Separability in Büchi VASS and Singly Non-Linear Systems of Inequalities

Pascal Baumann
Max Planck Institute for Software Systems (MPI-SWS), Kaiserslautern, Germany

Eren Keskin
TU Braunschweig, Germany

Roland Meyer
TU Braunschweig, Germany

Georg Zetzsche
Max Planck Institute for Software Systems (MPI-SWS), Kaiserslautern, Germany

Abstract

The ω-regular separability problem for Büchi VASS coverability languages has recently been shown to be decidable, but with an EXPSPACE lower and a non-primitive recursive upper bound – the exact complexity remained open. We close this gap and show that the problem is EXPSPACE-complete. A careful analysis of our complexity bounds additionally yields a PSPACE procedure in the case of fixed dimension ≥ 1, which matches a pre-established lower bound of PSPACE for one dimensional Büchi VASS. Our algorithm is a non-deterministic search for a witness whose size, as we show, can be suitably bounded. Part of the procedure is to decide the existence of runs in VASS that satisfy certain non-linear properties. Therefore, a key technical ingredient is to analyze a class of systems of inequalities where one variable may occur in non-linear (polynomial) expressions.

These so-called singly non-linear systems (SNLS) take the form $\boldsymbol{A}(x) \cdot \boldsymbol{y} \geq \boldsymbol{b}(x)$, where $\boldsymbol{A}(x)$ and $\boldsymbol{b}(x)$ are a matrix resp. a vector whose entries are polynomials in x, and $\boldsymbol{y}$ ranges over vectors in the rationals. Our main contribution on SNLS is an exponential upper bound on the size of rational solutions to singly non-linear systems. The proof consists of three steps. First, we give a tailor-made quantifier elimination to characterize all real solutions to x. Second, using the root separation theorem about the distance of real roots of polynomials, we show that if a rational solution exists, then there is one with at most polynomially many bits. Third, we insert the solution for x into the SNLS, making it linear and allowing us to invoke standard solution bounds from convex geometry.

Finally, we combine the results about SNLS with several techniques from the area of VASS to devise an EXPSPACE decision procedure for ω-regular separability of Büchi VASS.

2012 ACM Subject Classification Theory of computation → Logic; Theory of computation → Formal languages and automata theory

Keywords and phrases Vector addition systems, infinite words, separability, inequalities, quantifier elimination, rational, polynomials

Digital Object Identifier 10.4230/LIPIcs.ICALP.2024.126

Category Track B: Automata, Logic, Semantics, and Theory of Programming

Related Version *Full Version*: https://arxiv.org/abs/2406.01008

Funding Funded by the European Union (ERC, FINABIS, 101077902). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

51st International Colloquium on Automata, Languages, and Programming (ICALP 2024).
Editors: Karl Bringmann, Martin Grohe, Gabriele Puppis, and Ola Svensson;
Article No. 126; pp. 126:1–126:19
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Vector addition systems with states (VASS) are one of the most popular and well-studied models of concurrent systems. A d-dimensional VASS consists of finitely many control states and d counters. Transitions between control states can increment or decrement the d counters, but importantly, one can only take a transition if the new counter values remain non-negative.

Separability problems. In recent years, a strong focus of the research on VASS was on *separability problems* [2, 5, 6, 8–11, 15–17, 31]. Here, we label the transitions of the input VASS $\mathcal{V}_1$, $\mathcal{V}_2$ by letters, which gives rise to languages L_1 and L_2. Then, we ask whether there exists a language S, from some class $\mathcal{S}$ of allowed separators, such that $L_1 \subseteq S$ and $L_2 \cap S = \emptyset$. Here, $\mathcal{S}$ is typically the class of regular languages.

An important motivation for studying separability problems is that separators can be viewed as certificates for disjointness, and thus the non-existence of a run in the product of $\mathcal{V}_1$ and $\mathcal{V}_2$. Such certificates are crucial for understanding safety verification for infinite-state systems, where the difficult part is to prove the non-existence of a run (the existence of a run is usually easy to show). In particular, certificates for non-existence are often the ingredient that is conceptually hardest to come by. For example, in the case of reachability in VASS, the KLM decomposition [18, 19, 22, 26] and Leroux's Presburger-definable inductive invariants [21] can be viewed as such certificates. Regular separators could play a similar role in alternative approaches to reachability.

In addition to understanding certificates, the recent attention on separability has led to other applications. For example, work on separability by bounded languages has led to a general framework to address unboundedness problems for VASS [8]. Moreover, separability results were used in an algorithm for deciding inclusion between unambiguous VASS [7].

With the recent contribution by Keskin and Meyer [16] (together with earlier decidability results for subclasses and variants [2, 5, 6, 8–11]), proving regular separability decidable for (finite-word) VASS, the *decidability* status of regular separability has largely been settled. However, concerning *complexity*, regular separability is far from understood, with few results: So far, the only exact complexity results are PSPACE-completeness for (succintly represented) one-dimensional VASS [9], EXPSPACE-completeness for VASS coverability languages [10], and Ackermann-completeness for VASS reachability languages [16].

Büchi VASS. A particularly challenging problem is (ω-)regular separability in Büchi VASS [2]. In a Büchi VASS $\mathcal{V}$, the language $L(\mathcal{V})$ consists of *infinite words* induced by runs that visit some final state infinitely often. As demonstrated by Baumann, Meyer, and Zetzsche [2], Büchi VASS behave quite differently in terms of regular separability from their finite-word counterpart, coverability languages of VASS [10]. Nevertheless, Baumann, Meyer, and Zetzsche proved decidability of regular separability for Büchi VASS [2]. However, the complexity remained open: Their algorithm requires at least Ackermannian time (because it constructs Karp-Miller graphs), and the only known lower bound is EXPSPACE.

Challenge: Non-linear constraints. Improving the complexity established in [2] is challenging due to the characterization of inseparability there: Inseparability is equivalent to the existence of a constellation of runs, called an *inseparability flower*, that must satisfy a *non-linear constraint*, meaning a constraint that is not expressible in linear arithmetic (i.e. first-order logic of $(\mathbb{Z}; +, <, 0, 1)$ or $(\mathbb{Q}; +, <, 0, 1)$). Essentially, such a flower is a triple

(α, β, γ) of cyclic runs such that (among other linear inequalities) the counter effect of the combined run $\alpha\beta\gamma$ is a scalar multiple of the counter effect of just α. In other words, we are looking for runs with effects $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{Z}^n$ such that

$$\exists x \in \mathbb{Q}\colon \boldsymbol{v} = x \cdot \boldsymbol{u} \ . \tag{1}$$

Detecting runs with such constraints is difficult: There are powerful generic EXPSPACE algorithms for detecting runs that satisfy unboundedness conditions [12], linear constraints [1], or variants of computation tree logic (CTL) [4]. However, condition (1) falls in neither of those categories.

