

Hyper-Ackermannian Bounds for Pushdown Vector Addition Systems *

Jérôme Leroux

LaBRI, UMR CNRS 5800
University of Bordeaux
Talence, France
jerome.leroux@labri.fr

M. Praveen

LaBRI, UMR CNRS 5800
University of Bordeaux
Talence, France
praveen.manjunatha@labri.fr

Grégoire Sutre

LaBRI, UMR CNRS 5800
University of Bordeaux
Talence, France
gregoire.sutre@labri.fr

Abstract

This paper studies the boundedness and termination problems for vector addition systems equipped with one stack. We introduce an algorithm, inspired by the Karp & Miller algorithm, that solves both problems for the larger class of well-structured pushdown systems. We show that the worst-case running time of this algorithm is hyper-Ackermannian for pushdown vector addition systems. For the upper bound, we introduce the notion of bad nested words over a well-quasi-ordered set, and we provide a general scheme of induction for bounding their lengths. We derive from this scheme a hyper-Ackermannian upper bound for the length of bad nested words over vectors of natural numbers. For the lower bound, we exhibit a family of pushdown vector addition systems with finite but large reachability sets (hyper-Ackermannian).

Categories and Subject Descriptors Theory of computation [*Semantics and reasoning*]: Program reasoning

General Terms Theory

Keywords Vector addition systems, nested words, well-structured pushdown systems, boundedness problem, hyper-Ackermannian complexity.

1. Introduction

Vector addition systems (VAS) are an important class of well-structured transition systems. They are powerful enough to model concurrent infinite-state systems but still many important properties are decidable, like boundedness, termination and reachability. Adding more power to VAS while retaining decidability is an active area of research (see, e.g., [4, 10, 12, 16, 17, 24]). Techniques developed for VAS are often extended to the more general class of

well-structured systems. In this work, we extend the power of VAS by adding one stack. We give decidability and complexity results, and show how these fit in the general scheme of well-structured systems.

It is well known that a finite-state machine equipped with two stacks is Turing powerful. Replacing one of the stacks by multiple stacks over a singleton stack alphabet (and without bottom test) will result in VAS with one stack, which we study here. This model is close to the border of undecidability and very little is known about the decidability of many problems. Here we prove that the boundedness and termination problems are decidable for this model. The basis of these results is the reduced reachability tree, which is also the basis for many results about VAS and well-structured systems. Our main contributions are twofold.

First, we introduce well-structured pushdown systems, of which pushdown VAS are an important class. We extend the well-known *reduced reachability tree* technique for well-structured transition systems [14] to the pushdown case. The reduced reachability tree allows us to decide the boundedness and termination problems for well-structured pushdown systems.

Second, we give hyper-Ackermannian upper and lower bounds for the size of the reduced reachability trees of pushdown VAS. For the upper bound, we introduce the notion of bad nested words over a well-quasi-ordered set, and we provide a general scheme of induction for bounding their lengths. We derive from this scheme a hyper-Ackermannian upper bound for the length of bad nested words over vectors of natural numbers. For the lower bound, we exhibit a family of pushdown vector addition systems with finite but large reachability sets (hyper-Ackermannian).

Many decidability results about well-structured systems depend on the behavior of well-quasi-orders. In the presence of a stack, making these work require new insights, as done in our results mentioned above. While results like Dickson's lemma and Higman's lemma can prove decidability of some algorithms, deriving complexity bounds for such algorithms has recently received attention [9, 13, 18]. Such bounds help to compare the power of different models and understand the structure of the models under study. Again, making these work in the presence of a stack require substantial extensions, as we have done for deriving bounds for the size of the reduced reachability tree. In particular, we extend the framework for deriving bounds on bad sequences [13] to bad nested sequences [3].

*This work was supported by the ANR project REACHARD (ANR-11-BS02-001).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CSL-LICS 2014, July 14–18, 2014, Vienna, Austria.
Copyright © 2014 ACM 978-1-4503-2886-9...\$15.00.
<http://dx.doi.org/10.1145/2603088.2603146>

The upper and lower bounds that we derive in this paper for the size of the reduced reachability tree do not apply for the complexity of the problems themselves, since there may be better algorithms. The best known lower bound is the non-elementary lower bound given by [19]. Very few upper bounds are known and none of them match this lower bound. Closing these gaps are challenging open problems.

Related work. In [5], a model called RVASS is introduced, which allows vector additions during stack operations. However, the semantics of this model resembles that of branching VAS [10] rather than that of a VAS with a stack. In [6], a different version of well-structured pushdown systems are introduced, which are systems where both the set of control states and the stack alphabet can be infinite and are well-quasi-ordered. They consider the coverability problem for subclasses. In contrast, the model we consider here has a finite stack alphabet and we consider termination and boundedness for the whole class. In [2], the authors introduce a model called pushdown automata with data and prove that the control state reachability problem is decidable. In this model again, the set of messages that can be pushed onto the stack can be infinite. The authors introduce a well-quasi-order on this set to decide control state reachability. It is not clear whether the techniques developed for models with infinite stack alphabet can be used in our model and vice-versa. A general framework for well-structured transition systems with auxiliary storage has been introduced in [7]. Decidability is shown for the coverability problem in some subclasses that have a ranking function compatible with the transition relation. Our focus on the other hand is specifically on reduced reachability trees and deriving upper bounds for algorithms based on them. The authors of [7] extend their work in [8] to decide some branching properties of asynchronous programs. They prove that termination and other problems are decidable for our model, provided counters are decreased only when the stack is empty. A similar restriction is imposed by the authors of [26] on their model called multi-set pushdown systems, closely related to our model. We do not impose such restrictions. The fundamental techniques used for deriving Ackermannian and hyper-Ackermannian bounds are developed mostly in [13, 25], on top of which we develop techniques for working with the stack.

Outline. We present in Section 3 our reduced reachability tree technique¹ for well-structured pushdown systems. Section 4 provides a general scheme of induction for bounding the length of bad nested words. Basic notions on fast-growing functions are recalled in Section 5. We derive in Section 6 a multiply-recursive upper bound for the length of bad nested sequences over $\mathbb{N}^d$. Section 7 shows that the worst-case size of the reduced reachability tree is hyper-Ackermannian for pushdown VAS.

2. Preliminaries

A *quasi-order* on a set S is a binary relation $\preceq$ on S that is reflexive and transitive. As usual, we write $s \prec t$ when $s \preceq t$ and $t \not\preceq s$. A *partial-order* is an antisymmetric quasi-order.

A *well-quasi-order* (wqo) on S is a quasi-order $\preceq$ such that every infinite sequence $s_0, s_1, s_2, \dots$ in S contains an infinite nondecreasing subsequence $s_{i_0} \preceq s_{i_1} \preceq s_{i_2} \dots$ (with $i_0 < i_1 < i_2 \dots$).

¹This section motivates the developments of the following sections, but its results are only used in Section 7.

The concatenation of a finite sequence (resp. word) σ with a finite or infinite one w is denoted by σ, w (resp. $\sigma \cdot w$). The length of a finite sequence or word σ is denoted $|\sigma|$.

3. Well-Structured Pushdown Systems

A *pushdown system* is a 4-tuple $\mathcal{P} = \langle S, s_{\text{init}}, \Gamma, \Delta \rangle$ where S is a set of *states*, $s_{\text{init}} \in S$ is the *initial state*, Γ is a finite *stack alphabet*, and $\Delta \subseteq S \times \text{Op} \times S$ is a set of *rules*. The set Op of *operations* is given by $\text{Op} = \{\varepsilon\} \cup \{\text{push}(\gamma), \text{pop}(\gamma) \mid \gamma \in \Gamma\}$. Observe that we do *not* require S to be finite. Informally, a rule (s, op, t) means that the pushdown system can move from state s to state t by performing the operation op on the stack. We say that $\mathcal{P}$ is *finitely-branching* if the set $\{(\text{op}, t) \mid (s, \text{op}, t) \in \Delta\}$ is finite for every state $s \in S$. A *simulation relation* for $\mathcal{P}$ is any binary relation R on S satisfying the following condition: for every $(s, \text{op}, t) \in \Delta$ and $s' \in S$ such that $(s, s') \in R$, there exists $t' \in S$ such that $(s', \text{op}, t') \in \Delta$ and $(t, t') \in R$.

We define the operational semantics of a pushdown system $\mathcal{P} = \langle S, s_{\text{init}}, \Gamma, \Delta \rangle$ as follows. A *configuration* is a pair $\chi = (s, \sigma)$ where $s \in S$ is a state and $\sigma \in \Gamma^*$ is a word denoting the contents of the stack. The *initial configuration* is $\chi_{\text{init}} = (s_{\text{init}}, \varepsilon)$. The *transition relation* is the least binary relation $\rightarrow \subseteq (S \times \Gamma^*) \times (S \times \Gamma^*)$ satisfying the following rules:

$$\frac{(s, \varepsilon, t) \in \Delta}{(s, \sigma) \rightarrow (t, \sigma)}$$

$$\frac{(s, \text{push}(\gamma), t) \in \Delta}{(s, \sigma) \rightarrow (t, \sigma \cdot \gamma)} \quad \frac{(s, \text{pop}(\gamma), t) \in \Delta}{(s, \sigma \cdot \gamma) \rightarrow (t, \sigma)}$$

A *run* is a (finite or infinite) sequence $\chi_0, \chi_1, \chi_2, \dots$ of configurations such that $\chi_0 = \chi_{\text{init}}$ and $\chi_{i-1} \rightarrow \chi_i$ for every index $i > 0$ of the sequence. The *reachability set* is the set of all configurations that occur on some run.

