{\text{def}}{=} \gamma + \omega^{P_x(\beta)} \cdot x + P_x(\omega^{P_x(\beta)}). \quad (31)$$

For instance, $P_x(\omega^2) = \omega \cdot x + P_x(\omega) = \omega \cdot x + x + P_x(1) = \omega \cdot x + x$, and more generally $P_x(\omega^{d+1}) = \sum_{d \geq i \geq 0} \omega^i \cdot x$. Each Cichoń function h_α is monotone.

Proof of Theorem 3.7. By extracting proper ideals and applying Theorem A.1, we obtain an upper bound of

$$\ell = h_{o((\mathbb{N} + \Gamma_1)^d)}(n \cdot d) \quad (32)$$

on the length of (g, n) -controlled bad sequences over \mathbb{N}_ω^d , and thus of $\ell + 1$ on the length of (g, n) -controlled descending chains $D_0 \supsetneq D_1 \supsetneq \dots \supsetneq D_\ell$ of downwards-closed subsets of \mathbb{N}^d . As (29) is quite a mouthful, we are going to over-approximate this bound with a more readable one.

Recall that any ordinal α below ω^{d+1} can be written in Cantor normal form as $\alpha = \omega^d \cdot a_d + \dots + \omega^0 \cdot a_0$ where a_d, \dots, a_0 are coefficients in \mathbb{N} . We can refine the *structural ordering* of [26, Eq. 2.70] for ordinals below ω^{d+1} by:

$$\omega^d \cdot a_d + \dots + \omega^0 \cdot a_0 \sqsubseteq \omega^d \cdot b_d + \dots + \omega^0 \cdot b_0 \text{ if } \forall 1 \leq i \leq d. a_i \leq b_i. \quad (33)$$

By [26, Exercise 2.11], $\alpha \sqsubseteq \beta$ ensures $h_\alpha(n) \leq h_\beta(n)$ for all n .

Observe now that (29) is such that, for all $n > 0$,

$$o((\mathbb{N} + \Gamma_1)^d) \sqsubseteq \sum_{d \geq i \geq 0} \omega^i \cdot nd! = P_{nd!}(\omega^{d+1}). \quad (34)$$

Hence the result stated in Theorem 3.7, by (32) and

$$\begin{aligned} h_{o((\mathbb{N} + \Gamma_1)^d)}(nd) &\leq h_{o((\mathbb{N} + \Gamma_1)^d)}(h(nd!)) && \text{since } nd \leq h(nd!) \\ &\leq h_{P_{nd!}(\omega^{d+1})}(h(nd!)) && \text{by (34)} \\ &= h_{\omega^{d+1}}(nd!) - 1. && \square \end{aligned}$$

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] P. A. Abdulla, A. Bouajjani, and J. d’Orso. Monotonic and downward closed games. *J. Logic Comput.*, 18(1):153–169, 2008. doi:10.1093/logcom/exm062.
- [3] A. Arnold and M. Latteux. Récursivité et cônes rationnels fermés par intersection. *CALCOLO*, 15(4):381–394, 1978. doi:10.1007/BF02576519.
- [4] M. Blondin, A. Finkel, and P. McKenzie. Handling infinitely branching WSTS. In *ICALP 2014*, volume 8573 of *LNCS*, pages 13–25, 2014. doi:10.1007/978-3-662-43951-7_2.
- [5] 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*, Coll. Math. Soc. János Bolyai, pages 189–198. North-Holland, 1975.
- [6] R. Bonnet, A. Finkel, and M. Praveen. Extending the Rackoff technique to affine nets. In *FSTTCS 2012*, volume 18 of *LIPICs*, pages 301–312. LZI, 2012. doi:10.4230/LIPICs.FSTTCS.2012.301.
- [7] L. Bozzelli and P. Ganty. Complexity analysis of the backward coverability algorithm for VASS. In *RP 2011*, volume 6945 of *LNCS*, pages 96–109. Springer, 2011. doi:10.1007/978-3-642-24288-5_10.
- [8] J.-B. Courtois and S. Schmitz. Alternating vector addition systems with states. In *MFCS 2014*, volume 8634 of *LNCS*, pages 220–231. Springer, 2014. doi:10.1007/978-3-662-44522-8_19.
- [9] Ph. de Groote, B. Guillaume, and S. Salvati. Vector addition tree automata. In *LICS 2004*, pages 64–73. IEEE Computer Society, 2004. ISBN 0-7695-2192-4. doi:10.1109/LICS.2004.51.
- [10] 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.
- [11] D. Figueira, S. Figueira, S. Schmitz, and Ph. Schnoebelen. Ackermannian and primitive-recursive bounds with Dickson’s Lemma. In *LICS 2011*, pages 269–278. IEEE Computer Society, 2011. doi:10.1109/LICS.2011.39.