In fact, we are not aware of any algorithmic approach to solving systems of linear inequalities with constraints of type (1) (let alone inside algorithms for VASS). There is a result by Gurari & Ibarra [13] showing that integral feasibility of systems of equalities $\boldsymbol{A} \cdot \boldsymbol{y} = \boldsymbol{b}(x)$ can be decided in NP, where $\boldsymbol{b}(x)$ is a vector containing in each component a quotient of polynomials in x. However, these do not seem to capture (1): By moving the denominators from $\boldsymbol{b}(x)$ to the left-hand side, one obtains equations where every variable from $\boldsymbol{y}$ is multiplied with the same polynomial over x. However, for (1), we need to multiply a *subset* of the linear variables (namely, those in $\boldsymbol{u}$) with a polynomial (namely, x). The same is true for the logic of *almost linear arithmetic* due to Weispfenning [32], whose existential fragment is also solvable in NP. Here, the definable sets are finite unions of solution sets of Gurari & Ibarra.

Furthermore, it is not even clear how to detect inseparability flowers by invoking reachability in VASS (even though this would only yield an Ackermann upper bound): To some extent, algorithms for reachability permit non-linear constraints – for example, using standard tricks, it is decidable whether one can reach a configuration with counter values (m, n) such that $n \leq 2^m$. However, the condition in (1) does not even seem to be captured by such methods.

Contribution. Our main result is that regular separability in Büchi VASS is EXPSPACE-complete, and PSPACE-complete in fixed dimension ≥ 1. The key technical ingredient is a method that we expect to be of independent interest: We develop a procedure for solving systems of linear inequalities with a single non-linear variable, which we call *singly non-linear systems of inequalities* (SNLS). We use our results about SNLS to show that if an inseparability witness exists, then there is one where all runs have at most doubly exponential length, yielding an EXPSPACE procedure. In fixed dimension, we obtain singly exponential bounds, leading to a PSPACE procedure.

Step I: Singly non-linear systems of inequalities. Intuitively, a singly non-linear system of inequalities (SNLS) is a system of inequalities that is linear in all but one variable. This means, there is one variable x that may appear in arbitrary polynomials, but all others can only occur linearly. More precisely, an SNLS is a system of inequalities of the form

$$\boldsymbol{A}(x) \cdot \boldsymbol{y} \geq \boldsymbol{b}(x), \quad \boldsymbol{y} \geq \boldsymbol{0} \ , \tag{2}$$

where $\boldsymbol{A}(x) \in \mathbb{Z}[x]^{m \times n}$ is an $m \times n$ matrix over the ring $\mathbb{Z}[x]$ of integer polynomials in x, $\boldsymbol{b}(x) \in \mathbb{Z}[x]^m$ is a vector of polynomials from $\mathbb{Z}[x]$, and $\boldsymbol{y}$ ranges over $\mathbb{Q}^n$. Notice that here indeed, x can be freely multiplied with itself and other variables, whereas the expression on the left-hand side must be linear in each component of $\boldsymbol{y}$.

Our main result about SNLS is that if a system as in (2) has a solution $(x, \boldsymbol{y}) \in \mathbb{Q} \times \mathbb{Q}^n$, then it has a solution where all numbers (numerators and denominators) are bounded exponentially in the description size of $\boldsymbol{A}(x)$ and $\boldsymbol{b}(x)$, even if numbers in the description are encoded in binary. This implies in particular that feasibility of SNLS is in NP.

In the proof, we first show that the set of all $x \in \mathbb{Q}$ for which there is a solution $(x, \boldsymbol{y})$ can be described by a Boolean combination Φ of polynomial constraints of the form $p(x) \geq 0$, for polynomials $p \in \mathbb{Z}[x]$. This amounts to a quantifier elimination procedure for a class of first-order formulas in the ordered field $(\mathbb{Q}; +, \cdot, <, 0, 1)$. This is perhaps surprising, since this structure does not admit quantifier elimination in general [24, Theorem 2].

Let us give a geometric explanation how we arrive at the constraints $\Phi(x)$: For each choice of x, the SNLS $\boldsymbol{A}(x) \cdot \boldsymbol{y} \geq \boldsymbol{b}(x),\ \boldsymbol{y} \geq \boldsymbol{0}$ defines a polyhedron. It is a standard fact in convex geometry that such a polyhedron has a point on a minimal face, and moreover this point can be expressed as the inverse of a submatrix of $\boldsymbol{A}(x)$ multiplied with $\boldsymbol{b}(x)$. This expression can then be plugged back into $\boldsymbol{A}(x) \cdot \boldsymbol{y} \geq \boldsymbol{b}(x)$ to obtain a set of polynomial constraints on x, subject to a particular determinant being non-zero. The latter non-zero condition can as well be expressed as a polynomial constraint.

We then show that Φ has a small solution: In one case, a rational root of one of the polynomials p is a solution – these can be bounded by the *Rational Root Theorem*. The other case is that the solution x lies strictly between two roots $r_1 < r_2$ of participating polynomials. But then one can observe that any rational number between those roots is a solution (if no other root lies between r_1 and r_2). Using the *Root Separation Theorem* (specifically, Rump's Bound [27]), which lower-bounds the size of such intervals (r_1, r_2), we can then conclude that such an interval must contain a rational number with small numerator and denominator.

Once we exhibit a small x, we can plug it into $\boldsymbol{A}(x) \cdot \boldsymbol{y} \geq \boldsymbol{b}(x)$ to obtain a system of linear inequalities. Then we use standard bounds to obtain a small (i.e. exponential) solution $\boldsymbol{y} \in \mathbb{Q}^n$. It should be noted that while our result about SNLS concerns rational solutions, we apply it in the case where $\boldsymbol{b}(x) \geq \boldsymbol{0}$, which means a rational solution can be turned into an integral solution by multiplying a common denominator.

Step II: Rackoff-like bounds. After establishing the solution bound for SNLS, we use this result in the context of Büchi VASS to show the existence of inseparability witnesses that are small, i.e. consist of runs that are at most doubly exponential in length. Here, we use an adaptation of the Rackoff technique [28] similar to the proofs of Habermehl [14] and Atig & Habermehl [1]. In [1], it is shown that runs satisfying (restricted) linear inequalities can be detected in EXPSPACE. For this, they use a Rackoff-style induction to bound the length of such runs. We devise a similar Rackoff-style induction to work with SNLS instead of ordinary linear inequalities. Different compared to the earlier works is the fact that our witnesses contain ω-counters, which may change when invoking the induction hypothesis. Moreover, we need to use a result of Demri on selective unboundedness [12, Theorem 4.6] (in EXPSPACE in the general case and PSPACE in fixed dimension) to check the coverability of our witnesses.

2 Preliminaries

Büchi VASS. A *Büchi vector addition system with states (Büchi VASS)* of dimension $d \in \mathbb{N}$ over an alphabet Σ is a tuple $\mathcal{V} = (Q, q_0, \Sigma, T, F)$. It consists of a finite set of states Q, an initial state $q_0 \in Q$, a set of final states $F \subseteq Q$, and a finite set of transitions $T \subseteq Q \times \Sigma^* \times \mathbb{Z}^d \times Q$. The size of the Büchi VASS is $|\mathcal{V}| := |Q| + |F| + \sum_{(q,w,\delta,q') \in T} \big(|w| + \|\delta\|_2\big)$. By $\|\delta\|_2$, we mean the size of the binary encoding of δ. Since we only consider Büchi VASS in this paper, we often simply call them VASS. If $d = 0$, we call $\mathcal{V}$ a *Büchi automaton*.