We investigate two verification questions on pushdown systems, namely *termination* and *boundedness*. The former asks whether all runs of a given pushdown system are finite, and the latter asks whether its reachability set is finite. Recall that these questions are decidable for pushdown systems with finitely many states. In our setting, the set of states may be infinite (e.g., $Q \times \mathbb{N}^d$ for pushdown vector addition systems with states, see Section 7). By extending well-known techniques from the theory of well-structured transition systems [1, 14, 15], we obtain a class of pushdown systems for which termination and boundedness are decidable.

Definition 3.1. A well-structured pushdown system is a *finitely-branching pushdown system* $\mathcal{P} = \langle S, s_{\text{init}}, \Gamma, \Delta \rangle$ equipped with a well-quasi-order $\preceq \subseteq S \times S$ on states such that $\preceq$ is a simulation relation for $\mathcal{P}$.

Remark 3.2. The monotonicity condition of Definition 3.1 matches the one given in [1] for infinite-state *labeled* transition systems. In our setting, the labels of [1] are operations $\text{op} \in \text{Op}$. More general monotonicity conditions have been proposed in [15] for *unlabeled* infinite-state transition systems. The condition of Definition 3.1 corresponds to the *strong compatibility* requirement of [15].

The *reachability tree* of a well-structured pushdown system $\mathcal{P} = \langle S, s_{\text{init}}, \Gamma, \Delta, \preceq \rangle$ is the directed unordered tree defined as follows. Nodes of the tree are labeled by configurations of $\mathcal{P}$. The root r is labeled by the initial configuration χ_{init} , written $r : \chi_{\text{init}}$. Each node $n : \chi$ has one child $m : \psi$ for

each configuration ψ such that $\chi \rightarrow \psi$. Note that the reachability tree of $\mathcal{P}$ is finitely branching.

The reachability tree is usually infinite for well-structured pushdown systems. We show how to truncate it while preserving enough information to decide termination and boundedness. A node $n : (s, \sigma)$ *subsumes* a node $m : (t, \tau)$ if n is a proper ancestor of m , $s \preceq t$ and σ is a prefix of σ' for every node $n' : (s', \sigma')$ on the path from n to m . The *reduced reachability tree* of $\mathcal{P}$ is the largest prefix of the reachability tree such that every subsumed node has no child.

Proposition 3.3. *The reduced reachability tree of any well-structured pushdown system is finite.*

Proof. Consider a well-structured pushdown system $\mathcal{P} = \langle S, s_{\text{init}}, \Gamma, \Delta, \preceq \rangle$. By contradiction, assume that its reduced reachability tree $\mathcal{T}$ is infinite. Recall that this tree is finitely branching. By König's lemma, $\mathcal{T}$ contains an infinite branch $n_0 : (s_0, \sigma_0) \rightarrow n_1 : (s_1, \sigma_1) \rightarrow n_2 : (s_2, \sigma_2) \rightarrow \dots$. Observe that $\sigma_0 = \varepsilon$ since n_0 is the root of the tree. We extract from $n_0, n_1, n_2, \dots$ an infinite subsequence $n_{i_0}, n_{i_1}, n_{i_2}, \dots$ defined as follows:

$$\begin{aligned} i_0 &= 0 \\ i_{k+1} &= \min \{i > i_k \mid \forall j > i_k, |\sigma_i| \leq |\sigma_j|\} \end{aligned}$$

By definition, $|\sigma_{i_k}| \leq |\sigma_j|$ for all $j, k \in \mathbb{N}$ with $i_k \leq j$. It follows from the operational semantics of pushdown systems that σ_{i_k} is a prefix of σ_j for all $j, k \in \mathbb{N}$ with $i_k \leq j$. Since $\preceq$ is a well-quasi-order, there exists $k < l$ such that $\sigma_{i_k} \preceq \sigma_{i_l}$. We obtain that the node n_{i_k} subsumes the node n_{i_l} . This entails that n_{i_l} has no child in $\mathcal{T}$, which is impossible. $\square$

Proposition 3.4. *A well-structured pushdown system $\mathcal{P}$ has an infinite run if, and only if, its reduced reachability tree contains a subsumed node.*

Proof. If $\mathcal{P}$ has an infinite run, then its reachability tree contains an infinite branch. Since the reduced reachability tree of $\mathcal{P}$ is finite, this infinite branch necessarily contains a subsumed node. Hence, the reachability tree of $\mathcal{P}$ contains a subsumed node, and so does its reduced reachability tree. For the converse, assume that the reachability tree of $\mathcal{P}$ contains a subsumed node. The path from the root to this subsumed node yields a run

$$(s_0, \sigma_0) \xrightarrow{\text{op}_1} \dots \xrightarrow{\text{op}_k} (s_k, \sigma_k) \xrightarrow{\text{op}_{k+1}} \dots \xrightarrow{\text{op}_l} (s_l, \sigma_l)$$

such that $k < l$, $s_k \preceq s_l$ and σ_k is a prefix of σ_j for all $k \leq j \leq l$. Here, we have decorated each transition of the run with its associated operation. Let $z_k, \dots, z_l$ be the words satisfying $\sigma_j = \sigma_k \cdot z_j$ for all $k \leq j \leq l$. Define $(t_k, \tau_k) = (s_l, \sigma_l)$. Since $s_k \preceq t_k$ and $\preceq$ is a simulation relation for $\mathcal{P}$, we derive that there exists $t_{k+1}, \dots, t_l$ such that, for every $k < j \leq l$, $s_j \preceq t_j$ and $(t_{j-1}, \text{op}_j, t_j) \in \Delta$. Observe that the sequence of operations $\text{op}_{k+1}, \dots, \text{op}_l$ never pops from σ_k . Hence, we may also perform $\text{op}_{k+1}, \dots, \text{op}_l$ starting from $\sigma_l = \tau_k \cdot z_k$. So the above run may be extended by appending

$$(s_l, \sigma_l) = (t_k, \tau_k \cdot z_k) \xrightarrow{\text{op}_{k+1}} \dots \xrightarrow{\text{op}_l} (t_l, \tau_k \cdot z_l)$$

We obtain an infinite run by repeating this construction. $\square$

The boundedness problem requires stronger monotonicity requirements. A well-structured pushdown system $\mathcal{P} = \langle S, s_{\text{init}}, \Gamma, \Delta, \preceq \rangle$ is called *strict* when $\preceq$ is a partial order and $\prec$ is a simulation relation. Similarly, in the reachability tree of $\mathcal{P}$, a node $n : (s, \sigma)$ *strictly subsumes* a node $m : (t, \tau)$ if n subsumes m , and $s \prec t$ or $|\sigma| < |\tau|$.

Proposition 3.5. *A strict well-structured pushdown system $\mathcal{P}$ has an infinite reachability set if, and only if, its reduced reachability tree contains a strictly subsumed node.*

Proof. Assume that $\mathcal{P}$ has an infinite reachability set. Let $\mathcal{T}$ be the largest prefix of its reachability tree such that, on each branch, all nodes have distinct labels. The tree $\mathcal{T}$ is infinite since every configuration in the reachability set is the label of some node in $\mathcal{T}$. By König's lemma, it follows that the reachability tree of $\mathcal{P}$ contains an infinite branch where all nodes have distinct labels. Since the reduced reachability tree of $\mathcal{P}$ is finite, this infinite branch necessarily contains two nodes $n : (s, \sigma)$ and $m : (t, \tau)$, both contained in the reduced reachability tree, such that n subsumes m . Recall that $(s, \sigma) \neq (t, \tau)$. If $|\sigma| < |\tau|$ then n strictly subsumes m . Otherwise, $\sigma = \tau$ since σ is a prefix of τ . It follows that $s \neq t$, which entails that $s \prec t$ as $\preceq$ is a partial order. We get, again, that n strictly subsumes m .

For the converse, assume that the reduced reachability tree of $\mathcal{P}$ contains a strictly subsumed node. The path from the root to this subsumed node yields a run

$$(s_0, \sigma_0) \xrightarrow{\text{op}_1} \dots \xrightarrow{\text{op}_k} (s_k, \sigma_k) \xrightarrow{\text{op}_{k+1}} \dots \xrightarrow{\text{op}_l} (s_l, \sigma_l)$$

such that $k < l$, $s_k \preceq s_l$, σ_k is a prefix of σ_j for all $k \leq j \leq l$, and $s_k \prec s_l$ or $|\sigma_k| < |\sigma_l|$. With the same arguments as in the proof of Proposition 3.4, we extend this run into an infinite one by iterating the sequence of operations $\text{op}_{k+1}, \dots, \text{op}_l$. If $s_k \prec s_l$ then the resulting infinite run visits infinitely many states (as $\preceq$ is a partial order). If $|\sigma_k| < |\sigma_l|$ then it visits infinitely many stack contents. In both cases, we obtain that $\mathcal{P}$ has an infinite reachability set. $\square$

4. Bad Nested Sequences Over a WQO Set

Nested words [3] offer a convenient formalism to study some properties of pushdown systems. Since we are dealing with sets of reachable configurations, our notation is slightly different from that of [3], where the main concern is the theory of languages. For us, an alphabet S is usually the set of states of a pushdown system, and the content of the stack is abstracted by its length.

