Efficient Analysis of VASS Termination Complexity

Antonín Kučera, Jérôme Leroux, Dominik Velan

► To cite this version:

Antonín Kučera, Jérôme Leroux, Dominik Velan. Efficient Analysis of VASS Termination Complexity. LICS '20: 35th Annual ACM/IEEE Symposium on Logic in Computer Science, Jul 2020, Saarbrücken, Germany. pp.676-688. hal-03052592

HAL Id: hal-03052592

<https://hal.science/hal-03052592>

Submitted on 19 Nov 2021

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.

Efficient Analysis of VASS Termination Complexity

Antonín Kučera
Faculty of Informatics
Masaryk University
Czechia
kucera@fi.muni.cz

Jérôme Leroux
LaBRI
France
leroux@labri.fr

Dominik Velan
Faculty of Informatics
Masaryk University
Czechia
xvelan1@fi.muni.cz

Abstract

The *termination complexity* of a given VASS is a function $\mathcal{L}$ assigning to every n the length of the longest non-terminating computation initiated in a configuration with all counters bounded by n . We show that for every VASS with demonic nondeterminism and every fixed k , the problem whether $\mathcal{L} \in \mathcal{G}_k$, where $\mathcal{G}_k$ is the k -th level in the Grzegorz hierarchy, is decidable in polynomial time. Furthermore, we show that if $\mathcal{L} \notin \mathcal{G}_k$, then $\mathcal{L}$ grows at least as fast as the generator F_{k+1} of $\mathcal{G}_{k+1}$. Hence, for every terminating VASS, the growth of $\mathcal{L}$ can be reasonably characterized by the *least* k such that $\mathcal{L} \in \mathcal{G}_k$.

Furthermore, we consider VASS with both angelic and demonic nondeterminism, i.e., VASS games where the players aim at lowering/raising the termination time. We prove that for every fixed k , the problem whether $\mathcal{L} \in \mathcal{G}_k$ for a given VASS game is **NP**-complete. Furthermore, if $\mathcal{L} \notin \mathcal{G}_k$, then $\mathcal{L}$ grows at least as fast as F_{k+1} .

CCS Concepts: • Theory of computation $\rightarrow$ Abstract machines.

Keywords: Vector Addition Systems, Termination

ACM Reference Format:

Antonín Kučera, Jérôme Leroux, and Dominik Velan. 2020. Efficient Analysis of VASS Termination Complexity. In *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS '20)*, July 8–11, 2020, Saarbrücken, Germany. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3373718.3394751>

1 Introduction

Vector addition systems with states (VASS) are a generic computational model of discrete systems operating over unbounded data domains. More precisely, a d -dimensional

VASS is a directed graph with a finite set of vertices Q (called *control states*) where every edge is assigned an *update vector* $\mathbf{u} \in \mathbb{Z}^d$. A *configuration* is a pair $p\mathbf{v}$ where p is the current control state and $\mathbf{v} \in \mathbb{N}^d$ is the vector of current counter values. Transitions between configurations are defined in the natural way. When considering the encoding size $\|\mathcal{A}\|$ of a given VASS $\mathcal{A}$, we assume that the update vectors are written in binary.

One of the fundamental question studied in program analysis is *termination*. For non-deterministic computational models, each choice is considered either *angelic* or *demonic* [6], i.e., favoring either termination or non-termination, respectively.

Although VASS are widely accepted as a natural model for programs operating over integer variables, the principal limits of effective VASS *termination complexity analysis* are not yet fully understood. Assume that Q is split into two subsets of angelic and demonic control states. A computation initiated in a configuration $p\mathbf{v}$ is then determined by two players, A (angel) and D (daemon) who resolve nondeterminism in configurations with angelic and demonic control states, aiming at minimizing and maximizing the length of the computation, respectively. By applying standard results about games with accumulated reward objectives (see, e.g., [5]), we obtain that for every initial configuration $p\mathbf{v}$ there exists a unique *termination value* $Tval(p\mathbf{v}) \in \mathbb{N}_\infty$ defined by the equality

$$\sup_{\pi} \inf_{\sigma} |Comp^{\sigma, \pi}(p\mathbf{v})| = \inf_{\sigma} \sup_{\pi} |Comp^{\sigma, \pi}(p\mathbf{v})|$$

where σ and π range over the strategies of the players A and D , and $|Comp^{\sigma, \pi}(p\mathbf{v})|$ is the length of the maximal computation obtained by applying the strategies σ and π in $p\mathbf{v}$. The *termination complexity* is a function $\mathcal{L} : \mathbb{N} \rightarrow \mathbb{N}_\infty$ assigning to every $n \in \mathbb{N}$ the maximal $Tval(p\mathbf{v})$, where $p \in Q$ and $\mathbf{v} \leq (n, \dots, n)$.

The existing works on VASS termination complexity have concentrated mainly on recognizing VASS with *low* asymptotic growth of $\mathcal{L}$ where all control states are *demonic*. More concretely, the *linearity* of $\mathcal{L}$, i.e., the question whether $\mathcal{L} \in \mathcal{O}(n)$, is solvable in polynomial time for demonic VASS. Furthermore, if $\mathcal{L}$ is *not* linear, then it is at least quadratic [3]. If a given demonic VASS is *positive-normal*¹, then there

¹A VASS is *positive normal* if there is a quasi-ranking function such that each component of its normal is positive (this condition can be verified in polynomial time). We refer to [3] for details.

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 the author(s) 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. LICS '20, July 8–11, 2020, Saarbrücken, Germany
© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-7104-9/20/07...\$15.00
<https://doi.org/10.1145/3373718.3394751>

exists $k \leq d$ computable in polynomial time (d is the dimension) such that $\mathcal{L} \in \Theta(n^k)$ [3]. The *polynomiality* of $\mathcal{L}$, i.e., the question whether $\mathcal{L} \in \mathcal{O}(n^k)$ for some $k \in \mathbb{N}$, is also decidable in polynomial time for demonic VASS [16]. In the same paper, it is also shown that if $\mathcal{L}$ is *not* polynomial, then it is at least *exponential*, i.e., $\mathcal{L} \in 2^{\Omega(n)}$. A recent result of [20] shows that if the termination complexity of a given demonic VASS is polynomial, then there is $k \in \mathbb{N}$ computable in polynomial time such that $\mathcal{L}$ is $\Theta(n^k)$. A polynomial-time algorithm deciding the linearity of $\mathcal{L}$ for *probabilistic* VASS, i.e., VASS with demonic non-determinism and probabilistic choice², is given in [2].

Our contribution. In this work, we extend the scope of efficient VASS termination analysis to general complexity classes determined by the standard hierarchy of fast growing functions, and also to VASS with both demonic and angelic non-determinism (i.e., VASS *games*). More precisely, we prove the following.

(A) For all $F : \mathbb{N} \rightarrow \mathbb{N}$ and $\ell \in \mathbb{N}$, let $F^{(\ell)}(x)$ denote the ℓ -th iterate of F , i.e., $F(\dots(F(x)\dots))$ composed ℓ times, where $F^{(0)}(x) = x$. Let $F_k(n)$, where $k \geq 1$, be the hierarchy of fast-growing functions³ defined by

$$\begin{aligned} & F_1(n) = 2n + 1, \quad F_2(n) = n^2, \quad F_3(n) = 2^n, \\ & F_{k+1}(n) = F_k^{(n)}(n) \text{ for } k \geq 3. \end{aligned}$$

Hence, F_4 is already non-elementary. For every $k \geq 1$, the class of functions $\mathcal{G}_k$ is defined by

$$\mathcal{G}_k = \{f : \mathbb{N} \rightarrow \mathbb{N} \mid f \leq F_k^{(\mu)} \text{ for some } \mu \in \mathbb{N}\}.$$

For every $k \geq 4$, the class $\mathcal{G}_k$ exactly captures the growth of functions in the class $\mathcal{E}_k$ of the standard Grzegorzczuk hierarchy [12] (each function of $\mathcal{E}_k$ is bounded by some function of $\mathcal{G}_k$ for all sufficiently large arguments, and vice versa). Hence, we also refer to $\mathcal{G}_k$ as the *Grzegorzczuk hierarchy* in the rest of this paper. Note that $\mathcal{G}_1$, $\mathcal{G}_2$, and $\mathcal{G}_3$ contain all linear, polynomial, and elementary functions, respectively.

Let $\mathcal{A}$ be a d -dimensional demonic VASS. We show that for every $k \geq 3$, the problem whether $\mathcal{L} \in \mathcal{G}_k$ is solvable in time polynomial in $\|\mathcal{A}\|$. We also show that if $\mathcal{L} \notin \mathcal{G}_k$, then $\mathcal{L}(n) \geq F_{k+1}(\lfloor n/c \rfloor)$ for all $n \in \mathbb{N}$ and some constant c depending only on $\mathcal{A}$, i.e., the growth of $\mathcal{L}$ indeed reaches⁴ the next level $\mathcal{G}_{k+1}$. Due to this gap, the growth of $\mathcal{L}$ for a terminating⁵ $\mathcal{A}$ is reasonably characterized by the least $\lambda \in \mathbb{N}$ such that $\mathcal{L} \in \mathcal{G}_\lambda$ (it follows already from the results of [19])

²Actually, this result is obtained for a subclass of probabilistic VASS with tree-like MEC decomposition.

³The standard definition of fast-growing hierarchy is simpler and sets $F_0(n) = n + 1$, $F_{k+1}(n) = F_k^{(n)}(n)$. However, the Grzegorzczuk hierarchy is based on the functions $E_0(m, n) = m + n$, $E_1(n) = n^2 + 2$, and $E_{k+1}(n) = E_k^{(n)}(2)$. Since we wish to keep a precise correspondence with the Grzegorzczuk hierarchy for $k \geq 4$, we prefer the presented variant of the F_k functions.

⁴Note that according to our result, it cannot happen that $\mathcal{L} \notin \mathcal{G}_k$ and $\mathcal{L} \leq F_k^{(g(n))}(n)$ for some sublinear function $g(n)$ such as $\sqrt{n}$ or $\log(n)$.

⁵A VASS $\mathcal{A}$ is *terminating* if $\mathcal{L}(n) < \infty$ for all $n \in \mathbb{N}$.

that $\lambda \leq d + 1$). There are known examples of d -dimensional VASS such that $\mathcal{L}$ grows at least as fast as F_d [17]. Given the enormous size of the state space of such VASS, it is perhaps surprising that the property $\mathcal{L} \in \mathcal{G}_k$ admits an efficient structural characterization.

Actually, our algorithm does a bit more than just deciding whether or not $\mathcal{L} \in \mathcal{G}_k$. For a given VASS $\mathcal{A}$, we say that a given transition (or a counter) of $\mathcal{A}$ is *beyond* F_j if there is a constant c depending only on $\mathcal{A}$ such that for all sufficiently large $n \in \mathbb{N}$ there is a computation initiated in a configuration $p\mathbf{v}$ where $\mathbf{v} \leq (n, \dots, n)$ along which the transition is executed at least $F_j(\lfloor n/c \rfloor)$ times (or the counter is pumped to at least $F_j(\lfloor n/c \rfloor)$, respectively). The very core of our algorithm is a novel procedure with input $(k, \mathcal{A})$ and output (T, I) , where

- $\mathcal{A}$ is a strongly connected VASS, $k \geq 3$ is an index,
- T and I are the sets of all transitions and counters beyond F_k .