The semantics of the Büchi VASS is defined over its *configurations*, which are elements of $Q \times \mathbb{N}^d$. The *initial configuration* of $\mathcal{V}$ is $(q_0, \boldsymbol{0})$. We lift the transitions of the Büchi VASS to a relation over configurations $\to \subseteq Q \times \mathbb{N}^d \times \Sigma^* \times Q \times \mathbb{N}^d$ as follows: $(q, \boldsymbol{m}) \xrightarrow{w} (q', \boldsymbol{m}')$ if there is $(q, w, \delta, q') \in T$ such that $\boldsymbol{m}' = \boldsymbol{m} + \delta$. A *run* of the Büchi VASS is a (possibly infinite) sequence of configurations of the form $\sigma = (p_0, \boldsymbol{m}_0) \xrightarrow{w_1} (p_1, \boldsymbol{m}_1) \xrightarrow{w_2} \cdots$.

A run σ is *accepting* if it starts from the initial configuration and visits final states infinitely often, meaning there are infinitely many configurations $(q, \boldsymbol{m})$ in σ with $q \in F$. The run is said to be *labeled* by the word $w = w_0 w_1 \cdots$ in Σ^ω. The *language* $L(\mathcal{V})$ of the Büchi VASS consists of all infinite words that label an accepting run.

An infinite-word language $L \subseteq \Sigma^\omega$ is called *regular* if it is accepted by a Büchi automaton. As we only consider infinite-word languages, we just call them languages.

Arithmetic. Our approach to regular separability in Büchi VASS rests on a result about solutions to singly non-linear systems of inequalities. This also requires some terminology.

We define the integers, rationals, polynomials, and matrices together with the operations we need to perform on them. Let $a \in \mathbb{Z}$ be an integer. Its size $\|a\|_2 = |bin(a)|$ is the length of its binary encoding. We also use $\|a\|_1$ to denote the size of the unary encoding. This is the absolute value plus an extra bit for the sign. A polynomial with integer coefficients $p \in \mathbb{Z}[x]$ is a sum $\sum_{i=0}^{k} a_i x^i$ with $a_0, \ldots, a_k \in \mathbb{Z}$ and $a_k \neq 0$ if $k > 0$. We define $\|p\|_1 = \sum_{i=0}^{k} \|a_i\|_1$ and similar for $\|p\|_2$. The *degree* of the polynomial is $\deg(p) = k$, its *maximal coefficient* is $\mathrm{maxc}(p) = \max_{i \in [0,k]} \|a_i\|_1$. Note that $\|p\|_1 \leq (\deg(p) + 1) \cdot \mathrm{maxc}(p)$. A real number $r \in \mathbb{R}$ with $p(r) = 0$ is called a *root* of the polynomial. Let S be a set with a size function $\| - \|$ defined on it. We consider matrices $\boldsymbol{A} \in S^{m \times n}$ over S, and define their size $\|\boldsymbol{A}\|$ by summing up the sizes of the entries. We use $row(\boldsymbol{A}) = m$ and $col(\boldsymbol{A}) = n$. When $S = \mathbb{Z}[x]$, we also use $\deg(\boldsymbol{A})$ for the highest degree of a polynomial in $\boldsymbol{A}$ and $\mathrm{maxc}(\boldsymbol{A})$ for the maximal coefficient of a polynomial in $\boldsymbol{A}$. Pairs $(s_1, s_2) \in S \times S$ form a special case with size $\|(s_1, s_2)\| = \|s_1\| + \|s_2\|$. In particular, a rational number $t \in \mathbb{Q}$ is a pair $t = \frac{a}{b}$ of integers $a, b \in \mathbb{Z}$ with $\|t\|_1 = \|a\|_1 + \|b\|_1$, and similar for $\|t\|_2$.

We perform addition $a + b$ and multiplication $a \cdot b$ among integers, rationals, polynomials, and matrices. These operations can be executed in time polynomial in $\|a\|_2 + \|b\|_2$. The same holds for the comparison $a \geq b$ among integers and rationals. We also add, multiply, and compare integers and rationals with $-\infty$ and ∞. The definitions are as expected.

3 Main results

A language $R \subseteq \Sigma^\omega$ is said to *separate* languages $L_1, L_2 \subseteq \Sigma^\omega$, if $L_1 \subseteq R$ and $R \cap L_2 = \emptyset$. We call L_1 and L_2 *regular separable*, denoted by $L_1 \mid L_2$, if there is a separator R that is a regular language. The problem we address is the *regular separability problem* for Büchi VASS:

Given Two VASS $\mathcal{V}_1$ and $\mathcal{V}_2$ over some alphabet Σ.

Question Does $L(\mathcal{V}_1) \mid L(\mathcal{V}_2)$ hold?

We also consider the variants of this problem where the inputs are of fixed dimension: For a fixed number $d \in \mathbb{N} \setminus \{0\}$, the *$d$-dimensional regular separability problem* for Büchi VASS is the same problem as above, except that the input VASS $\mathcal{V}_1$ and $\mathcal{V}_2$ are restricted to be of dimension at most d. Our first main result is the following:

▶ **Theorem 3.1.** *The regular separability problem for Büchi VASS is* EXPSPACE*-complete. Moreover, the d-dimensional regular separability problem is* PSPACE*-complete for all $d \geq 1$.*

As mentioned above, the proof is based on a small model property for what we call singly non-linear systems of inequalities. We expect this result to be of independent interest. Formally, a *singly non-linear system (SNLS)* is a system of inequalities of the form

$$\boldsymbol{A}(x) \cdot \boldsymbol{y} \geq \boldsymbol{b}(x) \wedge \boldsymbol{y} \geq \boldsymbol{0} \ .$$

Here, $\boldsymbol{A}(x) \in \mathbb{Z}[x]^{m \times n}$ is an $m \times n$ matrix over the set of polynomials with integer coefficients in variable x, and $\boldsymbol{b} \in \mathbb{Z}[x]^m$ is a vector of polynomials. We also write an SNLS as $\mathcal{S} = (\boldsymbol{A}(x), \boldsymbol{b}(x))$, or $\mathcal{S}(x, \boldsymbol{y})$ to emphasize the variables. A *solution* to $\mathcal{S}$ is a pair $(t, \boldsymbol{s}) \in \mathbb{Q} \times \mathbb{Q}^n$ that satisfies $\boldsymbol{A}(t) \cdot \boldsymbol{s} \geq \boldsymbol{b}(t) \wedge \boldsymbol{s} \geq \boldsymbol{0}$. If a solution exists, we call the system *feasible*.

Our second main result is a bound on the size of least solutions.

▶ **Theorem 3.2.** *If the SNLS $\mathcal{S}$ is feasible, then it has a solution $(t, \boldsymbol{s})$, where all components of $\boldsymbol{s}$ have the same denominator, with $\|t\|_1, \|\boldsymbol{s}\|_1 \in (col(\mathcal{S}) \cdot \deg(\mathcal{S}) \cdot \mathrm{maxc}(\mathcal{S}))^{\mathcal{O}(\deg(\mathcal{S})^2 \cdot row(\mathcal{S})^4)}$.*

Theorem 3.2 implies that a feasible system $\mathcal{S}$ always has a solution of size at most singly exponential in $\|\mathcal{S}\|_1$. This gives an upper bound on the complexity of feasibility.

▶ **Corollary 3.3.** *Feasibility of SNLS is in* NP.