Definition 4.1. *A nested sequence over a set S is a (finite or infinite) sequence $(s_0, h_0), (s_1, h_1), \dots$ of elements in $S \times \mathbb{N}$ satisfying $h_0 = 0$ and $h_j \in h_{j-1} + \{-1, 0, 1\}$ for every index $j > 0$ of the sequence.*

Example 4.2. The empty sequence is a nested sequence. The nested sequence $(a, 0), (b, 1), (c, 2), (d, 1)$ over $\{a, b, c, d\}$ can be depicted as follows:

$$\begin{array}{cccc} 2- & & c & \\ 1- & & b & d \\ 0- & a & & \end{array}$$

Remark 4.3. In a well-structured pushdown system, every run can be seen as a nested sequence, by mapping each configuration (s, σ) to the pair $(s, |\sigma|)$. Similarly, every branch of the reachability tree can be seen as a nested sequence.

Definition 4.4. *A nested sequence $(s_0, h_0), (s_1, h_1), \dots$ over a quasi-ordered set $(S, \preceq)$ is said to be good if there exists $i < j$ such that $s_i \preceq s_j$ and $h_{i+1}, \dots, h_j \geq h_i$. A bad nested sequence is one that is not good.*

Example 4.5. The following nested sequence over $(\mathbb{N}, \leq)$ is bad:

```

3-      0
2-      1      0      0
1-      2      1      1      0      0
0-  3      2      1      0

```

Remark 4.6. A nested sequence $(s_0, 0), (s_1, 0), \dots$ is good iff $s_i \leq s_j$ for some $i < j$. Hence good/bad nested sequences are generalizations of good/bad (non-nested) sequences [13].

Remark 4.7. Consider a branch in the reachability tree of some well-structured pushdown system. The nested sequence associated with the branch is good if, and only if, the branch contains a subsumed node.

Lemma 4.8. *Bad nested sequences over a wqo are finite.*

Proof. Similar to the proof of Proposition 3.3. $\square$

The finiteness of the reduced reachability tree presented in the previous section relies on the finiteness of bad nested sequences. In order to derive complexity arguments, bounds on these sequences must be obtained. *Normed* wqo sets provide a framework for deriving such bounds. We first present some simple examples of (non-nested) bad sequences that motivate normed wqo sets.

The following nested sequence shows that the length of bad nested sequences over a wqo cannot be bounded since there exist bad nested sequences over $(\mathbb{N}, \leq)$ of length n for every n .

```

0-  100  99  98  ...  1  0

```

Hence, the initial element must be bounded somehow. In the algorithms presented in the previous section, nested sequences starts from the same element, the initial state of the well-structured pushdown system. However, even if the initial element is fixed, the problem is not yet solved as shown by the following bad nested sequence over $(\mathbb{N}^2, \leq)$.

```

0-  (0, 1)  (100, 0)  (99, 0)  (98, 0)  ...  (1, 0)  (0, 0)

```

This problem appears because the second element of the sequence is arbitrarily large compared to the first one. In order to control the growth of values in a nested sequence, we use the notion of *normed* wqo sets.

Definition 4.9. A norm for a wqo set $(S, \preceq)$ is a function $\|\cdot\| : S \mapsto \mathbb{N}$ such that $\{s \in S \mid \|s\| \leq n\}$ is finite for every n . The structure $(S, \preceq, \|\cdot\|)$ is called a normed wqo.

Example 4.10. The wqo set $(\mathbb{N}^d, \leq)$ is normed by the function that maps any vector to its largest component.

Definition 4.11. A nested sequence $(s_0, h_0), (s_1, h_1), \dots$ over a normed wqo set $(S, \preceq, \|\cdot\|)$ is said to be n -controlled for some $n \in \mathbb{N}$ if $\|s_j\| \leq n + j$ for every index j of the sequence.

While there exist bad nested sequences of arbitrarily long lengths, König's lemma shows that there exists a maximal length for n -controlled bad nested sequences. Indeed, these sequences can be organized in a forest where nodes are labeled by elements in $S \times \mathbb{N}$. By definition of the norm function, this forest is finite branching and the number of roots is finite. Hence, if the forest is infinite, the König's lemma shows that it contains an infinite branch. Since infinite nested sequences are good we get a contradiction.

Let $\text{BAD}_S(n)$ be the set of n -controlled bad nested sequences over S . This set is non empty since it contains the

empty sequence. The *maximal length* function $L_S : \mathbb{N} \mapsto \mathbb{N}$ is defined as follows:

$$L_S(n) = \max\{|w| \mid w \in \text{BAD}_S(n)\}$$

The maximal length function L_S satisfies a recurrence relation that we call the *descent equation*. For this, we need the set $S|_{\leq n}$ defined as follows:

$$S|_{\leq n} = \{s \in S \mid \|s\| \leq n\}$$

Given $s \in S|_{\leq n}$, let us denote by $\text{BAD}_{S,s}(n)$ the set of n -controlled bad nested sequences starting from $(s, 0)$. This set is non-empty since it contains $(s, 0)$. Let $L_{S,s}(n)$ be the number defined as follows:

$$L_{S,s}(n) = \max\{|w| \mid w \in \text{BAD}_{S,s}(n)\}$$

The following equality is immediate for every $n \in \mathbb{N}$ such that $S|_{\leq n}$ is non empty:

$$L_S(n) = \max_{s \in S|_{\leq n}} L_{S,s}(n)$$

Let $(S/s, \preceq, \|\cdot\|)$ be the normed wqo set where S/s (called the *quotient* of S by $s \in S$) is defined as $S/s = \{t \in S \mid s \not\preceq t\}$. In the following, we prove the following equality, called the *descent equation* for every $n \in \mathbb{N}$ and for every $s \in S|_{\leq n}$:

$$L_{S,s}(n) = 1 + L_{S/s}(n+1) + L_{S/s}(n+1 + L_{S/s}(n+1))$$

Let us first prove that $L_{S,s}(n)$ is bounded by the above term. We pick a n -controlled bad nested sequence w over S starting from $(s, 0)$ with $s \in S|_{\leq n}$. The nested sequence w can be decomposed into $w = (s, 0)w'$ where w' is a sequence of elements in $S \times \mathbb{N}$. Since w is bad, w' is in fact a sequence of elements in $(S/s) \times \mathbb{N}$. Let us introduce the maximal prefix w'' of w' such that w'' is a sequence of elements in $S \times (\mathbb{N} \setminus \{0\})$. Let us show that $|w'| \leq L_{S/s}(n+1)$. The sequence w'' can be decomposed into $w'' = (s_1, h_1), \dots, (s_k, h_k)$ where $k = |w''|$. Note that $(s_1, h_1 - 1), \dots, (s_k, h_k - 1)$ is a $(n+1)$ -controlled bad nested sequence over the wqo set S/s . Thus, $|w''| \leq L_{S/s}(n+1)$. Now, let us consider the sequence w''' such that $w' = w'', w'''$. Let us prove that $|w'''| \leq L_{S/s}(n+1+k)$. Observe that if w''' is empty the relation is immediate. So, let us assume that w''' is non empty. By maximality of w'' , we deduce that w''' starts with an element in $S \times \{0\}$. Thus w''' is a $(n+1+k)$ -controlled bad nested sequence over S/s . We get $|w'''| \leq L_{S/s}(n+1+k)$. Since $L_{S/s}$ is monotonic and $k \leq L_{S/s}(n+1)$, we deduce that $|w''|$ is bounded by $L_{S/s}(n+1 + L_{S/s}(n+1))$. Hence $|w|$ is bounded by $1 + L_{S/s}(n+1) + L_{S/s}(n+1 + L_{S/s}(n+1))$. We have proved one inequality of the descent equation, namely that $L_{S,s}(n)$ is bounded by $1 + L_{S/s}(n+1) + L_{S/s}(n+1 + L_{S/s}(n+1))$ for every $n \in \mathbb{N}$ and for every $s \in S|_{\leq n}$.

To prove the converse inequality, we need to concatenate bad nested sequences. The concatenation will be possible thanks to the following lemma.

Lemma 4.12. *Every $w \in \text{BAD}_S(n)$ with $|w| = L_S(n)$ is either empty or the last element is in $S \times \{0\}$.*

Proof. If $L_S(n) = 0$, then the proof is immediate. Assume that $L_S(n) > 0$ and let us consider a n -controlled bad nested sequence w such that $|w| = L_S(n)$. Let (s, h) be the last element of w . By contradiction, suppose that $h > 0$. It is readily seen that the sequence obtained by appending $(s, h-1)$ to w is also a n -controlled bad nested sequence. This contradicts the maximality of w . We derive that $h = 0$, which concludes the proof of the lemma. $\square$

Let us now complete the proof of the descent equation, by showing that $1 + L_{S/s}(n+1) + L_{S/s}(n+1 + L_{S/s}(n+1))$ is bounded by $L_{S,s}(n)$ for every $n \in \mathbb{N}$ and for every $s \in S|_{\leq n}$. Pick $n \in \mathbb{N}$ and $s \in S|_{\leq n}$. Let $n_1 = L_{S/s}(n+1)$. There exists $w_1 \in \text{BAD}_{S/s}(n+1)$ with $|w_1| = n_1$. According to Lemma 4.12, if $n_1 > 0$ then the last element of w_1 is in $S \times \{0\}$. Let $n_2 = L_{S/s}(n+1 + n_1)$. By definition, there exists w_2 in $\text{BAD}_{S/s}(n+1+n_1)$ with $|w_2| = n_2$. We introduce the nested sequence $w = (s, 0), w'_1, w_2$ where w'_1 is obtained from w_1 by replacing every occurrence of $(t, h) \in S \times \mathbb{N}$ by $(t, h+1)$. Observe that w is a bad nested sequence that is n -controlled. Since $w \in \text{BAD}_{S,s}(n)$, we get $L_{S,s}(n) \geq |w|$. It follows that $1 + L_{S/s}(n+1) + L_{S/s}(n+1 + L_{S/s}(n+1))$ is bounded by $L_{S,s}(n)$, which concludes the proof of the descent equation.