For $k = 3$, the procedure invokes the algorithm of [16]. For $k > 3$, the procedure starts by computing the sets T' and I' of all transitions and counters beyond F_{k-1} recursively, and proceeds by eliminating more and more transitions and counters in T' and I' that are bounded by $F_{k-1}^{(\mu)}$ for some fixed $\mu \in \mathbb{N}$. This goes on until a fixed point T, I is found. In the elimination phase, recursive calls are applied also to certain sub-VASS of $\mathcal{A}$.

Using the above procedure, we not only determine the existence of a transition beyond F_k (this is equivalent to $\mathcal{L} \notin \mathcal{G}_{k-1}$), but also identify all transitions and counters beyond F_k . Furthermore, for all transitions and counters that are *not* beyond F_k we obtain the *largest* index ℓ such that the transition/counter is beyond F_ℓ . In particular, for $k = d + 1$, our algorithm outputs the corresponding ℓ for *all* counters and transitions of $\mathcal{A}$, i.e., we get a precise classification of the “Grzegorzczuk complexity” for every counter and transition of a strongly connected demonic VASS $\mathcal{A}$.

The running time of the procedure is polynomial in $\|\mathcal{A}\|$ for every fixed k . Our current upper bound on the degree of the corresponding polynomial increases with k . When setting $k = d + 1$, the bound becomes exponential in d (for a fixed VASS dimension, it stays polynomial). The problem is that our procedure performs recursive calls not only for smaller indexes but also for smaller VASS. It is still possible that the number of recursive calls with distinct inputs is actually polynomial in d , and our algorithm becomes polynomial in d just by employing dynamic programming. This question seems to require additional insights and it is left for future work.

(B) Let $\mathcal{A}$ be a general VASS with both angelic and demonic non-determinism (i.e., $\mathcal{A}$ is a VASS game). We show that for every fixed $k \geq 1$, the problem whether $\mathcal{L} \in \mathcal{G}_k$ is NP-complete. Furthermore, if $\mathcal{L} \notin \mathcal{G}_k$, then $\mathcal{L}$ is beyond F_{k+1} .

Thus, we extend the dichotomy of termination complexity for demonic VASS established in (A) to VASS games.

The membership to **NP** is obtained by proving that $\mathcal{L}_{\mathcal{A}} \in \mathcal{G}_k$ iff there exists a *counterless*⁶ strategy η of player A in $\mathcal{A}$ such that $\mathcal{L}_{\mathcal{A}^\eta} \in \mathcal{G}_k$. Here, $\mathcal{A}^\eta$ is the VASS obtained by “applying” η to $\mathcal{A}$, i.e., for every angelic control state p , the only out-going transition of p in $\mathcal{A}^\eta$ is the transition selected by η in p . Since $\mathcal{A}^\eta$ can be seen as demonic VASS (player A has no real choice in $\mathcal{A}^\eta$), the problem whether $\mathcal{L}_{\mathcal{A}^\eta} \in \mathcal{G}_k$ is decidable in polynomial time by applying the results of (A). The **NP**-hardness follows easily from the results of [9].

Related work. Deciding if a demonic VASS has a finite termination value is exactly the same problem as deciding the classical termination problem for VASS. Let us recall that for every $k \in \mathbb{N}$, there exists a demonic VASS of size $O(k)$ with termination complexity beyond F_k [17]. Nevertheless, the problem of deciding the termination problem for an initialized VASS is exponential space complete [18]. A variant of that problem, called the structural termination problem, consists of deciding if a VASS is terminating from any initial configuration. This problem is equivalent to decide if the *termination complexity* is always finite. This property can be decided in polynomial time thanks to the Kosaraju-Sullivan algorithm [15] that decides in polynomial time the presence of cycles in a VASS with non-negative weights.

The model of VASS games with termination objectives has been studied in [4, 13, 14]. Here, the main question is whether player D has a strategy preventing player A from reaching a terminal configuration from a given initial configuration. The existence and structure of winning strategies in a closely related model of multi-dimensional energy games (with or without a fixed initial credit in the counters and for various objectives) has been considered in, e.g., [7, 8, 10, 11]). However, the problem of asymptotic termination complexity for VASS games has not yet been addressed in previous works.

2 Preliminaries

We use $\mathbb{Z}$ and $\mathbb{N}$ to denote the sets of integers and non-negative integers, respectively. Further, we use $\mathbb{N}_\infty$ to denote the set $\mathbb{N} \cup \{\infty\}$ where ∞ is treated according to the standard conventions. For a given rational r , we use $\lfloor r \rfloor$ to denote the lower integer part of r .

We use bold letters such as $\mathbf{v}, \mathbf{u}, \dots$ to denote the vectors of $\mathbb{Z}^d$ where $d \geq 1$. For every $n \in \mathbb{Z}$, the vector $(n, \dots, n)$ is denoted by $\vec{n}$. The i -th component of $\mathbf{v}$ is denoted by $\mathbf{v}(i)$, and we use $\mathbf{v}^-$ to denote the set of all $i \in \{1, \dots, d\}$ such that $\mathbf{v}(i) < 0$. A vector $\mathbf{v}$ is *non-negative* if $\mathbf{v}^- = \emptyset$.

Let $f, g : \mathbb{N} \rightarrow \mathbb{N}_\infty$. We write $f \leq g$ if $f(n) \leq g(n)$ for all $n \in \mathbb{N}$. We say that f is *beyond* g if there exist $c \geq 1$ and $n_0 \in \mathbb{N}$ such that $f(n) \geq g(\lfloor n/c \rfloor)$ for all $n \geq n_0$.

⁶A strategy is *counterless* if it depends just on the currently visited control state.

2.1 Vector Addition Systems with States (VASS)

A VASS is a finite state directed graph with transitions labelled by vectors of counter changes.

Definition 2.1. Let $d \geq 1$. A d -dimensional vector addition system with states (VASS) is a pair $\mathcal{A} = (Q, \text{Tran})$, where $Q \neq \emptyset$ is a finite set of *states* and $\text{Tran} \subseteq Q \times \mathbb{Z}^d \times Q$ is a finite set of *transitions* such that for every $q \in Q$ there exists $p \in Q$ and $\mathbf{u} \in \mathbb{Z}^d$ such that $(q, \mathbf{u}, p) \in \text{Tran}$.

We assume that Q is split into two disjoint subsets Q_A and Q_D of *angelic* and *demonic* states controlled by the players A and D , respectively. A VASS is *demonic* if $Q_A = \emptyset$. We say that $\mathcal{A}$ is *strongly connected* if the underlying directed graph of $\mathcal{A}$ is strongly connected.

A *configuration* of $\mathcal{A}$ is a pair $p\mathbf{v} \in Q \times \mathbb{N}^d$. A *computation* of $\mathcal{A}$ is a sequence of configurations $\alpha = p_1\mathbf{v}_1, p_2\mathbf{v}_2, \dots$ of length $m \in \mathbb{N}_\infty$ such that for every $1 \leq i < m$ there is a transition $(p_i, \mathbf{u}_i, p_{i+1})$ satisfying $\mathbf{v}_{i+1} = \mathbf{v}_i + \mathbf{u}_i$.

A *finite path* in $\mathcal{A}$ of length m is a finite sequence $\rho = p_1, \mathbf{u}_1, p_2, \mathbf{u}_2, \dots, p_m$ such that $(p_i, \mathbf{u}_i, p_{i+1}) \in \text{Tran}$ for all $1 \leq i < m$. We use $\Delta(\rho)$ to denote the *effect* of ρ , defined as $\sum_{i=1}^m \mathbf{u}_i$. Hence, $\Delta(\rho)^-$ is the set of all counters strictly decreased by ρ . A finite path is a *cycle* if $p_m = p_1$. An *infinite path* in $\mathcal{A}$ is an infinite sequence $\alpha = p_1, \mathbf{u}_1, p_2, \mathbf{u}_2, \dots$ such that every finite prefix $p_1, \mathbf{u}_1, \dots, p_m$ of α is a finite path in $\mathcal{A}$. Note that every computation determines the associated path in $\mathcal{A}$ in the natural way.

2.2 VASS Termination Complexity

A *strategy* for player A (or player D) in $\mathcal{A}$ is a function η assigning to every finite computation $p_1\mathbf{v}_1, \dots, p_m\mathbf{v}_m$ where $p_m \in Q_A$ (or $p_m \in Q_D$) a transition $(p_m, \mathbf{u}, q)$.

Every pair of strategies (σ, π) for players A and D and every initial configuration $p\mathbf{v}$ determine the unique *maximal*⁷ computation $\text{Comp}^{\sigma, \pi}(p\mathbf{v})$ initiated in $p\mathbf{v}$ in the natural way. We use $|\text{Comp}^{\sigma, \pi}(p\mathbf{v})|$ to denote the length of $\text{Comp}^{\sigma, \pi}(p\mathbf{v})$.

For every initial configuration $p\mathbf{v}$, the players A and D aim at minimizing and maximizing the length of a computation initiated in $p\mathbf{v}$, respectively. By applying standard arguments (see, e.g., [5]), we obtain

$$\sup_{\pi} \inf_{\sigma} |\text{Comp}^{\sigma, \pi}(p\mathbf{v})| = \inf_{\sigma} \sup_{\pi} |\text{Comp}^{\sigma, \pi}(p\mathbf{v})|$$

where σ and π range over all strategies for players A and D in $\mathcal{A}$, respectively. Hence, there exists a unique *termination value* of $p\mathbf{v}$, denoted by $\text{Tval}(p\mathbf{v})$, defined by the above equality. Furthermore, both players have *optimal* positional⁸

⁷A computation is maximal if it cannot be prolonged consistently with σ and π . That is, $\text{Comp}^{\sigma, \pi}(p\mathbf{v})$ is either infinite or ends in a configuration where the responsible player selects a transition making some of the counters negative.

⁸A strategy is *positional* if it depends only on the currently visited configuration.

strategies σ^* and π^* achieving the outcome $Tval(pv)$ or better in every configuration pv against every strategy of the opponent.

The *termination complexity* of $\mathcal{A}$ is a function $\mathcal{L} : \mathbb{N} \rightarrow \mathbb{N}_\infty$ defined as follows:

$$\mathcal{L}(n) = \max\{Tval(pv) \mid p \in Q, v \leq \vec{n}\}.$$

In situations when the underlying VASS $\mathcal{A}$ is not clearly determined by the context, we use a lower index $\mathcal{A}$ (writing, e.g., $\mathcal{L}_{\mathcal{A}}$, $Tval_{\mathcal{A}}(pv)$, etc.)

3 Results about Demonic VASS

Let $\mathcal{A} = (Q, Tran)$ be a d -dimensional demonic VASS. For every transition $t \in Tran$ and every counter $i \in \{1, \dots, d\}$, we define the *transition complexity* $\mathcal{T}_t : \mathbb{N} \rightarrow \mathbb{N}_\infty$ and the *counter complexity* $C_i : \mathbb{N} \rightarrow \mathbb{N}_\infty$ as follows:

$$\begin{aligned} \mathcal{T}_t(n) &= \sup_{\pi} \sup_{p \in Q} \{\#_t(Comp^\pi(p\vec{n}))\} \\ C_i(n) &= \sup_{\pi} \sup_{p \in Q} \sup \{v(i) \mid qv \text{ occurs in } Comp^\pi(p\vec{n})\} \end{aligned}$$

Here, $\#_t$ is a function assigning to every computation of $\mathcal{A}$ the (possibly infinite) number of occurrences of the transition t in the corresponding path in $\mathcal{A}$.