The reader may have noted that SNLS are more general than the non-linear systems we are confronted with when checking separability. There are at least two arguments in support of the generalization. First, non-linearity is not well-understood, and we believe a class of systems that admits an efficient algorithm for checking feasibility will find its applications. Second, the generalization only adds little complexity to the proof or, phrased differently, the special case already needs most considerations.

Organization. The remainder of the paper is organized as follows. In Section 4, we prove Theorem 3.2 and in Section 5, we show Theorem 3.1.

4 Singly Non-Linear Systems

In this section, we prove Theorem 3.2.

Some notation. By $\boldsymbol{A}(t)$ or $eval(\boldsymbol{A}(x), t)$ we mean the matrix with rational entries that results from $\boldsymbol{A}(x)$ by evaluating all polynomials at t. Let $\boldsymbol{A} \in R^{n \times n}$ be a square matrix over some ring R. In our exposition, we will consider matrices over the rings $\mathbb{Z}$, $\mathbb{Q}$, and $\mathbb{Z}[x]$. We write $\det(\boldsymbol{A})$ for the determinant, and recall that if R is a field (such as $\mathbb{Q}$), then $\boldsymbol{A}$ is invertible if and only if $\det(\boldsymbol{A}) \neq 0$. The *adjugate* (also called *classical adjoint*) of $\boldsymbol{A}$ is the matrix $\mathrm{adj}(\boldsymbol{A}) \in R^{n \times n}$ with $\mathrm{adj}(\boldsymbol{A})[j, i] = (-1)^{i+j} \det(\boldsymbol{A}_{ij})$, where $\boldsymbol{A}_{ij}$ is the matrix obtained from $\boldsymbol{A}$ by removing the i-th row and the j-th column. It is well-known that then $\boldsymbol{A} \cdot \mathrm{adj}(\boldsymbol{A}) = \det(\boldsymbol{A}) \cdot \boldsymbol{I}$, where $\boldsymbol{I}$ is the identity matrix in dimension n. In particular, if $\boldsymbol{A}$ is invertible, its inverse can be computed as $\boldsymbol{A}^{-1} = \frac{\mathrm{adj}(\boldsymbol{A})}{\det(\boldsymbol{A})}$ [20, Chapter XIII, Prop. 4.16].

In upper bound arguments, we will use the well-known Leibniz formula for determinants, which says $\det(\boldsymbol{A}) = \sum_{\sigma \in S_n} \mathrm{sgn}(\sigma) \cdot \prod_{i=1}^{n} \boldsymbol{A}[i, \sigma(i)]$ [20, Chapter XIII, Prop. 4.6]. Here, S_n is the set of all permutations of $[1, n]$ and $\mathrm{sgn}(\sigma) \in \{-1, 1\}$ is the sign of $\sigma \in S_n$.

Bounding solutions. For the proof of Theorem 3.2, we proceed in two steps. We first show that if an SNLS $\mathcal{S}(x, \boldsymbol{y})$ is feasible, then we find a small rational t for x such that the system $\mathcal{S}(t, \boldsymbol{y})$ is feasible. This system is the result of evaluating all polynomials in $\mathcal{S}$ at t, and thus having only $\boldsymbol{y}$ as the variables.

▶ **Lemma 4.1.** *If the SNLS $\mathcal{S}(x, \boldsymbol{y})$ is feasible, then there is a number $t \in \mathbb{Q}$ with $\|t\|_1 \in (col(\mathcal{S}) \cdot \deg(\mathcal{S}) \cdot \mathrm{maxc}(\mathcal{S}))^{\mathcal{O}(\deg(\mathcal{S}) \cdot row(\mathcal{S})^3)}$ such that $\mathcal{S}(t, \boldsymbol{y})$ is feasible.*

Lemma 4.1 is non-trivial and will occupy almost this entire section. To explain our approach, note that the feasibility of $\mathcal{S}(x, \boldsymbol{y})$ is equivalent to the feasibility of $\exists \boldsymbol{y}.\mathcal{S}(x, \boldsymbol{y})$. Our first step is to eliminate the quantifier and determine a new formula $\Phi(x)$ in which $\boldsymbol{y}$ no longer occurs and that is equivalent to the previous one over the rationals, $\exists \boldsymbol{y}.\mathcal{S}(x, \boldsymbol{y}) \models|_{\mathbb{Q}} \Phi(x)$. The equivalence says that for every $t \in \mathbb{Q}$, we have $t \models \exists \boldsymbol{y}.\mathcal{S}(x, \boldsymbol{y})$ if and only if $t \models \Phi(x)$.

The second step for Lemma 4.1 is to show that if the new formula holds, then we find a small solution for t. To this end, we will apply the Root Separation Theorem, which provides a lower bound on the distance between distinct real roots of polynomials. After establishing Lemma 4.1, we obtain Theorem 3.2 (at the end of this section) by taking the t provided by Lemma 4.1, and pair it with the $\boldsymbol{s} \in \mathbb{Q}^n$, which must exist according to the quantifier elimination done in the first step of Lemma 4.1.

4.1 Quantifier Elimination

We show how to remove the quantifier from $\exists \boldsymbol{y}.\mathcal{S}(x, \boldsymbol{y})$ with a tailor-made quantifier elimination algorithm. The fact that quantifier elimination is possible in this setting came as a surprise to us, given the non-linear nature and the setting of rationals. For example, the real closed field $(\mathbb{R}; +, \cdot, <, 0, 1)$ admits quantifier elimination by a classical result of Tarski [25, Theorem 3.3.15], but this is not true for the ordered field $(\mathbb{Q}; +, \cdot, <, 0, 1)$ of rationals [24, Theorem 2] (see [25, p. 71–72] for a simple example). This means, there are first-order formulas over $(\mathbb{Q}; +, \cdot, <, 0, 1)$ that have no quantifier-free equivalent. However, we show that if we existentially quantify the linear variables in the formulas induced by SNLS, then those quantifiers can be eliminated.

The precise formulation of the result needs some notation. A *lower bound constraint* has the form $p(x) \geq 0$ or $p(x) > 0$ with $p \in \mathbb{Z}[x]$ a polynomial with integer coefficients. The formula $\Phi(x)$ that we want to obtain takes the form $\bigvee_{i \in I} \bigwedge_{j \in J_i} \Phi_{i,j}(x)$, where the formulae $\Phi_{i,j}(x)$ are lower bound constraints. We call it a DNFLB, short for *disjunctive normal form with lower bound constraints as the literals*. We may also omit x and write Φ. We use $\deg(\Phi)$ and $\operatorname{maxc}(\Phi)$ for the maximal degree resp. coefficient of a polynomial in Φ.