In order to simplify the presentation of the descent equation, we introduce the function $K_S : \mathbb{N} \mapsto \mathbb{N}$ defined for every $n \in \mathbb{N}$ as follows:

$$K_S(n) = n + L_S(n)$$

Intuitively, $K_S(n)$ explicitly accounts for the starting value n by adding it to $L_S(n)$. It is understood that, like L_S , the function K_S is parametrized by a normed wqo set $(S, \preceq, \|\cdot\|)$.

Now, just observe that for every $n \in \mathbb{N}$ and every $s \in S|_{\leq n}$, the descent equation entails that:

$$\begin{aligned} n + L_{S,s}(n) &= n + 1 + L_{S/s}(n+1) + L_{S/s}(n+1 + L_{S/s}(n+1)) \\ &= K_{S/s}(K_{S/s}(n+1)) \end{aligned}$$

In particular, if $S|_{\leq n}$ is non empty, we obtain that:

$$\begin{aligned} K_S(n) &= n + L_S(n) \\ &= n + \max_{s \in S|_{\leq n}} L_{S,s}(n) \\ &= \max_{s \in S|_{\leq n}} n + L_{S,s}(n) \\ &= \max_{s \in S|_{\leq n}} K_{S/s}(K_{S/s}(n+1)) \end{aligned}$$

We have proved the following theorem.

Theorem 4.13. *For every normed wqo set $(S, \preceq, \|\cdot\|)$ and for every $n \in \mathbb{N}$, we have:*

$$K_S(n) = \begin{cases} \max_{s \in S|_{\leq n}} K_{S/s}(K_{S/s}(n+1)) & \text{if } S|_{\leq n} \neq \emptyset \\ n & \text{otherwise} \end{cases}$$

5. Fast-Growing Functions

The next sections classify the maximal length function in the *fast-growing hierarchy* [20]. We shall work with set-theoretic ordinals less than ω^ω , written in Cantor's Normal Form. Each limit ordinal λ has a *canonical fundamental sequence* $(\lambda_n)_{n \in \mathbb{N}}$, satisfying $\lambda_0 < \lambda_1 < \dots < \lambda_n < \lambda_{n+1} < \dots$, with λ being the supremum of this sequence. For limit ordinals below ω^ω , this fundamental sequence is given by

$$\begin{aligned} (\omega^d \cdot a_d + \dots + \omega^r \cdot a_r)_n &= \\ \omega^d \cdot a_d + \dots + \omega^{r+1} \cdot a_{r+1} + \omega^r \cdot (a_r - 1) + \omega^{r-1} \cdot (n+1) \end{aligned}$$

assuming $d \geq r$ and $a_r > 0$.

The family of *fast-growing functions* $(F_\alpha)_\alpha$, indexed by ordinals $\alpha \leq \omega^\omega$, is defined as follows:

$$F_0(n) = n + 1,$$

$$F_{\alpha+1}(n) = F_\alpha^{n+1}(n) = \overbrace{F_\alpha(F_\alpha(\dots F_\alpha(n) \dots))}^{n+1 \text{ times}}$$

$$F_\lambda(n) = F_{\lambda_n}(n) \quad \text{if } \lambda < \omega^\omega \text{ is a limit ordinal}$$

$$F_{\omega^\omega}(n) = F_{\omega^{n+1}}(n)$$

The value $F_1(n)$ is linear in n , $F_2(n)$ is exponential in n , and $F_3(n)$ is a tower of exponential of height n . The function F_ω is an Ackermannian function. This function is defined by diagonalization over the ordinals $\alpha < \omega$, since $F_\omega(n) = F_{n+1}(n)$. The hyper-Ackermannian function F_{ω^ω} used in this paper is defined in the same fashion, by diagonalization over the ordinals $\alpha < \omega^\omega$.

Example 5.1. The above definition induces:

$$\begin{aligned} F_{\omega^{2.4}}(8) &= F_{\omega^{2.3+\omega.9}}(8) \\ &= F_{\omega^{2.3+\omega.8+9}}(8) \\ &= \underbrace{F_{\omega^{2.3+\omega.8+8}}(\dots(F_{\omega^{2.3+\omega.8+8}}(8))\dots)}_{9 \text{ times}} \end{aligned}$$

In addition to the usual linear order $\leq$ on ordinals, we will also use a partial order, called *embedding* in [9], and defined as follows. For two ordinals α, β such that $\alpha = \omega^d \cdot a_d + \dots + \omega^0 \cdot a_0$ and $\beta = \omega^d \cdot b_d + \dots + \omega^0 \cdot b_0$, α *embeds* in β if $a_i \leq b_i$ for all $i \in \{0, \dots, d\}$. We write $\alpha \sqsubseteq \beta$ when α embeds in β . We will use the following monotonicity properties of fast-growing functions. These properties are shown by induction over ordinals

Lemma 5.2 (Monotonicity, [9]). *For all ordinals $\alpha, \beta < \omega^\omega$ and for all natural numbers $m, n \in \mathbb{N}$,*

$$\alpha \sqsubseteq \beta \text{ and } m \leq n \quad \Rightarrow \quad m < F_\alpha(m) \leq F_\beta(n)$$

Proof. Our definition of fast-growing functions, which comes from [25], slightly deviates from [9]. The difference lies in the definition of the canonical fundamental sequence. The last term $\omega^{r-1} \cdot (n+1)$ in our definition is replaced by $\omega^{r-1} \cdot n$ in [9]. However, the latter definition would introduce unnecessary complications in the following section, as we would have to distinguish the cases where n is zero or positive. This the reason why we choose the definition of [25]. Still, the proofs of Lemmas 2.1 and 2.2 in [9] can be easily adapted to our setting for proving this lemma. $\square$

6. Bad Nested Sequences Over $\mathbb{N}^d$

In this section, we derive upper bounds for the length of bad nested sequences over multiple disjoint copies of $\mathbb{N}^d$. The sum of $c \in \mathbb{N}$ distinct copies of $\mathbb{N}^d$ is the Cartesian product $\{1, \dots, c\} \times \mathbb{N}^d$, denoted by $c \times \mathbb{N}^d$. This set naturally corresponds to states of vector addition systems with c locations and d counters ranging over the natural numbers. The set $c \times \mathbb{N}^d$ is equipped with the wqo $\preceq_{c \times \mathbb{N}^d}$ defined by $(p, \mathbf{u}) \preceq_{c \times \mathbb{N}^d} (q, \mathbf{v})$ if $p = q$ and $\mathbf{u} \leq \mathbf{v}$, and the norm $\|\cdot\|_{c \times \mathbb{N}^d}$ defined by $\|(q, \mathbf{v})\|_{c \times \mathbb{N}^d} = \|\mathbf{v}\|_\infty$, where $\|\mathbf{v}\|_\infty$ maps $\mathbf{v}$ to its largest component (to zero for the empty vector). In this section, we prove the following theorem.

Theorem 6.1. *For every $c, d \geq 1$ and $n \geq 2$, we have:*

$$L_{c \times \mathbb{N}^d}(n) \leq F_{\omega^d \cdot c}(d \cdot n)$$

Remark 6.2. The previous theorem requires $d \geq 1$. When $d = 0$, the norm of any element in $c \times \mathbb{N}^0$ is zero, so the value

$L_{c \times \mathbb{N}^0}(n)$ does not depend on n . Let us denote by k_c this value. As the quotient of $c \times \mathbb{N}^0$ by an element $v \in c \times \mathbb{N}^0$ is equal to $(c-1) \times \mathbb{N}^0$, Theorem 4.13 shows that $k_c = 2k_{c-1} + 1$ for every $c \geq 1$. From $k_0 = 0$, we derive $k_c = 2^c - 1$ which proves the following equality:

$$L_{c \times \mathbb{N}^0}(n) = 2^c - 1$$

The rest of this section is the technical part for proving Theorem 6.1. In particular, we assume that $d \geq 1$. Our approach is based on the descent equation introduced in Theorem 4.13. This approach introduces sums of multiple disjoint copies of sets $\mathbb{N}^i$ for various dimensions $i \in \{0, \dots, d\}$. Let us show this property on the normed wqo set $\mathbb{N}^d$ (no multiple copy). The descent equation introduces the quotients $\mathbb{N}^d / \mathbf{v}$ where $\mathbf{v} \in \mathbb{N}^d$. A vector $\mathbf{x}$ in such a quotient satisfies $\mathbf{x}(j) \in \{0, \dots, \mathbf{v}(j) - 1\}$ for some index j . Hence, the quotient $\mathbb{N}^d / \mathbf{v}$ can be embedded into $(\mathbf{v}(1) + \dots + \mathbf{v}(d)) \times \mathbb{N}^{d-1}$. By applying recursively the descent equation, starting from $c \times \mathbb{N}^d$, finite sums of disjoint copies of sets $\mathbb{N}^i$ for various dimensions $i \in \{0, \dots, d\}$ are introduced.