For every $\Omega \subseteq \{1, \dots, d\}$, we define the VASS $\mathcal{A}[\Omega]$ obtained from $\mathcal{A}$ by *ignoring* the counters in Ω . That is, every transition (p, v, q) is changed so that $v_i = 0$ for all $i \in \Omega$.

Recall the hierarchy of fast-growing functions F_k and the Grzegorzczuk hierarchy $\mathcal{G}_k$ given in Section 1. We say that a given $f : \mathbb{N} \rightarrow \mathbb{N}$ is *below* $F_k^{(\mu)}$, where k, μ are positive integers, if there exists $g \in \mathcal{G}_{k-1}$ such that $f \leq F_k^{(\mu-1)} \circ g$.

We give a polynomial-time algorithm computing the sets of all counters/transitions whose complexity is beyond F_k for a given $k \geq 3$. At the same time, we also show that if the complexity of a given counter/transition is *not* beyond F_k , then it belongs to $\mathcal{G}_{k-1}$.

We start by observing that for every $k \geq 3$, the set of all transitions t such that $\mathcal{T}_t$ is beyond F_k is easily determined by the set of all counters i such that C_i is beyond F_k . The next lemma is actually formulated in a more precise way, which becomes useful when proving the correctness of our algorithm.

Lemma 3.1. *Let $k \geq 3$, and let I be the set of all $i \in \{1, \dots, d\}$ such that C_i is beyond $F_k^{(\mu)}$ in $\mathcal{A}$ for some positive $\mu \in \mathbb{N}$. Suppose that for every counter $i \notin I$ we have that C_i is below $F_k^{(\mu)}$ in $\mathcal{A}$. Then*

$T = \{t \in Tran \mid \exists \text{ cycle } \gamma \text{ such that } t \in \gamma \text{ and } \Delta(\gamma)^- \subseteq I\}$
is the set of all transitions with complexity beyond $F_k^{(\mu)}$ in $\mathcal{A}$.
If $t \notin T$, then $\mathcal{T}_t$ is below $F_k^{(\mu)}$ in $\mathcal{A}$.

Proof. First we show that the complexity of every $t \in T$ is beyond $F_k^{(\mu)}$ in $\mathcal{A}$. Here we use the following claim, which is also used later.

Claim 1. *Let $\mathcal{B}$ be a strongly connected VASS of dimension d , and let I be the set of all counters beyond $F_k^{(\mu)}$ for some $k, \mu \in \mathbb{N}$. Then there exists constant $c \in \mathbb{N}$ and a control state p such that for all sufficiently large n there is a computation initiated in $p\vec{n}$ visiting a configuration qv such that $v(i) \geq F_k^{(\mu)}(\lfloor n/c \rfloor)$ for all $i \in I$, and $v(i) \geq \lfloor n/c \rfloor$ for all counters i not in I .*

Proof of the claim. Let us fix some control state p of $\mathcal{A}$. Assume $I = \{i_1, \dots, i_w\}$. For every $j \in \{1, \dots, w\}$ and all sufficiently large n , there exists a computation β_j initiated in a configuration $q_j\vec{m}$, where $m = \lfloor n/(w+1) \rfloor$, reaching a configuration where the i_j -th counter is at least $F_k^{(\mu)}(\lfloor m/c \rfloor)$ for some constant $c \geq 1$ independent of m . Let γ_j be the corresponding path in $\mathcal{A}$. Consider a path ϱ obtained by concatenating

$$\tau_1, \gamma_1, \tau_2, \gamma_2, \dots, \tau_w, \gamma_w$$

where τ_1 is a path from p to the initial state of γ_1 , and τ_{j+1} is a path from the last state of γ_j to the first state of γ_{j+1} . Furthermore, the length of every τ_j is at most $|Q|$.

Note that by executing ϱ from the initial configuration $p\vec{n}$, we reach a configuration where all counters of I are pumped to $F_k^{(\mu)}(\lfloor m/c \rfloor) - b \geq F_k^{(\mu)}(\lfloor n/c' \rfloor)$ for some constant c' independent of n , and the counters not in I stay above $\lfloor n/(w+1) \rfloor - b$, where $b = |Q| \cdot w \cdot \kappa$ and κ is the maximal decrease of a counter by a transition of $\mathcal{A}$. Since $w \leq d$, the claim follows. $\square$

Let $t \in T$, and let γ be a cycle containing t such that $\Delta(\gamma)^- \subseteq I$. Using Claim 1, we immediately obtain that for all sufficiently large n , there exists a computation initiated in $p\vec{n}$ along which the cycle γ is executed at least $\lfloor n/c' \rfloor$ times, where c' is a suitable constant.

Now let $t \notin T$. Let V_n be the set of all v such that $v(i) = 0$ for all $i \in I$, and there exists a computation initiated in $q\vec{n}$ (for some q) visiting a configuration ru such that $u(i) = v(i)$ for every $i \notin I$. Since C_i is below $F_k^{(\mu)}$ for every $i \notin I$, the size of V_n is also below $F_k^{(\mu)}$. That is, there is $g \in \mathcal{G}_{k-1}$ such that the size of V_n is bounded by $F_k^{(\mu-1)}(g(n))$ for all $n \in \mathbb{N}$. Now suppose there exists $n \in \mathbb{N}$ and a computation initiated in $q\vec{n}$ along which the transition t is executed at least $F_k^{(\mu-1)}(g(n)) + 1$ times. That is, the computation visits configurations

$$p\mathbf{v}_1, p\mathbf{v}_2, \dots, p\mathbf{v}_m$$

each of which executes t as the next transition and $m \geq F_k^{(\mu-1)}(g(n)) + 1$. Since the size of V_n is bounded by $F_k^{(\mu-1)}(g(n))$, there are $1 \leq k < k' \leq m$ such that $\mathbf{v}_k(i) = \mathbf{v}_{k'}(i)$ for all counters $i \notin I$. Furthermore, the sub-computation from $p\mathbf{v}_k$ to $p\mathbf{v}_{k'}$ determines a cycle γ on p starting with t such that $\Delta(\gamma)(i) = 0$ for every counter $i \notin I$. Hence, $\Delta(\gamma)^- \subseteq I$, which is a contradiction. $\square$

```

1 procedure LargeCounters( $k, \mathcal{A}$ )
2 input:  $k \geq 3$ , a strongly connected VASS  $\mathcal{A} = (Q, Tran)$ 
3 output: the set  $I$  of all counters  $i$  such that  $C_i$  is beyond  $F_k$  in  $\mathcal{A}$ 
4
5 if  $k == 3$ :
6   return IterationSchemeAlgorithm( $\mathcal{A}$ )
7  $I = \{1, \dots, d\}$ 
8 repeat
9    $T = \{t \in Tran \mid \exists \text{ cycle } \gamma \text{ such that } t \in \gamma \text{ and } \Delta(\gamma)^- \subseteq I\}$ 
10   $\Omega = \emptyset$ 
11  repeat
12     $Aux = \emptyset$ 
13    for all  $\mathcal{B}$  determined by  $\mathcal{A}_T$ 
14       $Aux = Aux \cup \text{LargeCounters}(k-1, \mathcal{B}[\Omega])$ 
15     $\Omega = \Omega \cup Aux$ 
16  until  $\Omega$  is unchanged
17   $I = \Omega$ 
18 until  $I$  is unchanged
19 return  $I$ 

```

Figure 1. A recursive procedure for computing all counters beyond F_k in $\mathcal{A}$.

According to Lemma 3.1, for every $k \geq 3$, the set of all transitions t such that $\mathcal{T}_t$ is beyond F_k is precisely the set

$$\{t \in Tran \mid \exists \text{ cycle } \gamma \text{ such that } t \in \gamma \text{ and } \Delta(\gamma)^- \subseteq I\}$$

where I is the set of all counters i such that C_i is beyond F_k . Hence, it suffices to compute the set I . This is achieved by the recursive procedure **LargeCounters** shown in Fig. 1. The procedure inputs an index $k \geq 3$ and a strongly connected d -dimensional VASS $\mathcal{A}$, and outputs the set I of all counters whose complexity is beyond F_k in $\mathcal{A}$.

The correctness of the procedure **LargeCounters** is proven by the induction on $k \geq 3$. In induction step, we need an assumption about the structure of computations of $\mathcal{A}$ that “pump” the counters with complexity beyond F_{k-1} to large values. This assumption is formulated in the next definition (in our inductive proof, we show the assumption is valid for all $k \geq 3$).

Definition 3.2. Let $\mathcal{A}$ be a d -dimensional strongly connected VASS, $k \geq 3$, and I the set of all counters with complexity beyond F_k . We say that $\mathcal{A}$ is *well-behaving for k* if there exist positive constants $c, \kappa, \alpha \in \mathbb{N}$ and a control state p such that for all sufficiently large n , there is a finite computation $\beta_n = p\vec{n}, \dots, p\vec{v}$ of length at most $F_k(\alpha \cdot n)$ where $v(i) \geq F_k(\lfloor n/c \rfloor)$ for all $i \in I$, and the associated path ϱ_n in $\mathcal{A}$ is a cycle with a special structure shown in Fig. 2 satisfying the following conditions:

- The length of the “inner cycle” obtained by concatenating $\tau_1, \dots, \tau_k$ is $\Theta(n)$.

Figure 2. The structure of the cycle ϱ_n associated to β_n .

- For every “outer” cycle $\gamma_j = q_1, u_1, q_2, u_2, \dots, q_m$, the following two conditions are satisfied:
 - For every counter $i \notin I$, we have that $\Delta(\gamma_j)(i) = 0$.
 - For every counter $i \notin I$ and every $m' \leq m$, we have that $\Delta(q_1, u_1, \dots, q_{m'})(i) \geq -\kappa$. That is, the intermediate decrease of every counter $i \notin I$ is bounded by κ along γ_j .

Let $\mathcal{A}$ be a strongly connected d -dimensional VASS. For $k = 3$, the procedure **LargeCounters** invokes the **IterationSchemeAlgorithm** of [16], returning the set I of all counters i such that C_i is beyond F_3 . The complexity of every counter $i \notin I$ is polynomial, i.e., $C_i \in \mathcal{G}_2$. So, the procedure is correct for $k = 3$. On top of that, every $\mathcal{A}$ is well-behaving for $k = 3$.

Now we show that the procedure is correct for a given $k \geq 4$, assuming the following:

- (A1) **LargeCounters** is correct for $k - 1$ and an arbitrary strongly connected VASS on input. Furthermore, **LargeCounters** terminates in time polynomial in the size of the VASS.
- (A2) Every strongly connected VASS is well-behaving for $k - 1$.
- (A3) If the complexity of a given counter in a given strongly connected VASS is not beyond F_{k-1} , then it is in $\mathcal{G}_{k-2}$.

For every $\mu \in \mathbb{N}$, let $I(\mu)$ be the set of all $i \in \{1, \dots, d\}$ such that C_i is beyond $F_{k-1}^{(\mu)}$ in $\mathcal{A}$. Let us introduce the set

$$T(\mu) = \{t \in Tran \mid \exists \text{ cycle } \gamma \text{ s.t. } t \in \gamma \text{ and } \Delta(\gamma)^- \subseteq I(\mu)\}$$

Remark 3.1. Note that if we further assume that C_i is below $F_{k-1}^{(\mu)}$ for every $i \notin I(\mu)$, then we can apply Lemma 3.1 and conclude that $T(\mu)$ is the set of all transitions with complexity beyond $F_{k-1}^{(\mu)}$, and the complexity of all transitions $t \notin T(\mu)$ is below $F_{k-1}^{(\mu)}$. We use this argument in the proof of Lemma 3.3.