▶ **Theorem 4.2.** *For every SNLS $\mathcal{S}(x, \boldsymbol{y})$, there is a DNFLB $\Phi(x)$ with $\exists \boldsymbol{y}.\mathcal{S}(x, \boldsymbol{y}) \models|_{\mathbb{Q}} \Phi(x)$, $\deg(\Phi) \in \mathcal{O}(row(\mathcal{S}) \cdot \deg(\mathcal{S}))$, and $\operatorname{maxc}(\Phi) \in (col(\mathcal{S}) \cdot \deg(\mathcal{S}) \cdot \operatorname{maxc}(\mathcal{S}))^{\mathcal{O}(row(\mathcal{S})^2)}$.*

Since our intention is to bound the solutions t to variable x, the given estimations on the degree and the maximal coefficient suffice for us. The proof actually gives an algorithm to compute Φ which runs in time exponential in the dimension of $\mathcal{S}$, but we do not need the effectiveness here. In the proof of Theorem 4.2, we will use a standard fact about polyhedra:

▶ **Lemma 4.3.** *Suppose $\boldsymbol{D} \in \mathbb{Q}^{m \times n}$ and $\boldsymbol{c} \in \mathbb{Q}^m$. If the system $\boldsymbol{D} \cdot \boldsymbol{x} \geq \boldsymbol{c}$ has a solution in $\mathbb{Q}^n$, then there is a solution $\boldsymbol{s} \in \mathbb{Q}^n$ that also satisfies $\boldsymbol{D}' \cdot \boldsymbol{s} = \boldsymbol{c}'$, where $(\boldsymbol{D}', \boldsymbol{c}')$ is a subset of the rows of $(\boldsymbol{D}, \boldsymbol{c})$ such that $\operatorname{rank}(\boldsymbol{D}') = \operatorname{rank}(\boldsymbol{D})$.*

Proof. By well-known decomposition theorems about polyhedra, a polyhedron $P = \{\boldsymbol{s} \in \mathbb{Q}^n \mid \boldsymbol{D} \cdot \boldsymbol{s} \geq \boldsymbol{c}\}$ is non-empty if and only if it has a non-empty minimal face [30, Theorem 8.5]. Moreover, minimal faces can be characterized as exactly the sets of the form $\{\boldsymbol{s} \in \mathbb{Q}^n \mid \boldsymbol{D}' \cdot \boldsymbol{s} = \boldsymbol{c}'\}$, where $(\boldsymbol{D}', \boldsymbol{c}')$ is a subset of the rows of $(\boldsymbol{D}, \boldsymbol{c})$ such that $\boldsymbol{D}'$ has the same rank as $\boldsymbol{D}$ [30, Theorem 8.4]. ◀

We are ready to prove Theorem 4.2:

Proof of Theorem 4.2. Let $\mathcal{S}(x, \boldsymbol{y}) = \boldsymbol{A}(x) \cdot \boldsymbol{y} \geq \boldsymbol{b}(x) \wedge \boldsymbol{y} \geq \boldsymbol{0}$. To fix the dimension, let $\boldsymbol{A} \in \mathbb{Z}[x]^{m \times n}$. We can equivalently write the SNLS as $\mathcal{S}'(x, \boldsymbol{y}) = \boldsymbol{D}(x) \cdot \boldsymbol{y} \geq \boldsymbol{c}(x)$ with

$$\boldsymbol{D}(x) \;=\; \begin{pmatrix} \boldsymbol{A}(x) \\ \boldsymbol{I}_n \end{pmatrix} \in \mathbb{Z}[x]^{(m+n)\times n} \qquad \boldsymbol{c}(x) \;=\; \begin{pmatrix} \boldsymbol{b}(x) \\ \boldsymbol{0}_n \end{pmatrix} \in \mathbb{Z}[x]^{m+n} \;,$$

i.e. we glue the $n \times n$ identity matrix $\boldsymbol{I}_n$ to the bottom of $\boldsymbol{A}(x)$ and extend $\boldsymbol{b}(x)$ by n zeros.

Assume $\mathcal{S}'$ is feasible and the solution for x is $t \in \mathbb{Q}$. By Lemma 4.3 and since $\boldsymbol{D}(t)$ has rank n, we can select a subset of n rows of $\boldsymbol{D}(t)$ and of $\boldsymbol{c}(t)$ such that the smaller system has a solution, even with equality. More formally, for any subset $R \subseteq [1, m+n]$, denote by $\boldsymbol{D}_R(t)$ (resp. $\boldsymbol{c}_R(t)$) the matrix (resp. vector) obtained by selecting only the rows in R from $\boldsymbol{D}(t)$ (resp. $\boldsymbol{c}(t)$). Then Lemma 4.3 tells us that there is an $\boldsymbol{s} \in \mathbb{Q}^n$ with $\boldsymbol{D}_R(t) \cdot \boldsymbol{s} = \boldsymbol{c}_R(t)$, where $\boldsymbol{D}_R(t)$ has rank n. In particular, $\boldsymbol{D}_R(t)$ is invertible and thus $\det(\boldsymbol{D}_R(t)) \neq 0$. The fact that $\boldsymbol{D}_R(t)$ is invertible allows us to write $\boldsymbol{s} = \boldsymbol{D}_R(t)^{-1} \cdot \boldsymbol{c}_R(t)$, which will be key for our quantifier elimination. The argumentation shows that for every $t \in \mathbb{Q}$, $\exists \boldsymbol{y}.\mathcal{S}'(t, \boldsymbol{y})$ is equivalent to the condition

$$\bigvee_{\substack{R \subseteq [1, m+n] \\ |R| = n}} \det(\boldsymbol{D}_R(t)) \neq 0 \;\wedge\; \boldsymbol{D}(t) \cdot \boldsymbol{D}_R(t)^{-1} \cdot \boldsymbol{c}_R(t) \geq \boldsymbol{c}(t) \;. \tag{3}$$

Here, of course, we only know that $\boldsymbol{D}_R(t)^{-1}$ exists when $\det(\boldsymbol{D}_R(t)) \neq 0$. To express (3) using polynomials, we employ the identity $\boldsymbol{D}_R(t)^{-1} = \frac{\operatorname{adj}(\boldsymbol{D}_R(t))}{\det(\boldsymbol{D}_R(t))}$ whenever $\boldsymbol{D}_R(t)$ is invertible (equivalently, whenever $\det(\boldsymbol{D}_R(t)) \neq 0$). Thus, the set of all t with $\exists \boldsymbol{y}\colon \mathcal{S}'(t, \boldsymbol{y})$ can be defined by the following DNFLB Φ:

$$\begin{aligned} \bigvee_{\substack{R \subseteq [1, m+n] \\ |R| = n}} & \Big(\det(\boldsymbol{D}_R(x)) > 0 \;\wedge\; \boldsymbol{D}(x) \cdot \operatorname{adj}(\boldsymbol{D}_R(x)) \cdot \boldsymbol{c}_R(x) \geq \det(\boldsymbol{D}_R(x)) \cdot \boldsymbol{c}(x) \Big) \\ & \vee \Big(\det(\boldsymbol{D}_R(x)) < 0 \;\wedge\; \boldsymbol{D}(x) \cdot \operatorname{adj}(\boldsymbol{D}_R(x)) \cdot \boldsymbol{c}_R(x) \leq \det(\boldsymbol{D}_R(x)) \cdot \boldsymbol{c}(x) \Big) \;, \end{aligned} \tag{4}$$

where indeed all conditions are of the form $p(x) \geq 0$ or $p(x) > 0$ for some polynomials p. Note that here, we distinguish the cases $\det(\boldsymbol{D}_R(x)) < 0$ and $\det(\boldsymbol{D}_R(x)) > 0$ because moving a negative $\det(\boldsymbol{D}_R(x))$ to the other side of the inequality changes $\geq$ to $\leq$. Moreover, note that (in contrast to (3)) in the formulation (4), all terms are well-defined, independenly of whether the current choice of R makes $\boldsymbol{D}_R(x)$ invertible or not. To be precise, we obtain the DNFLB by subtracting the right-hand sides of the inequalities from the left-hand sides and multiplying the result by -1 to invert the inequality where necessary. The above form will suffice to give an estimate on the maximal degree and the maximal coefficient.