- [12] A. Finkel and J. Goubault-Larrecq. Forward analysis for WSTS, part I: Completions. In *STACS 2009*, volume 3 of *LIPICs*, pages 433–444. LZI, 2009. doi:10.4230/LIPICs.STACS.2009.1844.
- [13] 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.
- [14] J. Goubault-Larrecq. On a generalization of a result by Valk and Jantzen. Research report, LSV, ENS Cachan, 2009. URL http://www.lsv.fr/Publis/RAPPORTS_LSV/PDF/rr-lsv-2009-09.pdf.
- [15] J. Goubault-Larrecq, P. Karandikar, K. Narayan Kumar, and Ph. Schnoebelen. The ideal approach to computing closed subsets in well-quasi-orderings. In preparation, 2016. See [14] for an earlier version.
- [16] J. Kochems and C.-H. L. Ong. Decidable models of recursive asynchronous concurrency. Preprint, 2015. URL <http://arxiv.org/abs/1410.8852>.
- [17] R. Lazić and S. Schmitz. Non-elementary complexities for branching VASS, MELL, and extensions. *ACM Trans. Comput. Logic*, 16(3:20):1–30, 2015. doi:10.1145/2733375.
- [18] R. Lazić and S. Schmitz. The complexity of coverability in ν -Petri nets. Preprint, 2016. URL <https://hal.inria.fr/hal-01265302>.
- [19] P. Lincoln, J. Mitchell, A. Scedrov, and N. Shankar. Decision problems for propositional linear logic. *Ann. Pure App. Logic*, 56(1–3):239–311, 1992. doi:10.1016/0168-0072(92)90075-B.
- [20] R. Lipton. The reachability problem requires exponential space. Technical Report 62, Yale University, 1976.
- [21] R. Majumdar and Z. Wang. Expand, enlarge, and check for branching vector addition systems. In *Concur 2013*, volume 8052 of *LNCS*, pages 152–166. Springer, 2013. doi:10.1007/978-3-642-40184-8_12.
- [22] 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.
- [23] F. Rosa-Velardo and D. 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.
- [24] S. Schmitz. On the computational complexity of dominance links in grammatical formalisms. In *ACL 2010*, pages 514–524. ACL Press, 2010.
- [25] S. Schmitz. Complexity hierarchies beyond Elementary. *ACM Trans. Comput. Theory*, 8(1):1–36, 2016. doi:10.1145/2858784.
- [26] S. Schmitz and Ph. Schnoebelen. Algorithmic aspects of WQO theory. Lecture notes, 2012. URL <http://cel.archives-ouvertes.fr/cel-00727025>.
- [27] Ph. Schnoebelen. Revisiting Ackermann-hardness for lossy counter machines and reset Petri nets. In *MFCS 2010*, volume 6281 of *LNCS*, pages 616–628. Springer, 2010. doi:10.1007/978-3-642-15155-2_54.
- [28] A. Urquhart. The complexity of decision procedures in relevance logic II. *J. Symb. Log.*, 64(4):1774–1802, 1999. doi:10.2307/2586811.

DIMAP, DEPARTMENT OF COMPUTER SCIENCE, UNIVERSITY OF WARWICK, UK

LSV, ENS CACHAN & CNRS & INRIA, UNIVERSITÉ PARIS-SACLAY, FRANCE