Figure 3. The structure of a compensating cycle

Observe that, by definition of $T\langle\mu\rangle$, the VASS $\mathcal{A}_{T\langle\mu\rangle}$ obtained from $\mathcal{A}$ by restricting transitions to $T\langle\mu\rangle$ is a disjoint union of strongly connected components $\mathcal{B}_1^\mu, \dots, \mathcal{B}_{m_\mu}^\mu$, each of which can be seen as a strongly connected d -dimensional VASS. Of course, $I\langle\mu\rangle$, $T\langle\mu\rangle$, and $\mathcal{B}_1^\mu, \dots, \mathcal{B}_{m_\mu}^\mu$ depend also on k , but this parameter has been fixed for now, so there is no risk of confusion.

Lemma 3.3. *For all $\mu \in \mathbb{N}$, consider the sets $S_j^{\mu+1}$ defined inductively for all $j \geq 1$ as follows:*

- *For every $j \geq 1$, the set $S_j^{\mu+1}$ consists of all counters with complexity beyond F_{k-1} in $\mathcal{B}[S_1^{\mu+1} \cup \dots \cup S_{j-1}^{\mu+1}]$ for some $\mathcal{B} \in \{\mathcal{B}_1^\mu, \dots, \mathcal{B}_{m_\mu}^\mu\}$ (here the empty union denotes $\emptyset$).*

Then, the complexity of every counter in $\bigcup_{j=1}^\infty S_j^{\mu+1}$ is beyond $F_{k-1}^{(\mu+1)}$ in $\mathcal{A}$, and the complexity of all other counters is below $F_{k-1}^{(\mu+1)}$ in $\mathcal{A}$.

Proof. We prove the lemma by induction on μ . In the base case when $\mu = 0$, we have that $T\langle 0 \rangle = \text{Tran}$, i.e., $\{\mathcal{B}_1^0, \dots, \mathcal{B}_{m_0}^0\} = \{\mathcal{A}\}$. Hence, S_1^1 is the set of all counters with complexity beyond F_{k-1} in $\mathcal{A}[\emptyset] = \mathcal{A}$, i.e., $S_1^1 = I\langle 1 \rangle$. This means that $S_2^1 = \emptyset$, because for every counter i we have that if C_i is beyond F_{k-1} in $\mathcal{A}[I\langle 1 \rangle]$, then C_i is beyond F_{k-1} also in $\mathcal{A}^0$. Furthermore, we obtain that the complexity of every counter $i \notin I\langle 1 \rangle = \bigcup_{j=1}^\infty S_j^1$ is below F_{k-1} in $\mathcal{A}$ by assumption (A3).

Now let $\mu \geq 1$. By induction on $j \geq 1$, we prove that all counters of $S_1^{\mu+1} \cup \dots \cup S_{j-1}^{\mu+1}$ are beyond $F_{k-1}^{(\mu+1)}$ in $\mathcal{A}$.

For $j = 1$, the union is empty and the claim follows trivially. Now let $j \geq 2$, and let $S = S_1^{\mu+1} \cup \dots \cup S_{j-2}^{\mu+1}$. By induction

⁹If C_i is beyond F_{k-1} in $\mathcal{A}[I\langle 1 \rangle]$, we can construct a computation “pumping” the counter i beyond F_{k-1} in $\mathcal{A}$ by first increasing all counters of $I\langle 1 \rangle$ beyond F_{k-1} and then following the computation of $\mathcal{A}[I\langle 1 \rangle]$.

hypothesis, the complexity of all counters of S is beyond $F_{k-1}^{(\mu+1)}$ in $\mathcal{A}$. Let us fix a counter $i \in S_{j-1}^{\mu+1}$. By definition of $S_{j-1}^{\mu+1}$, there exists $\mathcal{B}^\mu \in \{\mathcal{B}_1^\mu, \dots, \mathcal{B}_{m_\mu}^\mu\}$ such that C_i is beyond F_{k-1} in $\mathcal{B}^\mu[S]$. We need to show that C_i is beyond $F_{k-1}^{(\mu+1)}$ in $\mathcal{A}$.

Since $T\langle\mu\rangle \subseteq T\langle\mu-1\rangle$, there exists a component $\mathcal{B}^{\mu-1}$ of $\mathcal{A}_{T\langle\mu-1\rangle}$ such that $\mathcal{B}^\mu \subseteq \mathcal{B}^{\mu-1}$, i.e., $\mathcal{B}^\mu$ is a sub-VASS of $\mathcal{B}^{\mu-1}$. By assumption (A2), $\mathcal{B}^\mu[S]$ is well-behaving for $k-1$. Hence, there exist positive constants $c, \kappa, \alpha \in \mathbb{N}$ and a control state p of $\mathcal{B}^\mu$ such that for all sufficiently large n , there is a finite computation $\beta_n = p\vec{n}, \dots, p\vec{v}$ of length at most $F_{k-1}(\alpha \cdot n)$ where $v(i) \geq F_{k-1}(\lfloor n/c \rfloor)$ for all counters i with complexity beyond F_{k-1} in $\mathcal{B}^\mu[S]$, and the associated cycle ϱ_n in $\mathcal{B}^\mu[S]$ satisfies the conditions of Definition 3.2 (see Fig. 2). In particular,

- the length of the “inner cycle” obtained by concatenating $\tau_1, \dots, \tau_k$ is $\Theta(n)$;
- for every “outer” cycle γ_j and every counter i with complexity not beyond F_{k-1} (i.e., below F_{k-1} by assumption (A3)), we have that $\Delta(\gamma_j)(i) = 0$, and the intermediate decrease of i is bounded by κ along γ_j .

Let J be the set of all counters whose complexity is beyond F_{k-1} in $\mathcal{B}^\mu[S]$. Observe the following:

- (1) The complexity of all counters in J is beyond F_{k-1} also in $\mathcal{B}^{\mu-1}[S]$, because $\mathcal{B}^\mu$ is a sub-VASS of $\mathcal{B}^{\mu-1}$. By induction hypothesis, this implies that the complexity of all counters in J is beyond $F_{k-1}^{(\mu)}$ in $\mathcal{A}$. Hence, $J \subseteq I\langle\mu\rangle$.
- (2) Since $J \subseteq I\langle\mu\rangle$, the effect of all “outer” cycles γ_j on all counters in $I\langle\mu\rangle \setminus J$ is zero, and the intermediate decrease of all these counters stays bounded by κ along γ_j . The effect of the “inner” cycle on the counters of $J \subseteq I\langle\mu\rangle$ may be negative.
- (3) The effect of the “inner” cycle may be negative also on the counters not included in $I\langle\mu\rangle$.

We show how to modify the “inner” cycle so that the total effect of the modified inner cycle on all counters not included in $I\langle\mu\rangle$ is non-negative, and the intermediate decrease of these counters along the modified inner cycle is bounded by some constant. Recall that all transitions of $\mathcal{B}^\mu$ are included in $T\langle\mu\rangle$, and for every $t \in T\langle\mu\rangle$ there exists a cycle of the form t, ϱ_t initiated by the transition t such that $\Delta(t, \varrho_t)^- \subseteq I\langle\mu\rangle$.

The “inner cycle” is modified in the following way. We pick an initial state and follow the inner cycle until we encounter the first simple cycle $\xi = t_1, \dots, t_k$. If $\Delta(\xi)$ is non-negative in the counters not included in $J \cup S$, then all transitions ξ are declared as “processed”. Otherwise, we construct a “compensating” cycle by concatenating $\varrho_{t_k}, \dots, \varrho_{t_1}$ and insert this cycle right after ξ (see Fig. 3). Transitions of this compensating cycle are also declared as “processed”. Then, we restart the procedure, but we skip all “processed” transitions when searching for a simple cycle from the initial state. Thus, we

produce a modified inner cycle whose length is still $\Theta(n)$ and the above properties are satisfied.

Now we show that if C_i is beyond F_{k-1} in $\mathcal{B}^\mu[S]$, then C_i is beyond $F_{k-1}^{(\mu+1)}$ in $\mathcal{A}$. To see this, for every (sufficiently large) $n \in \mathbb{N}$, consider a computation which first pumps all counters of S beyond $F_{k-1}^{(\mu+1)}(\lfloor n/c \rfloor)$ and all counters of $I(\mu)$ beyond $F_{k-1}^{(\mu)}(\lfloor n/c \rfloor)$ simultaneously, keeping the other counters above $\lfloor n/c \rfloor$. Here, c is a suitable constant, see Claim 1. Then, the computation continues by executing the computation β_m where $m = F_{k-1}^{(\mu)}(\lfloor n/c' \rfloor)$ and the “inner cycle” is modified in the way described above (the counter updates of every transition in β_m are the ones of $\mathcal{A}$, not the ones of $\mathcal{B}^\mu[S]$). The constant c' is chosen so that all counters stay positive when executing β_m (such a c' certainly exists due to the properties of Definition 3.2). Thus, all counters of J are pumped to a value beyond $F_{k-1}^{(\mu+1)}$ in $\mathcal{A}$.

Let $S = \bigcup_{j=1}^\infty S_j^{\mu+1}$. Clearly, the complexity of every counter outside S is not above F_{k-1} in every $\mathcal{B}^\mu[S]$, where $\mathcal{B}^\mu \in \{\mathcal{B}_1^\mu, \dots, \mathcal{B}_m^\mu\}$ (otherwise, the counter would be included in S). Hence, the complexity of each such counter is below F_{k-1} in every $\mathcal{B}^\mu[S]$ by Assumption (A3). Hence, the complexity is bounded by a function of the form $F_{k-2}^{(b)}$ for some $b \in \mathbb{N}$. Since there are only finitely many $\mathcal{B}^\mu$ and finitely many counters outside S , we can choose the same b for all $\mathcal{B}^\mu$ and all counters outside S .

Recall that the complexity of all transitions not in $T(\mu)$ is below $F_{k-1}^{(\mu)}$ in $\mathcal{A}$. Hence, for every n and every computation of $\mathcal{A}$ initiated in a configuration $q\vec{n}$, the computation may execute at most $F_{k-1}^{(\mu-1)}(F_{k-2}^{(a)}(n))$ transitions not included in $T(\mu)$, where $a \in \mathbb{N}$ is some constant. This means that the computation can be split in continuous sub-computations each of which is either outside or inside some VASS of $\{\mathcal{B}_1^\mu, \dots, \mathcal{B}_m^\mu\}$. The structure is shown in Fig. 4, where the sub-computations inside some VASS of $\{\mathcal{B}_1^\mu, \dots, \mathcal{B}_m^\mu\}$ are depicted as loops, and the sub-computations outside $\{\mathcal{B}_1^\mu, \dots, \mathcal{B}_m^\mu\}$ as connecting paths among these loops. Our aim is to show that the value of all counters outside S is bounded by $f(n)$ along this computation, where $f : \mathbb{N} \rightarrow \mathbb{N}$ is a function below $F_{k-1}^{(\mu+1)}$.