Those finite sums are denoted formally thanks to the notion of *types* introduced in [13]. We refer to this paper for further explanations. A type is a function $\tau : \{0, \dots, d\} \mapsto \mathbb{N}$, also viewed as a multiset. We associate to a type τ the set $\mathbb{N}^\tau$ defined as follows:

$$\mathbb{N}^\tau = \bigcup_{i=0}^d \tau(i) \times \mathbb{N}^i$$

The set $\mathbb{N}^\tau$ is equipped with a wqo $\preceq_\tau$ and a norm $\|\cdot\|_\tau$ by observing that the sets $\tau(i) \times \mathbb{N}^i$ are pairwise disjoint. More formally, the wqo $\preceq_\tau$ on $\mathbb{N}^\tau$ is defined by $u \preceq_\tau v$ if there exists i such that $u, v \in \tau(i) \times \mathbb{N}^i$ and $u \preceq_{\tau(i) \times \mathbb{N}^i} v$, and the norm $\|\cdot\|_\tau$ is defined by $\|v\|_\tau = \|v\|_{\tau(i) \times \mathbb{N}^i}$ where i is such that $v \in \tau(i) \times \mathbb{N}^i$.

Since the structure $(\mathbb{N}^\tau, \preceq_\tau, \|\cdot\|_\tau)$ is a normed wqo, the function $K_{\mathbb{N}^\tau}$ is well defined. To simplify notations, we denote this function by K_τ . In the sequel, we assume that the reader is familiar with multiset notations. In particular, since a type is a multiset of natural numbers in $\{0, \dots, d\}$, we denote by $\emptyset$ the empty type that maps any i on zero. Given a type τ such that $\tau(i) \geq 1$, we denote by $\tau - \{i\}$ the type ϱ defined by $\varrho(i) = \tau(i) - 1$ and $\varrho(j) = \tau(j)$ if $j \neq i$.

Following the approach introduced in [13] (explained in the beginning of this section on the quotients $\mathbb{N}^d / \mathbf{v}$), the quotients $\mathbb{N}^\tau / v$, where $v \in \mathbb{N}^\tau$, can be embedded into $\mathbb{N}^\varrho$ for types ϱ derived from τ and v . These types are defined as follows. Given a type τ , a dimension $i \in \{0, \dots, d\}$ such that $\tau(i) \geq 1$ and a natural number $n \geq 0$, the derivation type $\partial_n^i(\tau)$ is defined by:

$$\partial_n^i(\tau) = \tau - \{i\} + n \cdot \sum_{0 \leq j < i} \{j\}$$

We obtain the following lemma, whose proof (which can be found in the full version of the paper) is similar to that of [13, Lemma A.3]:

Lemma 6.3. *It holds that:*

$$K_{\mathbb{N}^\tau / v}(m) \leq K_{\partial_n^i(\tau)}(m)$$

for every type τ , for every $i \in \{1, \dots, d\}$, for every $v \in \tau(i) \times \mathbb{N}^i$, for every $n \geq \|v\|_\tau$, and for every $m \in \mathbb{N}$.

Combining this lemma with Theorem 4.13 and the fact that K_τ is monotonic for all τ , we get the following inequalities,

which hold for every natural number $n \geq 0$, and for every type $\tau \neq \emptyset$:

$$K_\tau(n) \leq \max_{i: \tau(i) \geq 1} \{K_{\partial_n^i(\tau)}(K_{\partial_n^i(\tau)}(n+1))\}$$

For a type τ , we denote by F_τ the function defined by:

$$F_\tau = F_{\omega^d \cdot \tau(d) + \dots + \omega^0 \cdot \tau(0)}$$

We associate to a type $\tau \neq \emptyset$ and a natural number $n \in \mathbb{N}$ the type $\tau_n = \partial_n^i(\tau)$ where i is the minimal index such that $\tau(i) > 0$. An immediate induction shows the following equality for every type $\tau \neq \emptyset$:

$$F_\tau(n) = F_{\tau_n}^{n+1}(n)$$

In the sequel, we prove that for every type τ , for every $n \geq 2$ we have:

$$K_\tau(n) \leq F_\tau(d \cdot n) \quad (1)$$

Note that Theorem 6.1 is a direct consequence of this inequality with the type $\tau = c \cdot \{d\}$.

We write $\tau \sqsubseteq \varrho$ if $\tau(i) \leq \varrho(i)$ for all i . We will use the lexicographic order $<$ on types, with $\tau(i+1)$ being more significant than $\tau(i)$ for all i . This is a well-founded linear order, which is the same as the order used for types in [11, 13]. We first prove some results on the family (F_τ) .

Lemma 6.4. *For every type τ , for every $i \in \{0, \dots, d\}$ such that $\tau(j) > 0$ implies $j \geq i$, and for every $n \geq 0$:*

$$F_{\tau + \{i\}}(n) \geq F_\tau(2n)$$

Proof. The case $n = 0$ comes from $F_\varrho(0) = 1$ for every type ϱ (the equality is obtained by induction on ϱ with the lexicographic order). We introduce $\varrho = \tau + \{i\}$. Since $\tau(j) > 0$ implies $j \geq i$, we have $\varrho_n = \tau + n \cdot (\{i-1\} + \dots + \{0\})$. Lemma 5.2 provides $F_{\varrho_n}^n(n) \geq F_\varrho^n(n) = 2n$. We have:

$$\begin{aligned} F_\varrho(n) &= F_{\varrho_n}^{n+1}(n) \\ &\geq F_{\varrho_n}(F_\varrho^n(n)) && [\text{Lemma 5.2}] \\ &\geq F_\tau(2n) && [\text{Lemma 5.2}] \end{aligned}$$

This concludes the proof of the lemma. $\square$

Notice that for every type $\tau \neq \emptyset$ and for every $n \geq 0$, the type τ_n is equal to $\partial_n^i(\tau)$ where i is the minimal index such that $\tau(i) > 0$. The following lemma shows that if we consider a larger index, we get a smaller value.

Lemma 6.5. *For every type $\tau \neq \emptyset$, for every i such that $\tau(i) > 0$, and for every $n \geq 2$, for every $k \geq n$, we have:*

$$F_{\partial_n^i(\tau)}(k) \leq F_{\tau_n}(k)$$

Proof. We decompose τ into $\tau = \tau' + \varrho$ such that $\tau'(j) = 0$ if $j < i$ and $\varrho(j) = 0$ if $j \geq i$. Observe that $\partial_n^i(\tau) = \tau'_n + \varrho$. If $\varrho = \emptyset$ then i is the minimal index such that $\tau(i) > 0$. In this case the lemma is immediate. So, we can assume that $\varrho \neq \emptyset$. In this case $\tau_n = \tau'_n + \varrho_n$. If $\varrho_n = \emptyset$ then $\varrho = \{0\}$ and $F_{\tau_n}(k) = F_{\tau'}(k)$. Thus

$$\begin{aligned} F_{\partial_n^i(\tau)}(k) &= F_{\tau'_n + \varrho}(k) \\ &= F_{\tau'_n + \{0\}}(k) \\ &= F_{\tau'_n}^{k+1}(k) \\ &\leq F_{\tau'_k}^{k+1}(k) && [\text{Lemma 5.2}] \\ &= F_{\tau'}(k) = F_{\tau_n}(k) \end{aligned}$$

Hence, we can assume that $\varrho_n \neq \emptyset$. We introduce the norms $\|\varrho_n\|_1 = \sum_{j=0}^d \varrho_n(j)$ and $\|\varrho\|_1 = \sum_{j=0}^d \varrho(j)$. Thus $\|\varrho_n\|_1 \geq 1$. Lemma 6.4 and Lemma 5.2 show by induction that $F_{\tau_n}(k) \geq F_{\tau'}(m)$ where $m = k \cdot 2^{\|\varrho_n\|_1}$. Note that $m \geq k \cdot (1 + \|\varrho_n\|_1) \geq k + (k-1) \cdot \|\varrho_n\|_1 + \|\varrho_n\|_1 \geq k + 1 + \|\varrho_n\|_1$ from $k-1 \geq 1$ and $\|\varrho_n\|_1 \geq 1$. As $\|\varrho_n\|_1 \geq \|\varrho\|_1 - 1$, we get $m \geq k + \|\varrho\|_1$. Now, just observe that $F_{\tau'}(m) = F_{\tau'_m}^{m+1}(m) \geq F_{\tau'_m}(k)$ from Lemma 5.2. For every $j < i$, we have $\partial_n^i(\tau)(j) = \tau(j) + n = \varrho(j) + n \leq \|\varrho\|_1 + k \leq m = \tau'_m(j)$. Hence we infer that $\partial_n^i(\tau) \sqsubseteq \tau'_m$ and so Lemma 5.2 provides $F_{\partial_n^i(\tau)}(k) \leq F_{\tau'_m}(k)$. The lemma follows from this inequality, and the inequalities $F_{\tau'}(m) \geq F_{\tau'_m}(k)$ and $F_{\tau_n}(k) \geq F_{\tau'}(m)$ previously proved. $\square$

Lemma 6.6. $x.F_{\tau}(n) + y \leq F_{\tau}(x.n + x - 1 + y)$ for every $x \geq 1$, $y, n \geq 0$, and for every type τ .