It is now clear that the coefficients (resp. degrees) appearing in Φ are exponential (resp. polynomial) in the bitsize of $\mathcal{S}$. The precise bounds promised in the Theorem are straightforward to deduce from standard bounds on determinants, see the full version for details. ◀

4.2 Root Separation

To show Lemma 4.1, it remains to be shown that any feasible DNFLB $\Phi(x)$ has a solution that is exponentially bounded. The key observation is that if r and r' are adjacent roots of a polynomial $p(x) \in \mathbb{Z}[x]$ and a constraint $p(x) \geq 0$ or $p(x) > 0$ is satisfied for some t for x with $r < t < r'$, then any number t' in the open interval (r, r') will also satisfy the constraint: The polynomial does not change its sign between r and r'. Thus, we can think of $\mathbb{R}$ as being split into (i) roots of p and (ii) intervals between roots of p (and the infinite intervals below

the smallest and above the largest root). Then whether $t \in \mathbb{Q}$ satisfies $p(x) \geq 0$ or $p(x) > 0$ only depends on which of those parts of $\mathbb{R}$ the number t belongs to. This remains true if we refine this decomposition of $\mathbb{R}$ according to *all* polynomials occurring in Φ.

In order to construct rational numbers with small numerator and denominator in intervals (r, r'), we will rely on a Root Separation Theorem, saying that polynomial roots are not too close. More specifically, we use Rump's Bound [27, Theorem 8.5.5]:

▶ **Theorem 4.4** (Rump's Bound [27, Theorem 8.5.5]). *Suppose $r, r' \in \mathbb{R}$ are distinct roots of a polynomial $p(x) \in \mathbb{Z}[x]$ with degree $d \in \mathbb{N}$. Then $|r - r'| > (d^{d+1}(1 + \|p(x)\|_1)^{2d})^{-1}$.*

We will also use an elementary fact about rational roots of integral polynomials. It is known as the Rational Root Theorem or Integral Root Test [20, Chapter IV, Prop. 3.3]:

▶ **Lemma 4.5** (Rational Root Theorem [20, Chapter IV, Prop. 3.3]). *Let $p(x) = c_n x^n + \cdots + c_0 \in \mathbb{Z}[x]$ be a polynomial. If $r = a/b$ is a root of p with a, b co-prime, then a divides c_0 and b divides c_n. In particular, $|a|, |b| \leq \mathrm{maxc}(p)$.*

Finally, we need a standard bound on all real roots of a polynomial [27, Corollary 8.3.2]. This is known as Cauchy's bound.

▶ **Lemma 4.6** (Cauchy's Bound [27, Corollary 8.3.2]). *If $r \in \mathbb{R}$ is a root of a polynomial $p \in \mathbb{Z}[x]$, then $|r| \leq 1 + \|p\|_1$.*

Let $r_1 < \cdots < r_k \in \mathbb{R}$ be all the real roots of polynomials occurring in Φ. Observe that if $t \in \mathbb{Q}$ satisfies $\Phi(x)$ and $t \in (r_i, r_{i+1})$, then any rational number in (r_i, r_{i+1}) must satisfy $\Phi(x)$, because none of the polynomials in Φ changes its sign between r_i and r_{i+1}. This allows us to bound a rational solution, by distinguishing the following cases:

1. Suppose $\Phi(x)$ is satisfied by some rational root r_i of p in Φ. Write $r_i = \frac{a}{b}$ with a, b co-prime. Then the Rational Root Theorem (Lemma 4.5) implies $|a|, |b| \leq \mathrm{maxc}(p)$.
2. Suppose $\Phi(x)$ has a rational solution in some interval (r_i, r_{i+1}). Since r_i and r_{i+1} are the roots of some polynomials p, q in $\Phi(x)$, as observed above, any rational number in (r_i, r_{i+1}) is also a solution to $\Phi(x)$. Note that r_i, r_{i+1} are roots of $p(x) \cdot q(x)$ and thus by Theorem 4.4, we have $|r_i - r_{i+1}| > \frac{1}{b}$ for some $b \in \mathbb{Z}$ that is exponentially bounded. Thus, there is an integer $a \in (br_i, br_{i+1})$. Note that then $\frac{a}{b}$ belongs to the interval (r_i, r_{i+1}) and thus satisfies $\Phi(x)$. Moreover, by the Cauchy Bound (Lemma 4.6), we also have an exponential bound $U \in \mathbb{R}$ on $|r_i|, |r_{i+1}|$ and thus on $|a| \leq |b|U$.
3. Suppose $\Phi(x)$ has a rational solution t outside of $[r_1, r_k]$. If $t > r_k$ then every rational number in $[r_k, \infty)$ is also a solution. Moreover, by Lemma 4.6, any rational number t' with $t' > 1 + \|p\|_1$ for every polynomial p occurring in Φ can be chosen, e.g. $t' = 2 + c$, where $c = \max\{\|p\|_1 \mid p \text{ polynomial in } \Phi\}$. On the other hand, if $t < r_1$, then $t' = -(2 + c)$ is a solution by an analogous argument.

This proves that any feasible $\Phi(x)$ has a rational solution that is exponentially bounded, which is what we will use in our application to Büchi VASS. The precise bounds of Lemma 4.1 are shown in the full version.

Proof sketch for Theorem 3.2. For showing Theorem 3.2, we can now use the fact that if $t \in \mathbb{Q}$ admits a solution $(t, \boldsymbol{s})$, then by our argument in the proof of Theorem 4.2, $\boldsymbol{s}^* := \frac{\mathrm{adj}(D_R(t))}{\det(D_R(t))} \cdot \boldsymbol{c}_R(t)$ is also a solution, for some subset $R \subseteq [1, m+n]$. This means that we can apply the bound on t and the bounds on $\mathrm{adj}(\boldsymbol{D}_R(x))$ and $\det(\boldsymbol{D}_R(x))$ established in the proof of Theorem 4.2 to bound the solution $\boldsymbol{s}^*$. We can ensure that all components of $\boldsymbol{c}_R(t)$ have the same denominator by increasing the bit size at most $\deg(\mathcal{S})$-fold. If we compute $\boldsymbol{s}^*$ starting from such a vector, we get an $\boldsymbol{s}^*$ where all components have the same denominator.

Since the entries in $\boldsymbol{D}_R$ all appear in $\mathcal{S}$ and it is well-known that the determinant has polynomial bit size in the bit size of a matrix, it follows that there exists a solution $(t, \boldsymbol{s})$ of polynomial bit size. The precise bounds promised in Theorem 3.2 are derived in the full version.

5 ω-Regular Separability

We use the results from Section 4 to prove Theorem 3.1. Note that the lower bounds in Theorem 3.1 easily follow from [2]. First, that paper already shows PSPACE-completeness of regular separability for one-dimensional Büchi VASS, which yields the PSPACE lower bound for fixed dimension ≥ 1. In fact, their argument also yields EXPSPACE-hardness in the general case: The full version [3, Appendix E.1] describes a simple reduction from intersection emptiness of one-dimensional VASS that accept by final state to regular separability of Büchi VASS, and the construction is the same in higher dimension. This yields EXPSPACE-hardness of regular separability of Büchi VASS, since intersection emptiness of VASS of arbitrary dimension that accept by final state is EXPSPACE-hard [23].