In the worst case, a sub-computation of some $\mathcal{B}_j^\mu$ is invoked after every execution of a transition not in $T(\mu)$. The sub-computation in $\mathcal{B}_j^\mu$ can be performed also in $\mathcal{B}_j^\mu[S]$ by following the same sequence of transitions, which implies that the increase of all counters outside S in this sub-computation is bounded by $F_{k-2}^{(b)}$ (see above). Hence, the total increase in every counter outside S along the whole computation is bounded by

$$\underbrace{F_{k-2}^{(b)} \circ \dots \circ F_{k-2}^{(b)}}_{F_{k-1}^{(\mu)}(F_{k-2}^{(a)}(n))} = F_{k-1}(b \cdot F_{k-1}^{(\mu-1)}(F_{k-2}^{(a)}(n))) \leq F_{k-1}^{(\mu)}(F_{k-2}^{(e)}(n))$$

where $e \in \mathbb{N}$ is a suitable constant. Thus, we obtain that the complexity of all counters outside S is below $F_{k-1}^{(\mu+1)}$ in $\mathcal{A}$. $\square$

Figure 4. The structure of a computation in $\mathcal{A}$.

Let $\mu \in \mathbb{N}$. Consider again the sets $S_j^{\mu+1}$ defined in Lemma 3.3, and observe that if $S_j^{\mu+1} = \emptyset$, then $S_{j'}^{\mu+1} = \emptyset$ for all $j' \geq j$. Furthermore, there exists $j \leq d$ such that $S_j^{\mu+1} = \emptyset$. The inner loop of the procedure **LargeCounters** at lines 11–16 computes the set $\bigcup_{j=1}^\infty S_j^{\mu+1}$ in polynomial time and stores this set into the variable Ω (observe that the content of Aux at line 15 after j iterations of the inner loop is precisely $S_j^{\mu+1}$, assuming that the variable T has been previously assigned the set $T(\mu)$ at line 9).

Now, let us consider the counters contained in the variable I after completing the outer loop at lines 8–18. A direct consequence of Lemma 4 is that all counters not contained in I are below F_k in $\mathcal{A}$. Also observe that for each such counter i , we obtain the largest μ such that C_i is beyond $F_{k-1}^{(\mu)}$, and the same holds for transitions whose complexity is below F_k in $\mathcal{A}$.

It remains to show that C_i is beyond F_k for every $i \in I$, and $\mathcal{A}$ is well-behaving for k . This is proven in the next lemma.

Lemma 3.4. *Let $\mu \in \mathbb{N}$ be the least index such that $I(\mu) = I(\mu+1)$. Then C_i is beyond F_k in $\mathcal{A}$ for every $i \in I(\mu)$. Furthermore, $\mathcal{A}$ is well-behaving for k .*

Proof. Since $I(\mu) = I(\mu+1)$, the set $I(\mu)$ can be split into pairwise disjoint subsets

$$I(\mu) = S_1 \uplus \dots \uplus S_\ell$$

so that for every counter $i \in S_j$, where $1 \leq j \leq \ell$, there exists $\mathcal{B} \in \{\mathcal{B}_1^\mu, \dots, \mathcal{B}_{m_\mu}^\mu\}$ such that C_i is beyond F_{k-1} in $\mathcal{B}[S_1 \cup \dots \cup S_{j-1}]$ (see Lemma 3.3). Hence, for every $j \in \{1, \dots, \ell\}$, there exists a sequence

$$\lambda_j = \mathcal{B}_1^\mu[S_1 \cup \dots \cup S_{j-1}], \dots, \mathcal{B}_{|S_j|}^\mu[S_1 \cup \dots \cup S_{j-1}]$$

such that for every counter $i \in S_j$ there is a $\mathcal{B}[S_1 \cup \dots \cup S_{j-1}]$ in λ_j where C_i is beyond F_{k-1} in $\mathcal{B}[S_1 \cup \dots \cup S_{j-1}]$. Let

$$\lambda = \mathcal{B}_1^\mu[K_1], \dots, \mathcal{B}_{|I(\mu)|}^\mu[K_{|I(\mu)|}]$$

be the sequence obtained by concatenating the sequences $\lambda_1, \dots, \lambda_\ell$. Recall that for every $\mathcal{B}_x^\mu[K_x]$ in the sequence λ there exist positive constants $c, \kappa, \alpha \in \mathbb{N}$ and a control state p such that for all sufficiently large n , there is a finite computation β_n initiated in $p\vec{n}$ satisfying the conditions of Definition 3.2. Furthermore, the computation β_n can be modified by inserting the “compensating cycles” in the way described in the proof of Lemma 3.3. The modified β_n does not decrease the counters not contained $I(\mu)$, and the intermediate decrease in these counters is bounded by a fixed constant

Figure 5. The arrangement of η_x^n cycles.

independent of n . Let η_x^n be the path in $\mathcal{A}$ associated to the modified β_n . Since $\mathcal{A}$ is strongly connected, we can arrange all η_x^n in the way shown in Fig. 5 and connect them by short paths whose length is bounded by a the number of control states of $\mathcal{A}$. For all sufficiently large n , the computation associated to this path increases all counters in $I(\mu)$ beyond F_{k-1} , and the other counters are (possibly) decreased only by transitions of the short connecting paths, i.e., just by a constant. Hence, we can perform $\Omega(n)$ such paths from a configuration $q(\vec{n})$. The first path increases the counters of $I(\mu)$ beyond F_{k-1} , the next path beyond $F_{k-1}^{(2)}$, etc., and after performing the last path, the counters of $I(\mu)$ are increased beyond $F_{k-1}^{(\lfloor n/c' \rfloor)}$ for some constant c' . This means that the complexity of all counters of $I(\mu)$ is beyond F_k . Furthermore, the path of Fig. 5 satisfies the requirements of Definition 3.2. $\square$

This completes the correctness proof for the procedure **LargeCounters**. The main repeat-until loop at lines 8–16 is a simple greatest fixed-point computation terminating after at most d iterations. The set T at line 9 can be computed in polynomial time using the algorithm of [15]. The inner repeat-until loop at lines 11–14 also requires at most d iterations, and the total number of all $\mathcal{B}$'s determined by $\mathcal{A}_T$ (see line 12) is bounded by $|Q|$. Since the total number of recursive calls to **LargeCounters** invoked by one iteration of the main repeat-until loop at lines 8–16 is bounded by $d^2 \cdot |Q|$, the total number of recursive calls invoked by the procedure is bounded by $(d^2 \cdot |Q|)^k$ for every k . Hence, the complexity of **LargeCounters** is polynomial for every fixed $k \in \mathbb{N}$.

It follows already from the results of [19] that the termination complexity of a terminating d -dimensional VASS is in $\mathcal{G}_{d+1}$. Hence, the least index k' such that $\mathcal{L} \in \mathcal{G}_{k'}$ is computable by running our algorithm for $k = d + 1$. This yields

complexity exponential in d , but polynomial for every fixed dimension.

Finally, we observe that the result about termination complexity obtained for strongly connected demonic VASS can be easily extended to arbitrary (not necessarily strongly connected) demonic VASS.

Lemma 3.5. *Let $\mathcal{A}$ be a d -dimensional demonic VASS and $k \geq 3$. We have that $\mathcal{L}_{\mathcal{A}} \in \mathcal{G}_k$ iff $\mathcal{L}_{\mathcal{B}} \in \mathcal{G}_k$ for every SCC $\mathcal{B}$ of $\mathcal{A}$.*

Proof. The “ $\Rightarrow$ ” direction is trivial. For the “ $\Leftarrow$ ” direction, we assume that in every SCC $\mathcal{B}$ of $\mathcal{A}$, every counter can be increased from $\vec{n}$ to at most $F_k^{(\mu_{\mathcal{B}})}(n)$ for some $\mu_{\mathcal{B}} \in \mathbb{N}$. Let $\mu = \max_{\mathcal{B}} \mu_{\mathcal{B}}$. Now consider some fixed computation initiated in a configuration $q\vec{n}$, and let $\mathcal{B}_1, \dots, \mathcal{B}_m$ be sequence of all SCCs visited by this computation in this given order (note that $m \leq |Q|$). For simplicity, we disregard the effects of all transitions outside of these SCCs (their total number is bounded by $|Q| - 1$, which is a constant independent of n).

By induction on r , we prove that the largest counter value reachable in $\mathcal{B}_r$ is at most $F_k^{(\mu \cdot r)}(n)$ for every $1 \leq r \leq m$. For $\mathcal{B}_1$ this is trivial. If the claim holds for $r-1$ then the value of all counters is at most $F_k^{(\mu \cdot (r-1))}(n)$ when leaving $\mathcal{B}_{r-1}$, and we may take this value as an overapproximate of the initial value of all counters when entering $\mathcal{B}_r$. Then the largest reachable value for any counter is $F_k^{(\mu)}(F_k^{(\mu \cdot (r-1))}(n)) = F_k^{(\mu \cdot r)}(n)$.

Since all counter values along our fixed computation are bounded by $F_k^{(\mu \cdot |Q|)}(n)$, the total number of configurations visited along the computation is at most $|Q| \cdot F_k^{(\mu \cdot |Q|)}(n)$. Since no configuration can be visited twice (otherwise $\mathcal{L}_{\mathcal{B}} \in \mathcal{G}_k$ does not hold), we are done. $\square$

The “precise” complexity of a given transition in non-strongly connected demonic VASS cannot be trivially deduced by analyzing the transition/counter complexities in the SCCs of $\mathcal{A}$. For example, if $\mathcal{T}_t \in \mathcal{G}_5$ in some SCC $\mathcal{B}$ of $\mathcal{A}$ and just some counters can be pumped beyond F_6 in some previous SCC $\mathcal{B}'$ of $\mathcal{A}$, this may but does not have to be sufficient for pushing the complexity of t beyond F_6 in $\mathcal{A}$.

In the next theorem, we summarize the results proven in this section.

Theorem 3.6. *Let $k \geq 3$. Then the problem whether $\mathcal{L}$ is beyond F_k for a given demonic VASS is in $\mathbf{P}$. Furthermore, if $\mathcal{L}$ is not beyond F_k , then $\mathcal{L} \in \mathcal{G}_{k-1}$.*

If $\mathcal{A}$ is strongly connected, we have the following additional results:

- For a given counter i , the problem whether C_i is beyond F_k is in $\mathbf{P}$. If C_i is not beyond F_k , then $C_i \in \mathcal{G}_{k-1}$, and the least $k' \leq k - 1$ such that $C_i \in \mathcal{G}_{k'}$ is computable in polynomial time. Furthermore, there exists the largest $\mu \in \mathbb{N}$ computable in polynomial time such that C_i is beyond $F_{k'}^{(\mu)}$.

- For a given transition t , the problem whether $\mathcal{T}_t$ is beyond F_k is in **P**. If $\mathcal{T}_t$ is not beyond F_k , then $\mathcal{T}_t \in \mathcal{G}_{k-1}$, and the least $k' \leq k-1$ such that $\mathcal{T}_t \in \mathcal{G}_{k'}$ is computable in polynomial time. Furthermore, there exists the largest $\mu \in \mathbb{N}$ computable in polynomial time such that $\mathcal{T}_t$ is beyond $F_{k'}^{(\mu)}$.

4 Results about VASS Games

In this section, we prove that for every $k \geq 1$, the problem whether $\mathcal{L}_{\mathcal{A}} \in \mathcal{G}_k$ for a given VASS game $\mathcal{A}$ is **NP**-complete. Furthermore, we show that if $\mathcal{L}_{\mathcal{A}} \notin \mathcal{G}_k$, then $\mathcal{L}_{\mathcal{A}}$ is beyond F_{k+1} . Hence, the dichotomy of termination complexity established previously for demonic VASS holds even for VASS games.