Proof. The lemma is proved by induction on types, with the lexicographic order. The base case $\tau = \emptyset$ is immediate since $x.F_{\emptyset}(n) + y = x(1+n) + y = 1 + x.n + x - 1 + y = F_{\emptyset}(x.n + x - 1 + y)$. Induction case for a type $\tau \neq \emptyset$. We have $x.F_{\tau}(n) + y = x.F_{\tau_n}^{n+1}(n) + y$. Since $\tau_n < \tau$, and F_{τ_n} is monotonic, by induction on $r \geq 0$, we get for every $k \geq 0$:

$$x.F_{\tau_n}^r(k) + y = F_{\tau_n}^r(x.k + r(x-1) + y)$$

By introducing $m = x.n + x - 1 + y$, we get:

$$\begin{aligned} x.F_{\tau_n}^{n+1}(n) + y &\leq F_{\tau_n}^{n+1}(m + n.(x-1)) \\ &\leq F_{\tau_m}^{n+1}(m + n.(x-1)) \quad [\text{Lemma 5.2}] \end{aligned}$$

Notice that $F_{\emptyset}^{n.(x-1)}(m) = m + n.(x-1)$. Thus $m + n.(x-1) \leq F_{\tau_m}^{n.(x-1)}(m)$ from Lemma 5.2. As $m - n = n.(x-1) + x - 1 + y \geq 0$, we deduce that $F_{\tau_m}^{n.(x-1)}(m) \leq F_{\tau_m}^{m-n}(m)$ from Lemma 5.2. Combined with the previous relations, we get:

$$\begin{aligned} x.F_{\tau_n}^{n+1}(n) + y &\leq F_{\tau_m}^{n+1+m-n}(m) \quad [\text{Lemma 5.2}] \\ &= F_{\tau_m}^{m+1}(m) \\ &= F_{\tau}(m) \\ &= F_{\tau}(x.n + x - 1 + y) \end{aligned}$$

We have proved the induction. $\square$

Lemma 6.7. We have $K_{\tau}(n) \leq F_{\tau}(d.n)$ for every $n \geq 2$.

Proof. We prove the lemma by induction on types with the lexicographic order. The base case $\tau = \emptyset$ is immediate. Let us consider the case $\tau = \{0\}$. In this case $K_{\tau}(n) \leq K_{\emptyset}(K_{\emptyset}(n+1)) = n+1$. Since $F_{\tau}(d.n) \geq F_{\emptyset}(d.n) = d.n+1$, we are done. Let us prove the induction step with a type τ such that $\tau \neq \emptyset$ and $\tau \neq \{0\}$. Let $n \geq 2$ and i be an index such that $K_{\tau}(n) \leq K_{\partial_{d,n}^i(\tau)}(K_{\partial_{d,n}^i(\tau)}(n+1))$. We have:

$$\begin{aligned} K_{\tau}(n) &\leq F_{\partial_{d,n}^i(\tau)}(d.F_{\partial_{d,n}^i(\tau)}(d.(n+1))) \\ &\quad [\text{IH and Lemma 5.2}] \\ &\leq F_{\tau_{d,n}}(d.F_{\tau_{d,n}}(d.(n+1))) \\ &\quad [\text{Lemma 6.5 and Lemma 5.2}] \\ &\leq F_{\tau_{d,n}}(F_{\tau_{d,n}}(d^2.(n+1) + d - 1)) \\ &\quad [\text{Lemma 6.6 and Lemma 5.2}] \end{aligned}$$

Since $\tau \neq \emptyset$ and $\tau \neq \{0\}$, we deduce that $\tau_{d,n} \neq \emptyset$. With an immediate induction, and Lemma 6.4, we get:

$$\begin{aligned} F_{\tau_{d,n}}^{dn-1}(dn) &\geq 2^{dn-1}(dn+1) - 1 \\ &\geq dn(dn+1) - 1 \\ &\geq 2d(dn+1) - 1 \\ &\geq d^2.(n+1) + d - 1 \end{aligned}$$

Therefore $K_{\tau}(n) \leq F_{\tau_{d,n}}(F_{\tau_{d,n}}(F_{\tau_{d,n}}^{dn-1}(dn))) = F_{\tau_{d,n}}^{d.n+1}(d.n) = F_{\tau}(d.n)$. We have proved the induction step. $\square$

Since $L_{\tau}(n) \leq K_{\tau}(n) \leq F_{\tau}(dn)$, Theorem 6.1 is proved.

7. Application to Pushdown VASS

We now apply the results of the previous sections to vector addition systems with states (VASS) extended with a pushdown stack. Consider a *dimension* $d \in \mathbb{N}$, with $d \geq 1$. A *pushdown VASS* is a 5-tuple $\mathcal{A} = \langle Q, q_{\text{init}}, \mathbf{v}_{\text{init}}, \Gamma, \Delta \rangle$ where Q is a finite set of *locations*, $q_{\text{init}} \in Q$ is the *initial location*, $\mathbf{v}_{\text{init}} \in \mathbb{N}^d$ is the *initial vector*, Γ is a finite *stack alphabet*, and $\Delta \subseteq Q \times \mathbb{Z}^d \times \mathbf{Op} \times Q$ is a finite set of *transition rules*. As before, the set $\mathbf{Op}$ of *operations* is given by $\mathbf{Op} = \{\varepsilon\} \cup \{\text{push}(\gamma), \text{pop}(\gamma) \mid \gamma \in \Gamma\}$. Informally, a pushdown VASS can be viewed as a pushdown automaton equipped with d counters ranging over natural numbers. Operations on counters are limited to translations. A rule $(q, \mathbf{a}, \text{op}, r)$ means that the pushdown VASS can move from location q to location r by performing the operation op on the stack and adding the vector $\mathbf{a}$ to its counters, provided that they remain nonnegative.

The operational semantics of a pushdown VASS $\mathcal{A}$ is defined in two steps. First, we associate with $\mathcal{A}$ the pushdown system $\mathcal{P}^{\mathcal{A}} = \langle S^{\mathcal{A}}, s_{\text{init}}^{\mathcal{A}}, \Gamma, \Delta^{\mathcal{A}} \rangle$ given by:

$$\begin{aligned} S^{\mathcal{A}} &= Q \times \mathbb{N}^d \\ s_{\text{init}}^{\mathcal{A}} &= (q_{\text{init}}, \mathbf{v}_{\text{init}}) \\ ((q, \mathbf{u}), \text{op}, (r, \mathbf{v})) \in \Delta^{\mathcal{A}} &\Leftrightarrow (q, \mathbf{v} - \mathbf{u}, \text{op}, r) \in \Delta \end{aligned}$$

Note that $\mathcal{P}^{\mathcal{A}}$ is finitely-branching. Second, we define the operational semantics of $\mathcal{A}$ to be the operational semantics of $\mathcal{P}^{\mathcal{A}}$ (see Section 3). Let us introduce the partial-order $\preceq_{\mathcal{A}}$ on $S^{\mathcal{A}}$ defined by $(q, \mathbf{v}) \preceq_{\mathcal{A}} (q', \mathbf{v}')$ when $q = q'$ and $\mathbf{v} \leq \mathbf{v}'$. This partial-order makes $\mathcal{P}^{\mathcal{A}}$ a well-structured pushdown system (see Definition 3.1). According to Propositions 3.3, 3.4 and 3.5, the termination and boundedness problems for pushdown VASS can be solved by constructing the reduced reachability tree, which is finite. We show in this section that the worst-case size of this tree is hyper-Ackermannian.

For the remainder of this section, we restrict ourselves to pushdown VASS such that $\|\mathbf{a}\|_{\infty} \leq 1$ for each rule $(q, \mathbf{a}, \text{op}, r)$. The hyper-Ackermannian upper bound that we obtain below can be extended to the general case by splitting each rule $(q, \mathbf{a}, \text{op}, r)$ into $\|\mathbf{a}\|_{\infty}$ basic ones. The *size* of a pushdown VASS $\mathcal{A} = \langle Q, q_{\text{init}}, \mathbf{v}_{\text{init}}, \Gamma, \Delta \rangle$ is defined as

$$|\mathcal{A}| = d + |Q| + d \cdot \|\mathbf{v}_{\text{init}}\|_{\infty} + |\Gamma| + d \cdot |\Delta|$$

According to Remark 4.3, every run of $\mathcal{P}^{\mathcal{A}}$ can be seen as a nested sequence over $(S^{\mathcal{A}}, \preceq_{\mathcal{A}})$. Let us introduce the norm $\|\cdot\|_{S^{\mathcal{A}}}$ on $(S^{\mathcal{A}}, \preceq_{\mathcal{A}})$, defined by $\|(q, \mathbf{v})\|_{S^{\mathcal{A}}} = \|\mathbf{v}\|_{\infty}$. This nested sequence is n -controlled, for $n = \|\mathbf{v}_{\text{init}}\|_{\infty} + 2$, since the initial state of $\mathcal{P}^{\mathcal{A}}$ is $(q_{\text{init}}, \mathbf{v}_{\text{init}})$ and $\|\mathbf{a}\|_{\infty} \leq 1$ for each rule $(q, \mathbf{a}, \text{op}, r)$. Recall that every subsumed node in the reduced reachability tree of $\mathcal{P}^{\mathcal{A}}$ is necessarily a leaf. It follows from Remark 4.7 that the height of this tree is

bounded by the maximal length of n -controlled bad nested sequences, namely $L_{S^A}(n)$. Note that $(S^A, \preceq_A, \|\cdot\|_{S^A})$ is isomorphic to $(|Q| \times \mathbb{N}^d, \preceq_{|Q| \times \mathbb{N}^d}, \|\cdot\|_{|Q| \times \mathbb{N}^d})$. Since $Q \neq \emptyset$, $d \geq 1$ and $n \geq 2$, we get from Theorem 6.1 that the height of the reduced reachability tree is at most $F_{\omega^d, |Q|}(d \cdot n)$. Each node of the reduced reachability tree can have at most $|\Delta|$ children. We derive the following upper bound on its size.

Theorem 7.1. *The reduced reachability tree of a pushdown VASS $\mathcal{A}$ has at most $F_{\omega^\omega}(|A|)$ nodes.*

Remark 7.2. For classical VASS (i.e., $\text{op} = \varepsilon$ for each rule $(q, \mathbf{a}, \text{op}, r)$), the reduced reachability tree has at most $F_\omega(|A|)$ nodes [13, 22]. For pushdown finite-state systems, it follows from Remark 6.2 that its size is bounded by $|\Delta|^{2^{|Q|}}$.