It remains to prove the upper bounds in Theorem 3.1. To adequately formulate our proofs, we need to introduce additional VASS-related concepts.

More on Büchi VASS. Let $\mathcal{V} = (Q, q_0, \Sigma, T, F)$ be a Büchi VASS. Consider a (possibly infinite) run $\sigma = (p_0, \boldsymbol{m}_0) \xrightarrow{w_1} (p_1, \boldsymbol{m}_1) \xrightarrow{w_2} \cdots$ of $\mathcal{V}$. The sequence of transitions in σ is called a *path* and has the form $\rho = (p_0, w_1, \delta_1, p_1)(p_1, w_2, \delta_1, p_2)\ldots$. If a path is finite and the source state of its first transition coincides with the target state of its last transition, then we call it a *loop*. Since a run is uniquely determined by the start configuration and its sequence of transitions, we also denote a run by $\sigma = (p_0, \boldsymbol{m}_0).\rho$. If σ is finite and $(p_\ell, \boldsymbol{m}_\ell)$ is its last configuration, then we sometimes write $\sigma = (p_0, \boldsymbol{m}_0).\rho.(p_\ell, \boldsymbol{m}_\ell)$ to emphasize this. The *effect* $\delta(\rho)$ of some finite path $\rho = (p_0, w_1, \delta_1, p_1)\ldots(p_{\ell-1}, w_\ell, \delta_\ell, p_\ell)$ is the sum of all induced counter changes, formally $\delta(\rho) = \sum_{1 \leq i \leq \ell} \delta_i$.

Recall that configurations of the Büchi VASS $\mathcal{V}$ are elements of the set $Q \times \mathbb{N}^d$. We call the second component in a configuration the *counter valuation* and refer to the i-th entry as the *value of counter* i. For a configuration cf and a set of counters $I \subseteq [1, d]$ we also use $cf[I]$ to denote the counter valuation of cf restricted to the counters in I. A configuration $(q, \boldsymbol{m})$ is *coverable* in $\mathcal{V}$ if there is a run starting in the initial configuration $(q_0, \boldsymbol{0})$ and reaching a configuration $(q, \boldsymbol{m}')$ with $\boldsymbol{m}' \geq \boldsymbol{m}$. Here, $\geq$ is defined component-wise.

Moreover we also consider a set of *extended configurations* $Q \times \mathbb{N}_\omega^d$, where $\mathbb{N}_\omega = \mathbb{N} \cup \{\omega\}$. Here ω is used to represent a counter value that has become unbounded. For an extended configuration $(q, \boldsymbol{m})$ we use $\omega(q, \boldsymbol{m}) \subseteq [1, d]$ to denote the set of counters valued ω in $\boldsymbol{m}$. Comparisons and arithmetic operations between integer values and ω behave as expected, treating ω as ∞. Formally, $\omega \geq \omega$, $\omega \geq z$, and $\omega + z = \omega$ for all $z \in \mathbb{Z}$. The *size* of an extended configuration is $|(q, \boldsymbol{m})| = \log_2 |Q| + \|\boldsymbol{m}\|_2 + d$, where the extra bit per counter encodes whether it has value ω or not. We also use the size of a unary encoding $\|(q, \boldsymbol{m})\|_1 = |Q| + \|\boldsymbol{m}\|_1 + d$.

The transition relation is also lifted to extended configurations in the expected manner. Formally, for $(q, \boldsymbol{m})$, $(q', \boldsymbol{m}') \in Q \times \mathbb{N}_\omega^d$ we have $(q, \boldsymbol{m}) \xrightarrow{w} (q', \boldsymbol{m}')$ if there is a transition $(q, w, \delta, q') \in T$ such that $\boldsymbol{m}' = \boldsymbol{m} + \delta$, where addition between elements of $\mathbb{N}_\omega^d$ and $\mathbb{Z}^d$ is defined component-wise. Furthermore, our definitions of runs, paths, loops, etc. carry over to *extended* versions over the set of extended configurations in a straightforward way. More precisely, an *extended run* is a sequence of extended configurations $cf_1 \xrightarrow{w_1} cf_2 \xrightarrow{w_2} \cdots$, an *extended path* is the underlying sequence of transitions of an extended run, and an *extended*

loop is a finite extended path starting and ending in the same state. To cover an extended configuration, intuitively, the ω-counters need to become unbounded, and the remaining counters need to be covered. Formally, an extended configuration $(q, \boldsymbol{m})$ is *coverable* in $\mathcal{V}$ if for every $k \in \mathbb{N}$ there is a run starting in the initial configuration $(q_0, \mathbf{0})$ and reaching a configuration $(q, \boldsymbol{m}_k) \in Q \times \mathbb{N}^d$ such that $\boldsymbol{m}_k[j] \geq k$ for every counter $j \in \omega(q, \boldsymbol{m})$ and $\boldsymbol{m}_k[i] \geq \boldsymbol{m}[i]$ for every counter $i \in [1, d] \setminus \omega(q, \boldsymbol{m})$.

Finally, we sometimes want to restrict only some counters of the VASS to stay non-negative. In this case, we consider extended configurations in $Q \times \mathbb{Z}_\omega^d$, where $\mathbb{Z}_\omega = \mathbb{Z} \cup \{\omega\}$. We say an extended run $\sigma = cf.\rho$ *remains non-negative* on counters $I \subseteq [1, d]$ if $cf'[I] \subseteq \mathbb{N}^{|I|}$ for all extended configurations cf' on σ.

Dyck Language. Towards the EXPSPACE upper bound of Theorem 3.1, a first step is to reduce the separability problem to a variant where one language is fixed to the Dyck language. The Dyck language D_n with n-letters is defined over the alphabet $\Sigma_n = \{a_i, \bar{a}_i \mid i \in [1, n]\}$. It contains those words w where, for every prefix v with $w = v.u$, we have at least as many letters a_i as $\bar{a}_i$. Thus, the letters behave like VASS counters and, indeed, the Dyck language is accepted by a single-state VASS $\mathcal{D}_n$ with n counters that increments the i-th counter upon seeing letter a_i and decrements the i-th counter upon seeing $\bar{a}_i$. If a VASS is defined over the Dyck alphabet Σ_n, we also call it n-visible. We will sometimes treat an n-visible VASS of dimension d as a $(d+n)$-dimensional VASS, and refer to the additional n counters as external. Note that this amounts to forming the product with $\mathcal{D}_n$. Given a path ρ, we use $\varphi(\rho)$ for the effect on the external counters in this product construction. Moreover, we write $\delta\varphi(\rho)$ to denote the combined effect on both internal and external counters, i.e. the $(d+n)$-dimensional vector $(\delta(\rho), \varphi(\rho))$.