We start by introducing some notions. Let $\mathcal{A} = (Q, \text{Tran})$ be a VASS game. For every $q \in Q$, let $\text{Tval}_{\mathcal{A}}^q : \mathbb{N} \rightarrow \mathbb{N}_{\infty}$ be a function defined by

$$\text{Tval}_{\mathcal{A}}^q(n) = \text{Tval}_{\mathcal{A}}(q\vec{n}).$$

Furthermore, a strategy η for player A in $\mathcal{A}$ is *counterless* if it depends only on the currently visited control state. Formally, a counterless strategy is a function $\eta : Q_A \rightarrow \text{Tran}$ such that $\eta(p)$ is an out-going transition of p for every $p \in Q_A$. We use $\mathcal{A}^\eta$ to denote the VASS obtained by “applying” η to $\mathcal{A}$, i.e., restricting the set of out-going transitions in every $p \in Q_A$ to $\eta(p)$. Note that $\mathcal{A}^\eta$ can be seen as a *demonic* VASS because player A has no real choice in $\mathcal{A}^\eta$.

Our first lemma classifies the growth of $\text{Tval}_{\mathcal{A}}^q(n)$ in demonic VASS.

Lemma 4.1. *Let $k \geq 1$, and let $\mathcal{A}$ be a demonic VASS. For every control state q of $\mathcal{A}$, the function $\text{Tval}_{\mathcal{A}}^q(n)$ is either in $\mathcal{G}_k$ or beyond F_{k+1} .*

Proof. For every SCC $\mathcal{B}$ of $\mathcal{A}$ and all control states q, r of $\mathcal{B}$ we have that $\text{Tval}_{\mathcal{B}}^q(n)$ and $\text{Tval}_{\mathcal{B}}^r(n)$ are simultaneously either in $\mathcal{G}_k$ or beyond F_{k+1} . The rest of the argument is the same as in the proof of Lemma 3.5. Note that Lemma 3.5 holds also for $k = 1, 2$ (recall that if the termination complexity of a given demonic VASS is not linear/polynomial, then it is at least quadratic/exponential, respectively [3, 16]). $\square$

Now we prove the following:

Theorem 4.2. *Let $k \geq 1$, and let $\mathcal{A}$ be a VASS game. For every control state q of $\mathcal{A}$, the function $\text{Tval}_{\mathcal{A}}^q(n)$ is either in $\mathcal{G}_k$ or beyond F_{k+1} . Furthermore, there exists a counterless strategy η for player A in $\mathcal{A}$ such that for every control state q we have that $\text{Tval}_{\mathcal{A}}^q(n) \in \mathcal{G}_k$ iff $\text{Tval}_{\mathcal{A}^\eta}^q(n) \in \mathcal{G}_k$.*

Proof. We proceed by induction on the number of angelic control states with at least two outgoing transitions.

If every angelic control state has only one outgoing transition, then player A has only one trivial strategy η , which is counterless. The result follows by applying Lemma 4.1 to $\mathcal{A}^\eta = \mathcal{A}$.

Now assume there exists $q \in Q_A$ with a set of outgoing transitions $\text{Tran}_q = \{u_1, \dots, u_c\}$ such that $c \geq 2$. For every $i \in \{1, \dots, c\}$, consider a VASS $\mathcal{A}_i$ obtained from $\mathcal{A}$ by restricting the set Tran_q to $\{u_i\}$. Furthermore, let $\mathcal{A}_q$ be a VASS obtained from $\mathcal{A}$ by putting $\text{Tran}_q = \{(q, -\vec{\kappa}, q)\}$ where κ is the maximal absolute value of integer occurring in the transition update vectors of $\mathcal{A}$. Note that the induction hypothesis is applicable to all of the constructed $\mathcal{A}_q, \mathcal{A}_1, \dots, \mathcal{A}_c$. Furthermore, we have the following:

Claim 2. *For all $p \in Q$, $i \in \{1, \dots, c\}$, and $n \in \mathbb{N}$,*

- (A) *if $\text{Tval}_{\mathcal{A}_q}^p(n)$ is beyond F_{k+1} , then $\text{Tval}_{\mathcal{A}}^p(n)$ is beyond F_{k+1} ;*
- (B) *$\text{Tval}_{\mathcal{A}}^p(n) \leq \text{Tval}_{\mathcal{A}_q}^p(n) + \text{Tval}_{\mathcal{A}_i}^q(\kappa \cdot \text{Tval}_{\mathcal{A}_q}^p(n) + n)$.*

Proof of the claim. For a given pair of strategies (σ, π) and an initial configuration ru , we use $q\text{-Comp}^{\sigma, \pi}(\text{ru})$ to denote the maximal computation determined by (σ, π) initiated in ru such that the state q is either not visited along the computation or q appears only in the last configuration of the computation.

We prove (A) by showing that, for every $n \in \mathbb{N}$,

$$\text{Tval}_{\mathcal{A}}^p(n) \geq \frac{1}{2}(\text{Tval}_{\mathcal{A}_q}^p(n) - \frac{n}{\kappa}).$$

Let π^* be an optimal positional strategy for player D in $\mathcal{A}_q$ (see Section 2). Note that π^* is applicable also to $\mathcal{A}$, and $\text{Tval}_{\mathcal{A}}^p(n) \geq \inf_{\sigma} |\text{Comp}_{\mathcal{A}}^{\sigma, \pi^*}(p\vec{n})|$. Let σ be a strategy of player A in $\mathcal{A}$. Consider the corresponding strategy $\hat{\sigma}$ of player A in $\mathcal{A}_q$ behaving in the same way as σ until a configuration of the form qu is visited; from this point on, $\hat{\sigma}$ selects the only available transition $(q, -\vec{\kappa}, q)$. Clearly,

$$\begin{aligned} \text{Tval}_{\mathcal{A}}^p(n) &\geq |\text{Comp}_{\mathcal{A}}^{\sigma, \pi^*}(p\vec{n})| \\ &\geq |q\text{-Comp}_{\mathcal{A}}^{\sigma, \pi^*}(p\vec{n})| \\ &= |q\text{-Comp}_{\mathcal{A}_q}^{\hat{\sigma}, \pi^*}(p\vec{n})|. \end{aligned}$$

Furthermore, $|\text{Comp}_{\mathcal{A}_q}^{\hat{\sigma}, \pi^*}(p\vec{n})|$ is bounded from above by

$$|q\text{-Comp}_{\mathcal{A}_q}^{\hat{\sigma}, \pi^*}(p\vec{n})| + (\kappa \cdot |q\text{-Comp}_{\mathcal{A}_q}^{\hat{\sigma}, \pi^*}(p\vec{n})| + n)/\kappa$$

Note that the second summand bounds the number of executions of $(q, -\vec{\kappa}, q)$. Thus, we obtain

$$|q\text{-Comp}_{\mathcal{A}_q}^{\hat{\sigma}, \pi^*}(p\vec{n})| \geq \frac{1}{2}(|\text{Comp}_{\mathcal{A}_q}^{\hat{\sigma}, \pi^*}(p\vec{n})| - \frac{n}{\kappa}),$$

and finally

$$\text{Tval}_{\mathcal{A}}^p(n) \geq \frac{1}{2}(|\text{Comp}_{\mathcal{A}_q}^{\hat{\sigma}, \pi^*}(p\vec{n})| - \frac{n}{\kappa}) \geq \frac{1}{2}(\text{Tval}_{\mathcal{A}_q}^p(n) - \frac{n}{\kappa}).$$

In the last inequality, we use the optimality of π^* .

(B) Let π^* be an optimal positional strategy for player D in $\mathcal{A}_i$ (note that π^* is applicable also to $\mathcal{A}_q$), and let σ^* be an optimal strategy for player A in $\mathcal{A}_q$. Furthermore, let $\hat{\sigma}^*$ be a strategy for player A in $\mathcal{A}_i$ behaving like σ^* until a configuration of the form qu is visited, and then switching to an optimal strategy for player A in $\mathcal{A}_i$. Since π^* is optimal, for every $n \in \mathbb{N}$ we have that $\text{Tval}_{\mathcal{A}_i}^p(n) \leq$

$|Comp_{\mathcal{A}_i}^{\hat{\sigma}^*, \pi^*}(p\vec{n})|$, and the computation $Comp_{\mathcal{A}_i}^{\hat{\sigma}^*, \pi^*}(p\vec{n})$ can be split into a prefix $q\text{-}Comp_{\mathcal{A}_i}^{\hat{\sigma}^*, \pi^*}(p\vec{n})$ followed by a (possibly empty) suffix initiated in a configuration of the form qu . Clearly,

$$\begin{aligned} |q\text{-}Comp_{\mathcal{A}_i}^{\hat{\sigma}^*, \pi^*}(p\vec{n})| &= |q\text{-}Comp_{\mathcal{A}_q}^{\hat{\sigma}^*, \pi^*}(p\vec{n})| \\ &\leq |Comp_{\mathcal{A}_q}^{\hat{\sigma}^*, \pi^*}(p\vec{n})| \\ &\leq Tval_{\mathcal{A}_q}^p(n) \quad (\hat{\sigma}^* \text{ is optimal}) \end{aligned}$$

The length of the suffix initiated in qu is bounded by $Tval_{\mathcal{A}_i}^q(\kappa \cdot Tval_{\mathcal{A}_q}^p(n) + n)$, because every component of $\mathbf{u}$ is bounded by $\kappa \cdot Tval_{\mathcal{A}_q}^p(n) + n$ and the strategy $\hat{\sigma}^*$ switches to an optimal strategy for player A in $\mathcal{A}_i$ after visiting qu . $\square$

Now we continue with the main proof. We distinguish two possibilities.

Case (1) $Tval_{\mathcal{A}_i}^q(n) \in \mathcal{G}_k$ for some $i \in \{1, \dots, c\}$. We show that, for every $p \in Q$,

- (a) if $Tval_{\mathcal{A}_i}^p(n) \in \mathcal{G}_k$, then $Tval_{\mathcal{A}}^p(n) \in \mathcal{G}_k$;
- (b) if $Tval_{\mathcal{A}_i}^p(n)$ is beyond F_{k+1} , then $Tval_{\mathcal{A}}^p(n)$ is beyond F_{k+1} .

Note that (a) and (b) imply the first part of our theorem by applying induction hypothesis to $\mathcal{A}_i$, and the counterless strategy η for player A in $\mathcal{A}$ can be constructed simply by putting $\eta = \eta_i$ (note that $\eta(q) = u_i$). So, it remains to prove (a) and (b).

(a) Assume $Tval_{\mathcal{A}_i}^p(n) \in \mathcal{G}_k$. By induction hypothesis, there exists a counterless strategy η_i for player A in $\mathcal{A}_i$ such that $Tval_{\mathcal{A}_i}^p(n) \in \mathcal{G}_k$. The strategy η_i is applicable also to $\mathcal{A}$. Since $\mathcal{A}^{\eta_i} = \mathcal{A}_i^{\eta_i}$, we are done.

(b) Assume $Tval_{\mathcal{A}_i}^p(n)$ is beyond F_{k+1} . By induction hypothesis, we have that either $Tval_{\mathcal{A}_q}^p(n) \in \mathcal{G}_k$ or $Tval_{\mathcal{A}_q}^p(n)$ is beyond F_{k+1} . The first possibility implies $Tval_{\mathcal{A}_i}^p(n) \in \mathcal{G}_k$ by using Claim 2(B), which is a contradiction. Hence, $Tval_{\mathcal{A}_q}^p(n)$ is beyond F_{k+1} , and by applying Claim 2(A) we obtain that $Tval_{\mathcal{A}}^p(n)$ is beyond F_{k+1} .