We next give an almost matching lower bound for the above-mentioned size of the reduced reachability tree. We introduce a pushdown VASS $\mathcal{A}_d(n)$, depicted in Figure 1, that is parametrized by a “target” dimension $d \geq 1$ and an initial value $n \in \mathbb{N}$. The lower bound is obtained as follows. First, we show that $\mathcal{A}_d(n)$ computes, in a weak sense, the value $F_{\omega^d}(n)$. This computation requires, in $\mathcal{A}_d(n)$, a run of length at least $F_{\omega^d}(n)$. Second, we prove that all runs of $\mathcal{A}_d(n)$ are finite. Therefore, its reduced reachability tree is identical to its reachability tree, and so its height is at least $F_{\omega^d}(n)$. The remainder of this section presents these ideas in detail. To simplify notation, we will shortly write $\mathcal{A}_d$ when the parameter n is understood from the context.

The dimension of $\mathcal{A}_d$ is $d + 2$. But, for clarity, we will use counter names to refer to individual coordinates of the vectors in the configurations of $\mathcal{A}_d$. The parameter n is the initial value of the first counter, called r , which will also contain the result at the end. For intermediate steps, we use the counters $\bar{r}, \kappa_0, \dots, \kappa_{d-1}$. The notation $r \downarrow$ in Figure 1 is short for r^- , $\bar{r}^{++}$, i.e., a decrement of r followed by an increment of $\bar{r}$. Similarly, the notation $r \uparrow$ is short for r^{++} , $\bar{r}^-$, i.e., an increment of r followed by a decrement of $\bar{r}$. The stack alphabet is $\Gamma = \{F, D_0, I_0, \dots, D_{d-1}, I_{d-1}\}$. Intuitively, the stack and counters of $\mathcal{A}_d$ are used to represent the recursive definition of $F_{\omega^d}(n)$. As expected, the stack is used to store the pending recursive calls. The innermost $F_\alpha(m)$ computation is encoded as $\alpha = \omega^{d-1} \cdot \kappa_{d-1} + \dots + \omega^0 \cdot \kappa_0$ and $r = m + 1$. For example, $F_{\omega^{2.3+0.5+2}}(8)$ is represented by having F on top of the stack, $\kappa_2 = 3$, $\kappa_1 = 5$, $\kappa_0 = 2$ and $r = 9$. The stack symbols D_i, I_i are used to restore the calling context (i.e., the counters κ_i) of pending recursive calls. More precisely, within the stack, D_i stands for “decrement κ_i ” and I_i stands for “increment κ_i ”.

The initial location of $\mathcal{A}_d$ is q_{init} , and its initial vector is $\mathbf{v}_{\text{init}} = (n, 0, \dots, 0)$. This means that the initial value of the counter r is n and all other counters have initial value 0. First, we increment r by one and move to q_0 . The loop at q_0 allows us to increment κ_{d-1} to the value $n + 1$ and the loop at q_1 allows us to set r back to $n + 1$. Then we push F on the stack and reach q_2 . At this point, the stack contents is $\sigma = F$, the counters $\kappa_{d-1}, \dots, \kappa_0$ contain $n + 1, 0, \dots, 0$, respectively, and $r = n + 1$. This configuration represents $F_{\omega^{d-1} \cdot (n+1)}(n)$, which is equal to $F_{\omega^d}(n)$, i.e., the number that we want to compute. From the location q_2 , $\mathcal{A}_d$ can either execute $\mathcal{A}_\lambda^i$ or $\mathcal{A}_{\alpha+1}$ or $\mathcal{A}_0$ or $\mathcal{B}^i$ to reach q_3 . From the location q_3 , $\mathcal{A}_d$ can either go back to q_2 to continue the computation, or move to the final location q_f .

At q_2 , if the top of the stack is F and $\kappa_0 = 0$, then it represents $F_{\omega^{d-1} \cdot \kappa_{d-1} + \dots + \omega^i \cdot \kappa_i}(r-1)$, assuming that $\kappa_i > 0$ and




Figure 1. Pushdown VASS $\mathcal{A}_d(n)$ that weakly computes $F_{\omega^d}(n)$.

$\kappa_{i-1} = \dots = \kappa_0 = 0$. In that case, the purpose of $\mathcal{A}_\lambda^i$ is to replace the top of the stack with something that represents $F_{\omega^{d-1} \cdot \kappa_{d-1} + \dots + \omega^i \cdot (\kappa_i - 1) + \omega^{i-1} \cdot r}(r-1)$. Accordingly, $\mathcal{A}_\lambda^i$ first pops F from the stack. Then it decrements κ_i and pushes I_i on the stack to “shield” the part of the stack below from this decreased value of κ_i . Then $\mathcal{A}_\lambda^i$ increments the counter κ_{i-1} r times, also pushing the symbol D_{i-1} each time. The purpose of D_{i-1} is to “guard” the part of the stack below from this increased value of κ_{i-1} . Finally, $\mathcal{A}_\lambda^i$ pushes F on the stack, restores the value of r and exits.

At q_2 , if the top of the stack is F and $\kappa_0 > 0$, then it represents $F_{\alpha+1}(r-1)$, where $\alpha = \omega^{d-1} \cdot \kappa_{d-1} + \dots + \omega^0 \cdot (\kappa_0 - 1)$. In that case, the purpose of $\mathcal{A}_{\alpha+1}$ is to replace the top of the stack with something that represents $F_\alpha(r-1)$. Accordingly, $\mathcal{A}_{\alpha+1}$ first pops F from the stack. Then it decrements κ_0 and pushes I_0 on the stack to “shield” the part of the stack below from this decreased value of κ_0 . Then it pushes r copies of the symbol F , restores the value of r and exits.

At q_2 , if the top of the stack is F and $\kappa_{d-1} = \dots = \kappa_0 = 0$, then it represents $F_0(r-1)$. Accordingly, $\mathcal{A}_0$ simply pops F from the stack, increments r and exits.

At q_2 , if the top of the stack is I_i or D_i , then sometime back, either 1) $\mathcal{A}_\lambda^i$ decremented κ_i , pushed one I_i to shield the bottom part of the stack from this decrement, incremented κ_{i-1} n times and pushed n copies of D_{i-1} to guard the bottom part of the stack from these increments, or 2)

$i = 0, n = 0, \mathcal{A}_{\alpha+1}$ decremented κ_0 once and pushed one I_0 to shield the bottom part of the stack from this decrement. The purpose of $\mathcal{B}^0, \dots, \mathcal{B}^{d-1}$ is to backtrack and access the bottom part of the stack by removing guards and shields, and restoring the counter values. Accordingly, $\mathcal{B}^i$ either pops the symbol D_i from the top of the stack, decrements κ_i and exits, or it pops I_i , increments κ_i and exits.

Lemma 7.3. *Suppose $\mathcal{A}_d$ is in location q_2 , F is on the top of the stack and the counter r has value $m + 1$. Then there is a partial² run in $\mathcal{A}_d$ that reaches q_3 , removes F from the top of the stack, sets the counter r to the value $F_{\omega^{d-1}, \kappa_{d-1} + \dots + \omega^0, \kappa_0}(m) + 1$, and does no other change.*

Corollary 7.4. *There is a run in $\mathcal{A}_d(n)$ that reaches q_f with the empty stack and the counter r having the value $F_{\omega^d}(n)$.*

Proof. A witnessing run is provided in the table below. The partial run from q_2 to q_3 is derived from Lemma 7.3.

Loc.	r	κ_{d-1}	Stack
$(q_{\text{init}}, n,$		0,	$\varepsilon)$
$(q_0, n + 1,$		0,	$\varepsilon)$
$(q_1, 0,$		$n + 1,$	$\varepsilon)$
$(q_2, n + 1,$		$n + 1,$	$F)$
$(q_3, F_{\omega^{d-1}, (n+1)}(n) + 1,$		$n + 1,$	$\varepsilon)$
$(q_f, F_{\omega^{d-1}, (n+1)}(n),$		$n + 1,$	$\varepsilon)$

The observation that $F_{\omega^d}(n) = F_{\omega^{d-1}, (n+1)}(n)$ concludes the proof of the corollary. $\square$

Next, we prove that all runs of $\mathcal{A}_d$ are finite. The intuition is to measure, by an ordinal, the “size” of both the stack contents and the values of the counters $\kappa_{d-1}, \dots, \kappa_0$, and prove that the size decreases at each step. Consider a vector $\mathbf{v} = (v_{d-1}, \dots, v_0)$ in $\mathbb{N}^d$. The ordinal associated with $\mathbf{v}$ is defined as $\text{Ord}(\mathbf{v}) = \omega^{d-1} \cdot v_{d-1} + \dots + \omega^0 \cdot v_0$. Let $\sigma \in \Gamma^*$ be a string over the stack alphabet. The ordinal associated with σ and $\mathbf{v}$ is defined as follows, by induction on $|\sigma|$.

$$\begin{aligned}
\text{Ord}(\varepsilon, \mathbf{v}) &= 0 \\
\text{Ord}(\sigma F, \mathbf{v}) &= \text{Ord}(\sigma, \mathbf{v}) + \omega^{\text{Ord}(\mathbf{v})} \\
\text{Ord}(\sigma I_i, \mathbf{v}) &= \text{Ord}(\sigma, \mathbf{v} + \mathbf{e}_i) \\
\text{Ord}(\sigma D_i, \mathbf{v}) &= \begin{cases} \text{undefined} & \text{if } v_i = 0 \\ \text{Ord}(\sigma, \mathbf{v} - \mathbf{e}_i) & \text{if } v_i > 0 \end{cases}
\end{aligned}$$

The ordered pair $\text{dr}(\sigma, \mathbf{v}) = (\text{Ord}(\sigma, \mathbf{v}), |\sigma|)$ is the *depth of remaining recursion* associated with a configuration of $\mathcal{A}_d$ whose stack contents is σ and where the counter values of $\kappa_{d-1}, \dots, \kappa_0$ form the vector $\mathbf{v}$. We use the lexicographic order on ordered pairs to compare depths of remaining recursion.