Case (2) $Tval_{\mathcal{A}_i}^q(n)$ is beyond F_{k+1} for all $i \in \{1, \dots, c\}$. We show that, for every $p \in Q$,

- (c) if $Tval_{\mathcal{A}_i}^p(n) \in \mathcal{G}_k$ for some $i \in \{1, \dots, c\}$, then $Tval_{\mathcal{A}}^p(n) \in \mathcal{G}_k$ for every $j \in \{1, \dots, c\}$.
- (d) if $Tval_{\mathcal{A}_i}^p(n)$ is beyond F_{k+1} for all $i \in \{1, \dots, c\}$, then $Tval_{\mathcal{A}}^p(n)$ is beyond F_{k+1} .

Note that (c) and (d) imply the first part of our theorem, and the required counterless strategy η for player A in $\mathcal{A}$ can be constructed by putting $\eta = \eta_j$ (for an arbitrary $j \in \{1, \dots, c\}$). Now we prove (c) and (d).

(c) Suppose $Tval_{\mathcal{A}_i}^p(n) \in \mathcal{G}_k$ for some $i \in \{1, \dots, c\}$, and consider a counterless strategy η_i for player A in $\mathcal{A}_i$ obtained by applying the induction hypothesis to $\mathcal{A}_i$ (note that

$Tval_{\mathcal{A}_i}^p(n) \in \mathcal{G}_k$). We show that the state q is *not* reachable from the state p in the underlying graph of $\mathcal{A}_i^{\eta_i}$. Suppose the converse. For every initial configuration $p\vec{n}$, consider the strategy π_n of player D defined as follows: For all $n < \kappa \cdot |Q|$, the strategy π_n is defined arbitrarily. For every $n \geq \kappa \cdot |Q|$, player D first follows the shortest path from p to q in the underlying graph of $\mathcal{A}_i^{\eta_i}$. Thus, he reaches a configuration qu where every component of $\mathbf{u}$ is at least $n - \kappa \cdot |Q|$. Then, player D switches to his *optimal* strategy in qu . Clearly, $|Comp_{\mathcal{A}_i}^{\pi_n}(p\vec{n})|$, when interpreted as a function of n , is beyond F_{k+1} , hence $Tval_{\mathcal{A}_i}^p(n)$ is beyond F_{k+1} , and we have a contradiction.

For an arbitrary $j \in \{1, \dots, c\}$, consider a counterless strategy $\hat{\eta}_j$ for player A in $\mathcal{A}_j$ obtained from η_i by redefining $\hat{\eta}_j(q)$ to u_j . For every initial configuration of the form $p\vec{n}$, there is no difference between using the strategies η_i and $\hat{\eta}_j$ in $\mathcal{A}_i$ and $\mathcal{A}_j$, respectively, because the only state q where these strategies differ is never reached. Consequently, $Tval_{\mathcal{A}_j}^p(n) \in \mathcal{G}_k$, which proves (c).

(d) Now assume $Tval_{\mathcal{A}_i}^p(n)$ is beyond F_{k+1} for all $i \in \{1, \dots, c\}$. We say that a cycle $q, u, \dots, q$ in $\mathcal{A}$ is *i-offending* if $u \neq u_i$ and all control states in the cycle except for the first and the last one are different from q .

For every $i \in \{1, \dots, c\}$, let π_i^* be an optimal positional strategy of player D in $\mathcal{A}_i$ (see Section 2.2). Consider the strategy π for player D in $\mathcal{A}$ defined as follows: for a given finite computation $p_1\mathbf{v}_1, \dots, p_m\mathbf{v}_m$ where $p_m \in Q_D$, let $p_1, t_1, \dots, p_m$ be the corresponding path in $\mathcal{A}$, and let $z < m$ be the maximal index such that $p_z = q$. If there is no such z , we put $j = 1$. Otherwise, $t_z \in \{u_1, \dots, u_c\}$, and j is the unique index such that $t_z = u_j$. Let ϱ be the finite path obtained from $p_1, t_1, \dots, p_m$ by deleting all j -offending cycles. Note that the first/last state of ϱ is p_1/p_m , and ϱ is a path in $\mathcal{A}_j$. Consider the configuration $p_1(\mathbf{v}_1/c)$ where $\mathbf{v}_1/c$ is the vector with i -th component equal to $\lfloor \mathbf{v}_1(i)/c \rfloor$. If ϱ is executable from $p_1(\mathbf{v}_1/c)$ in $\mathcal{A}_j$ and determines a computation α , we put $\pi(p_1\mathbf{v}_1, \dots, p_m\mathbf{v}_m) = \pi_j^*(\alpha)$. Otherwise, $\pi(p_1\mathbf{v}_1, \dots, p_m\mathbf{v}_m)$ is defined arbitrarily. We show the following:

Claim 3. *Let σ be a strategy of player A in $\mathcal{A}$. For every $n \in \mathbb{N}$, there are strategies $\sigma_1, \dots, \sigma_c$ for player A in $\mathcal{A}_1, \dots, \mathcal{A}_c$, respectively, such that*

$$|Comp_{\mathcal{A}}^{\sigma, \pi}(p\vec{n})| \geq \min_{i \in \{1, \dots, c\}} \left\{ |Comp_{\mathcal{A}_i}^{\sigma_i, \pi_i^*}(p(\vec{n}/c))| \right\}.$$

Proof of the claim. Let $Comp^{\sigma, \pi}(p\vec{n}) = p_1\mathbf{v}_1, \dots, p_m\mathbf{v}_m$, and let $p_1, t_1, \dots, p_m$ be the corresponding path in $\mathcal{A}$. Our aim is to show that $p_1, t_1, \dots, p_m$ contains a “scattered subsequence” corresponding to a maximal computation in some $\mathcal{A}_i$ initiated in $p(\vec{n}/c)$ induced by the optimal strategy π_i^* of player D and some strategy σ_i of player A .

We say that a configuration $r\vec{v}$ is *terminal* if for every/some outgoing transition $(r, \mathbf{u}, s)$ of r , the vector $\mathbf{v} + \mathbf{u}$ is negative in some component, depending on whether r

is demonic/angelic, respectively. For every $i \in \{1, \dots, c\}$, consider the path obtained from $p_1, t_1, \dots, p_m$ by deleting all i -offending cycles. If the last occurrence of q in this path is followed by a transition different from u_i , then the whole suffix starting with this transition is also deleted. The resulting path is denoted by α_i . Observe that α_i is a path in $\mathcal{A}_i$ initiated in p_1 . If some prefix of α_i determines a computation in $\mathcal{A}_i$ initiated in $p_1(\vec{n}/c)$ ending in a terminal configuration, then this computation is maximal w.r.t. the optimal strategy π_i^* of player D and some strategy σ_i of player A (see the definition of π above). Hence,

$$|\text{Comp}_{\mathcal{A}}^{\sigma, \pi}(p\vec{n})| = m \geq |\text{Comp}_{\mathcal{A}_i}^{\sigma_i, \pi_i^*}(p(\vec{n}/c))|$$

and we are done. We complete our proof by showing that if every α_i determines a computation initiated in $p_1(\vec{n}/c)$ ending in a non-terminal configuration, then the considered computation $p_1 v_1, \dots, p_m v_m$ can be prolonged consistently with σ and π , which contradicts its maximality.

For all $i \in \{1, \dots, c\}$ and $j \leq m$, let $\mathbf{v}_j^i$ be a vector defined inductively as follows:

- $\mathbf{v}_1^i = \vec{n}/c$;
- if the transition t_j of the path $p_1, t_1, \dots, p_m$ has been deleted when constructing the path α_i (i.e., t_j appears in some i -offending cycle or a deleted suffix), then $\mathbf{v}_{j+1}^i = \mathbf{v}_j^i$; otherwise, $\mathbf{v}_{j+1}^i = \mathbf{v}_j^i + \mathbf{u}$, where $t_j = (p_j, \mathbf{u}, p_{j+1})$.

Intuitively, $\mathbf{v}_1^i, \dots, \mathbf{v}_m^i$ is the sequence of vectors determined by executing α_i in $p_1(\vec{n}/c)$, but the same vector is repeated for the subsequences that have been deleted when constructing α_i . Hence, all $\mathbf{v}_1^i, \dots, \mathbf{v}_m^i$ are non-negative, and for all $j \leq m$ we have that $\mathbf{v}_j \geq \mathbf{v}_1^1 + \dots + \mathbf{v}_j^c$. We distinguish two possibilities:

- $p_m \neq q$. Let $i \in \{1, \dots, c\}$ be an index such that α_i ends in p_m . Since $\mathbf{v}_m \geq \mathbf{v}_m^i$ and $p_m \mathbf{v}_m^i$ is non-terminal in $\mathcal{A}_i$, we obtain that $p_m \mathbf{v}_m$ is non-terminal in $\mathcal{A}$ (note that the state p_m has the same set of out-going transitions in $\mathcal{A}$ and $\mathcal{A}_i$).
- $p_m = q$. Then α_i ends in q for all $i \in \{1, \dots, c\}$. Since $q \mathbf{v}_m^i$ is non-terminal in $\mathcal{A}_i$ and $\mathbf{v}_m \geq \mathbf{v}_m^i$ for every $i \in \{1, \dots, c\}$, the configuration $q \mathbf{v}_m$ is non-terminal in $\mathcal{A}$. $\square$

Using Claim 3, we can finish a proof of (d) as follows:

$$\begin{aligned} \text{Tval}_{\mathcal{A}}^p(n) &= \text{Tval}_{\mathcal{A}}(p\vec{n}) \\ &\geq \inf_{\sigma} |\text{Comp}_{\mathcal{A}}^{\sigma, \pi}(p\vec{n})| \quad (\text{for } \pi \text{ defined above}) \\ &\geq \min_{i \in \{1, \dots, c\}} \left\{ |\text{Comp}_{\mathcal{A}_i}^{\sigma_i, \pi_i^*}(p(\vec{n}/c))| \right\} \quad (\text{by Claim 3}) \\ &\geq \min_{i \in \{1, \dots, c\}} \left\{ \text{Tval}_{\mathcal{A}_i}^p(\lfloor n/c \rfloor) \right\} \quad (\pi_i^* \text{ is optimal}) \end{aligned}$$

Since $\text{Tval}_{\mathcal{A}_i}^p(n)$ is beyond F_{k+1} for every $i \in \{1, \dots, c\}$, we obtain that $\text{Tval}_{\mathcal{A}}^p(n)$ is also beyond F_{k+1} . $\square$

Theorem 4.3. *Let $k \geq 1$. The problem whether $\mathcal{L}_{\mathcal{A}} \in \mathcal{G}_k$ for a given VASS game $\mathcal{A}$ is **NP**-complete. Furthermore, if $\mathcal{L}_{\mathcal{A}} \notin \mathcal{G}_k$, then $\mathcal{L}_{\mathcal{A}}$ is beyond F_{k+1} .*