Lemma 7.5. *Every partial run of $\mathcal{A}_d(n)$ that starts and ends in q_2 , but does not visit q_2 in between, makes the depth of remaining recursion strictly decrease.*

Corollary 7.6. *All runs of $\mathcal{A}_d(n)$ are finite.*

Proof. Suppose, for the sake of contradiction, that there is an infinite run. The loops at q_{init} and q_1 can only be executed finitely many times, so the run must visit q_2 at some point.

The infinite run can not visit q_f since q_f does not have any successors. Hence, after a finite prefix, the run must stay within $q_2, q_3, \mathcal{A}_\lambda^i, \mathcal{A}_{\alpha+1}, \mathcal{A}_0$ and $\mathcal{B}^i$. Upon entering into $\mathcal{A}_\lambda^i, \mathcal{A}_{\alpha+1}, \mathcal{A}_0$ or $\mathcal{B}^i$, the run must exit after a finite number of steps since the loops inside these can only be executed finitely many times for each entry. Hence, the run must visit $\{q_2, q_3\}$ infinitely often. Since q_2 and q_f are the only successors of q_3 and q_f is never visited, the run must visit q_2 infinitely often. Now from Lemma 7.5, we infer that there is an infinite strictly decreasing sequence of depths of remaining recursion, which is a contradiction. Hence, there is no infinite run in $\mathcal{A}_d(n)$. $\square$

We obtain the following lower bound on the size of the reduced reachability tree for pushdown VASSs.

Theorem 7.7. *For every $n \in \mathbb{N}$, there exists a pushdown VASS $\mathcal{A}_n$, of size quadratic in n , such that the reduced reachability tree of $\mathcal{A}_n$ has at least $F_{\omega^\omega}(n)$ nodes.*

Proof. Corollary 7.6, combined with Proposition 3.4, shows that the reachability tree of $\mathcal{A}_{n+1}(n)$ contains no subsumed node. Therefore, the reachability tree of $\mathcal{A}_{n+1}(n)$ coincides with its reduced reachability tree, which is finite according to Proposition 3.3. The height of this tree is at least $F_{\omega^\omega}(n) = F_{\omega^{n+1}}(n)$ by Corollary 7.4. It is readily seen that $|\mathcal{A}_{n+1}(n)|$ is quadratic in n . $\square$

Remark 7.8. The hyper-Ackermannian lower and upper bounds of Theorems 7.1 and 7.7 also apply to the reachability set when it is finite.

8. Conclusion

Reachability sets of bounded VASS (without stack) are known to be Ackermannian large. The lower-bound was provided by [21] thanks to a family of VASS with finite, but large reachability sets (Ackermannian). The upper-bound was first derived in [22] from bounds on the length of bad sequences of vectors of natural numbers. However, the boundedness problem for VASS is decidable in exponential space. This complexity was proved by Rackoff in [23] by observing that unboundedness of a VASS can be witnessed by finite runs of length at most doubly exponential.

In this paper, reachability sets of bounded pushdown VASS are shown to be hyper-Ackermannian large. The lower-bound is obtained thanks to a family of bounded pushdown VASS with finite, but large reachability sets (hyper-Ackermannian). The upper-bound is derived from a precise complexity analysis of an adaptation of the Karp & Miller algorithm. Based on this algorithm, we deduce that the unboundedness of pushdown VASS can be witnessed by finite runs. We think that the precise complexity of the boundedness problem can be obtained by bounding the length of these witnesses. However, adapting Rackoff’s techniques to pushdown VASS is a difficult and challenging problem. Note that witnesses have a length that is at least a tower of exponentials, i.e., a length in $F_3(n)$, where n is the size of the pushdown VASS. Such a lower-bound can be easily obtained by adapting [19]. The complexity gap between the lower-bound $F_3(n)$ and the upper-bound $F_{\omega^\omega}(n)$ is an open problem.

Acknowledgments

The authors thank Sylvain Schmitz and Philippe Schnoebelen for fruitful discussions.

² By *partial run*, we mean a run that is not required to start from the initial configuration.

References

- [1] P. A. Abdulla, K. Čerāns, B. Jonsson, and Y.-K. Tsay. Algorithmic analysis of programs with well quasi-ordered domains. *Inform. Comput.*, 160(1–2):109–127, 2000.
- [2] P. A. Abdulla, M. F. Atig, G. Delzanno, and A. Podelski. Push-down automata with gap-order constraints. In *Selec. FSEN’13*, volume 8161 of *LNCS*, pages 199–216. Springer, 2013.
- [3] R. Alur and P. Madhusudan. Adding nesting structure to words. *J. ACM*, 56(3):1–16, 2009.
- [4] M. F. Atig and P. Ganty. Approximating petri net reachability along context-free traces. In *Proc. FSTTCS’11*, volume 13 of *LIPICs*, pages 152–163. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2011.
- [5] A. Bouajjani and M. Emmi. Analysis of recursively parallel programs. *ACM Trans. Prog. Lang. Syst.*, 35(3):1–49, 2013.
- [6] X. Cai and M. Ogawa. Well-structured pushdown systems. In *Proc. CONCUR’13*, volume 8052 of *LNCS*, pages 121–136. Springer, 2013.
- [7] R. Chadha and M. Viswanathan. Decidability results for well-structured transition systems with auxiliary storage. In *Proc. CONCUR’07*, volume 4703 of *LNCS*, pages 136–150. Springer, 2007.
- [8] R. Chadha and M. Viswanathan. Deciding branching time properties for asynchronous programs. *Theor. Comput. Sci.*, 410(42):4169–4179, 2009.
- [9] P. Chambart and Ph. Schnoebelen. The ordinal recursive complexity of lossy channel systems. In *Proc. LICS’08*, pages 205–216. IEEE, 2008.
- [10] S. Demri, M. Jurdziński, O. Lachish, and R. Lazic. The covering and boundedness problems for branching vector addition systems. *J. Comput. Syst. Sci.*, 79(1):23–38, 2013.
- [11] N. Dershowitz and Z. Manna. Proving termination with multiset orderings. *Commun. ACM*, 22(8):465–476, 1979.
- [12] C. Dufourd, A. Finkel, and Ph. Schnoebelen. Reset nets between decidability and undecidability. In *Proc. ICALP’98*, volume 1443 of *LNCS*, pages 103–115. Springer, 1998.
- [13] D. Figueira, S. Figueira, S. Schmitz, and Ph. Schnoebelen. Ackermannian and primitive-recursive bounds with Dickson’s lemma. In *Proc. LICS’11*, pages 269–278. IEEE, 2011.
- [14] A. Finkel. Reduction and covering of infinite reachability trees. *Inform. Comput.*, 89(2):144–179, 1990.
- [15] A. Finkel and Ph. Schnoebelen. Well-structured transition systems everywhere! *Theor. Comput. Sci.*, 256(1–2):63–92, 2001.
- [16] A. Finkel, P. McKenzie, and C. Picaronny. A well-structured framework for analysing Petri net extensions. *Inform. Comput.*, 195(1-2):1–29, 2004.
- [17] P. Ganty and R. Majumdar. Algorithmic verification of asynchronous programs. *ACM Trans. Prog. Lang. Syst.*, 34(1):6:1–6:48, 2012.
- [18] C. Haase, S. Schmitz, and Ph. Schnoebelen. The power of priority channel systems. In *Proc. CONCUR’13*, volume 8052 of *LNCS*, pages 319–333. Springer, 2013.
- [19] R. Lazic. The reachability problem for vector addition systems with a stack is not elementary. *CoRR*, abs/1310.1767, 2013. Presented at RP’12.
- [20] M. Löb and S. Wainer. Hierarchies of number-theoretic functions. I. *Archiv für mathematische Logik und Grundlagenforschung*, 13(1–2):39–51, 1970. <http://dx.doi.org/10.1007/BF01967649>.
- [21] E. Mayr and A. Meyer. The complexity of the finite containment problem for petri nets. *J. ACM*, 28(3):561–576, 1981.
- [22] K. McAloon. Petri nets and large finite sets. *Theor. Comput. Sci.*, 32(1–2):173–183, 1984.
- [23] C. Rackoff. The covering and boundedness problems for vector addition systems. *Theor. Comput. Sci.*, 6(2):223–231, 1978. .
- [24] K. Reinhardt. Reachability in petri nets with inhibitor arcs. *Electr. Notes Theor. Comput. Sci.*, 223:239–264, 2008.
- [25] S. Schmitz and Ph. Schnoebelen. Multiply-recursive upper bounds with Higman’s lemma. In *Proc. ICALP’11*, volume 6756 of *LNCS*, pages 441–452. Springer, 2011.
- [26] K. Sen and M. Viswanathan. Model checking multithreaded programs with asynchronous atomic methods. In *Proc. CAV’06*, volume 4144 of *LNCS*, pages 300–314. Springer, 2006.