Proof. Due to Theorem 4.2, we know that $\mathcal{L}_{\mathcal{A}}$ is either in $\mathcal{G}_k$ or beyond F_{k+1} . Furthermore, $\mathcal{L}_{\mathcal{A}} \in \mathcal{G}_k$ iff there is a counterless strategy η of player A such that $\mathcal{L}_{\mathcal{A}\eta} \in \mathcal{G}_k$. Hence, we can decide whether $\mathcal{L}_{\mathcal{A}} \in \mathcal{G}_k$ by guessing η , constructing $\mathcal{A}\eta$, and deciding whether $\mathcal{L}_{\mathcal{A}\eta} \in \mathcal{G}_k$, which proves the membership to **NP**. The **NP**-hardness follows immediately from the construction used [9] to prove **coNP**-hardness of the unknown initial credit problem for multi-weighted two-player game structures¹⁰. For a given 3SAT propositional formula φ in CNF, a VASS game $\mathcal{A}_{\varphi}$ is constructed in [9] such that

- if φ is satisfiable, then player A has a (counterless) strategy achieving termination in $O(n)$ transitions for every configuration $p\mathbf{v}$ such that $\mathbf{v} \leq \vec{n}$, i.e., $\mathcal{L}_{\mathcal{A}_{\varphi}} \in O(n) \subseteq \mathcal{G}_k$;
- if φ is not satisfiable, then for every counterless strategy η of player A , the VASS $\mathcal{A}_{\varphi}^{\eta}$ is non-terminating, which implies $\mathcal{L}_{\mathcal{A}_{\varphi}} \notin \mathcal{G}_k$. $\square$

A natural question is whether Theorem 4.2 can be extended to other complexity classes such as $O(n^k)$. Recall that the problem whether $\mathcal{L}_{\mathcal{A}} \in O(n^k)$ for a given demonic VASS $\mathcal{A}$ is decidable in polynomial time, and if $\mathcal{L}_{\mathcal{A}} \notin O(n^k)$, then $\mathcal{L}_{\mathcal{A}} \in \Omega(n^{k+1})$ [20]. However, our proof of Theorem 4.2 requires $\mathcal{G}_k$ being closed under function composition, which does not hold for $O(n^k)$ and $k \geq 2$. Still, one may conjecture that the termination complexity of a given VASS game $\mathcal{A}$ is $O(n^k)$ iff there is a finite-memory strategy η for player A depending only on some *bounded* information about the sequence of visited control states such that $\mathcal{L}_{\mathcal{A}} \in O(n^k)$ iff $\mathcal{L}_{\mathcal{A}\eta} \in O(n^k)$, where $\mathcal{A}_{\eta}$ is obtained from $\mathcal{A}$ by encoding the finite information used by η into the state space of $\mathcal{A}$. If the size of the memory used by η was polynomial in the size of $\mathcal{A}$, the problem whether $\mathcal{L}_{\mathcal{A}} \in O(n^k)$ would be in **NP**. The next proposition shows that no such η exists unless **NP** = **PSPACE**.

Proposition 4.4. *The problem whether $\mathcal{L}_{\mathcal{A}} \in O(n^2)$ for a given VASS game $\mathcal{A}$ is **PSPACE**-hard, even if counter updates are encoded in unary.*

Proof. We reduce the **PSPACE**-complete QBF problem. An instance of QBF is a quantified Boolean formula of the form $\varphi \equiv \forall x_1 \exists x_2 \dots \exists x_n C_1 \wedge \dots \wedge C_m$ over the variables $x_1, \dots, x_n$ where n is even and every C_i is a clause (i.e., a disjunction of literals of the form x_i or $\bar{x}_i$). The question is whether φ is valid. We show how to construct a VASS $\mathcal{A}$ whose size is polynomial in $\|\varphi\|$ such that

- if φ is not valid, then $\mathcal{L}_{\mathcal{A}} \in O(n^2)$;
- if φ is valid, then $\mathcal{L}_{\mathcal{A}} \in \Omega(n^3)$.

¹⁰We thank Krishnendu Chatterjee for pointing out this argument.

Figure 6. The problem whether $\mathcal{L}_{\mathcal{A}} \in O(n^2)$ for a given VASS game $\mathcal{A}$ is **PSPACE**-hard.

For every literal $\ell \in \{x_i, \bar{x}_i \mid 1 \leq i \leq n\}$, the VASS $\mathcal{A}$ has two counters $b[\ell]$ and $c[\ell]$. Furthermore, there are auxiliary counters a_1 , a_2 , and e . The structure of $\mathcal{A}$ is shown in Fig. 6. All transitions outside the gadgets $\langle y, z \rangle$ keep all counters unchanged. In the gadget $\langle y, z \rangle$, each transition increments or decrements only the counters a_1 , a_2 , y , and z in the way shown in Fig. 6, the other counters are unchanged. Intuitively, the purpose of $\langle y, z \rangle$ is to “pump” the counter z to a higher value constrained by y . The values of a_1 and a_2 stay linear, so if the size of y is linear/quadratic, the counter z can be made quadratic/cubic by the gadget.

Assume the initial configuration $p_1 \vec{n}$ (one can easily check that this configuration has the highest termination value among all configurations of the form $q \vec{n}$). Initially, the players guess an assignment by entering either the gadget $\langle b[x_i], c[x_i] \rangle$ or the gadget $\langle b[\bar{x}_i], c[\bar{x}_i] \rangle$ for each $i \in \{1, \dots, n\}$. Thus, the assignment is represented by the counters $c[\ell]$ pumped to a quadratic value by player D . After that, player A chooses a clause in the control state q (he wishes to choose a clause not satisfied by the previously guessed assignment). Then, player D selects a literal ℓ in the chosen clause. If the literal ℓ is set to true by the chosen assignment, the associated $c[\ell]$ counter has been pumped to a quadratic value. This means that player D can use the $\langle c[\ell], e \rangle$ gadget to pump the counter e even to a cubic value. If φ is not valid, player A can always choose a clause where all literals ℓ are invalid and the associated counters $c[\ell]$ are linear. This means that player D can only achieve a quadratic increase of counters in a quadratic number of steps. If φ is valid, then player D can make each clause valid and always choose a literal ℓ such that $c[\ell]$ has been pumped to a quadratic value. Hence, he can make the counter e cubic in a cubic number of transitions. $\square$

5 Conclusions

We have shown that the problem whether the termination complexity of a given VASS belongs to a given level in the

Grzegorzcyk hierarchy is solvable in polynomial time, and the same problem is **NP**-complete for VASS games.

There are several possible directions for future work. First, the decidability/complexity of the problem whether $\mathcal{L}(\mathcal{A}) \in O(n^k)$ for a given VASS game and a given $k \geq 2$ remains open. Our **PSPACE**-hardness result indicates that additional insights are needed to resolve this question. Another open problem is the precise complexity of determining the *least* k such that $\mathcal{L}(\mathcal{A}) \in \mathcal{G}_k$ for a given demonic VASS $\mathcal{A}$ (or a given VASS game $\mathcal{A}$). As we already indicated in Section 1, it is still possible that the procedure **LargeCounters** needs to initiate only polynomially many recursive calls with distinct parameters, which would imply the computability of the k in polynomial time for demonic VASS.

Acknowledgments

Antonín Kučera and Dominik Velan are supported by the Czech Science Foundation, grant No. 18-11193S. Jérôme Leroux is supported by the grant ANR-17-CE40-0028 of the French National Research Agency ANR (project BRAVAS). The work is also supported by the European Network for Game Theory, COST Action CA16228 (GAMENET).

References

- [1] 2010. *Proceedings of ICALP 2010, Part II*. LNCS, Vol. 6199. Springer.
- [2] T. Brázdil, K. Chatterjee, A. Kučera, P. Novotný, and D. Velan. 2019. Deciding Fast Termination for Probabilistic VASS with Nondeterminism. In *Proceedings of ATVA 2019 (LNCS)*, Vol. 11781. Springer, 462–478.
- [3] T. Brázdil, K. Chatterjee, A. Kučera, P. Novotný, D. Velan, and F. Zuleger. 2018. Efficient Algorithms for Asymptotic Bounds on Termination Time in VASS. In *Proceedings of LICS 2018*. ACM Press, 185–194.
- [4] T. Brázdil, P. Jančar, and A. Kučera. 2010. Reachability Games on Extended Vector Addition Systems with States, See [1], 478–489.
- [5] T. Brázdil, A. Kučera, and P. Novotný. 2012. Determinacy in Stochastic Games with Unbounded Payoff Functions. In *Proceedings of 8th Doctoral Workshop on Mathematical and Engineering Methods in Computer Science (MEMICS 2012) (LNCS)*, Vol. 7721. 94–105.
- [6] M. Broy and M. Wirsing. 1981. On the Algebraic Specification of Nondeterministic Programming Languages. In *Proceedings of CAAP'81 (LNCS)*, Vol. 112. Springer, 162–179.

- [7] K. Chatterjee and L. Doyen. 2010. Energy Parity Games, See [1], 599–610.
- [8] K. Chatterjee and L. Doyen. 2011. Energy and Mean-Payoff Parity Markov Decision Processes. In *Proceedings of MFCS 2011 (LNCS)*, Vol. 6907. Springer, 206–218.
- [9] K. Chatterjee, L. Doyen, T. Henzinger, and J.-F. Raskin. 2010. Generalized Mean-payoff and Energy Games. In *Proceedings of FST&TCS 2010 (LIPIcs)*, Vol. 8. Schloss Dagstuhl, 505–516.
- [10] K. Chatterjee, M. Henzinger, S. Krinninger, and D. Nanongkai. 2012. Polynomial-Time Algorithms for Energy Games with Special Weight Structures. In *Proceedings of ESA 2012 (LNCS)*, Vol. 7501. Springer, 301–312.
- [11] U. Fahrenberg, L. Juhl, K. Larsen, and J. Srba. 2011. Energy Games in Multiweighted Automata. In *Proceedings of the 8th International Colloquium on Theoretical Aspects of Computing (ICTAC'11) (LNCS)*, Vol. 6916. Springer, 95–115.
- [12] A. Grzegorzcyk. 1953. Some Classes of Recursive Functions. *Rozprawy Matematyczne* 4 (1953), 1–45.
- [13] P. Jančar. 2015. On Reachability-Related Games on Vector Addition Systems with States. In *Reachability Problems (LNCS)*, Vol. 9328. Springer, 50–62.
- [14] M. Jurdziński, R. Lazić, and S. Schmitz. 2015. Fixed-Dimensional Energy Games are in Pseudo-Polynomial Time. In *Proceedings of ICALP 2015 (LNCS)*, Vol. 9135. Springer, 260–272.
- [15] S.R. Kosaraju and G. Sullivan. 1988. Detecting Cycles in Dynamic Graphs in Polynomial Time. In *Proceedings of STOC'88*. ACM Press, 398–406.
- [16] J. Leroux. 2018. Polynomial Vector Addition Systems With States. In *Proceedings of ICALP 2018 (LIPIcs)*, Vol. 107. Schloss Dagstuhl, 134:1–134:13.
- [17] E.W. Mayr and A.R. Meyer. 1981. The Complexity of the Finite Containment Problem for Petri Nets. *JACM* 28, 3 (1981), 561–576.
- [18] C. Rackoff. 1978. The Covering and Boundedness Problems for Vector Addition Systems. *TCS* 6 (1978), 223–231.
- [19] S. Schmitz. 2019. The Parametric Complexity of Lossy Counter Machines. In *Proceedings of ICALP 2018 (LIPIcs)*, Vol. 132. Schloss Dagstuhl, 129:1–129:15.
- [20] F. Zuleger. 2020. The Polynomial Complexity of Vector Addition Systems with States. In *Proceedings of FoSSaCS 2020 (LNCS)*, Vol. 12077. Springer, 622–641.