To avoid an exponential blow-up, our reduction uses a variant of VASS whose transitions are labeled by compressed words. Essentially, the reduction takes $\mathcal{V}_1$ and $\mathcal{V}_2$ and produces a VASS $\mathcal{V}$ that is a product of $\mathcal{V}_1$ and $\mathcal{V}_2$. Moreover, it acts on its counters like $\mathcal{V}_1$; the input labels of $\mathcal{V}$ correspond to the counter updates of $\mathcal{V}_2$. Since the latter are binary-encoded, the new VASS will have binary encoded input words. Let us make this precise. A *label-compressed VASS* (lcVASS) $\mathcal{V}$ is a VASS, where the transitions are of the form $(p, a^m, \delta, q) \in Q \times \Sigma^* \times \mathbb{Z}^d \times Q$, where $a \in \Sigma$ and $m \in \mathbb{N}$ is given in binary. Thus, for an lcVASS $\mathcal{V}$, we define its *size* as $|\mathcal{V}| = |Q| + |F| + \sum_{(q,a^m,\delta,q') \in T} (\log_2(m) + \|\delta\|_2)$. The reduction that fixes the Dyck language is captured by the following lemma.

▶ **Lemma 5.1** ([2, Lemma 3.4]). *Given $\mathcal{V}_1$ and $\mathcal{V}_2$ over Σ, we can compute in time polynomial in $|\mathcal{V}_1| + |\mathcal{V}_2|$ an n-visible lcVASS $\mathcal{V}$ so that $L(\mathcal{V}_1) \,|\, L(\mathcal{V}_2)$ if and only if $L(\mathcal{V}) \,|\, D_n$. Here, n is the dimension of $\mathcal{V}_2$.*

The polynomial time bound is not mentioned in [2], but the simple construction they use (from [11]) clearly implies this bound. The latter separability problem $L(\mathcal{V}) \,|\, L(D_n)$ has been studied closely in [2] as well. They first show that $\mathcal{V}$ can be transformed so as to make the language $L(\mathcal{V})$ pumpable. For the resulting VASS, they show that $L(\mathcal{V}) \,|\, L(D_n)$ holds if and only if the so called Karp-Miller graph of $\mathsf{KM}(\mathcal{V})$ does not contain an inseparability witness. Unfortunately, the transformation required for pumpability involves another Karp-Miller graph construction, and therefore does not fit into the space bound we aim for (said graph can be of Ackermannian size in the worst case). Instead, we reformulate the witness.

▶ **Definition 5.2.** *Let $\mathcal{V}$ be an n-visible d-dimensional VASS. An* inseparability bloom *for $\mathcal{V}$ is a tuple $❀ = (q_f, I, \alpha, \beta, \gamma)$ with q_f a final state, loops α, β, γ starting and ending in q_f, and a partition of the counters $\Omega \uplus I = [1, d+n]$ so that*

(i) *for all* $\rho = \alpha, \beta, \gamma$*, we have* $\delta\varphi(\rho)[I] \geq 0$*,*
(ii) $\delta(\alpha) + \delta(\beta) + \delta(\gamma) \geq 0$
(iii) $\varphi(\alpha) + \varphi(\beta) \geq 0$*,*
(iv) *there is a* $t \in \mathbb{Q}$ *with* $\varphi(\alpha) + \varphi(\beta) + \varphi(\gamma) = t \cdot \varphi(\alpha)$*.*

The size *of the bloom is* $|❀| = \log_2 |Q| + |I| + |\alpha| + |\beta| + |\gamma|$*.*

A stem for ❀ *is an extended run* $cf.\sigma$ *of the product VASS* $\mathcal{V} \times \mathcal{D}_n$ *with the following properties: It ends in an extended configuration* $cf' = (q_f, \boldsymbol{m})$ *for some* $\boldsymbol{m} \in \mathbb{N}_\omega^{d+n}$ *with* $\Omega \subseteq \omega(cf')$*, and the counters in* $I \setminus \omega(cf)$ *remain non-negative when executing* σ *from* cf *resp.* α*,* β*, and* γ *from* cf'*.*

An inseparability flower $—❀ = (cf.\sigma, ❀)$ *consists of an inseparability bloom and a suitable stem. The size is* $|—❀| = \|cf\|_1 + |\sigma| + |❀|$*. The flower is* coverable *if the extended configuration* cf *is coverable in the product VASS* $\mathcal{V} \times \mathcal{D}_n$*.*

The following can be derived from the results in [2], refer to the full version for the details.

▶ **Lemma 5.3.** *Let* $\mathcal{V}$ *be* n*-visible. We have* $L(\mathcal{V}) \not| D_n$ *if and only if some inseparability flower is coverable.*

Our main result is the following. Note: the unary counter encoding strengthens the bound.

▶ **Theorem 5.4.** *If an inseparability flower is coverable in an* n*-visible lcVASS* $\mathcal{V}$ *of dimension* d*, then there is one of size at most* $|—❀| = 2^{|\mathcal{V}|^{\mathcal{O}((d+n)^2)}}$*.*

Main algortihm. We now have all ingredients to formulate the algorithm that proves the upper bounds in Theorem 3.1. We first describe the EXPSPACE upper bound. Given $\mathcal{V}_1$ and $\mathcal{V}_2$ whose languages we wish to separate, we first compute the lcVASS $\mathcal{V}$ using Lemma 5.1. This takes poly time. The task is to check $L(\mathcal{V}) \,|\, D_n$, where n is the dimension of $\mathcal{V}_2$. Using Lemma 5.3, we have to find an inseparability flower for $\mathcal{V}$ that is coverable. Theorem 5.4 bounds the size of the flowers we have to consider. We thus use non-determinism to find a flower of bounded size followed by a somewhat involved coverability check. Savitch's theorem [29] turns the non-deterministic algorithm into a deterministic one.

We detect a flower of bounded size as follows. We first guess the final state $q_f \in F$ and the partitioning of the counters $I \uplus \Omega$. With this information, we can guess the stem.

Towards obtaining a suitable stem, we start by guessing an extended configuration cf, whose non-ω-entries are at most doubly exponentially large. We can store such configurations in exponential space. If $\Omega \subseteq \omega(cf)$ fails, we abort. We now guess a path σ of doubly exponential length from cf to a configuration cf'. As we proceed, we store the length of the path, which only needs exponential space. We abort, if one of the following happens while guessing σ: a counter from $I \setminus \omega(cf)$ becomes negative, σ becomes too long, or the last state on σ is different from q_f. If we have not aborted until now, we have determined $cf.\sigma.cf'$ that may serve as a stem for a bloom with final state q_f and partition $I \uplus \Omega$.

Given the stem, we can finish the construction of the bloom by guessing the cycles $\rho = \alpha, \beta, \gamma$. The reason we proceed in this order is the following. The cycles are too long to be stored in exponential space. Instead, we check the non-negativity required by a stem on-the-fly, while constructing the cycles. To do so, we need the configuration cf', which we can store upon finishing the guess of σ above. While guessing the cycles, we store their length and the numbers $\delta(\rho)$ and $\varphi(\rho)$. It is readily checked that these numbers are bounded by

$$2^{|\mathcal{V}|} \cdot 2^{|\mathcal{V}|^{\mathcal{O}((d+n)^2)}} = 2^{|\mathcal{V}|^{\mathcal{O}((d+n)^2)}}.$$

This means we can store them in exponential space. We abort, if one of the following applies: an intermediate valuation becomes negative on $I \setminus \omega(cf)$, the path becomes too long, or the last state is different from q_f. We compute the operations and comparisons required by (i) to (iv) in Definition 5.2. For (iv), we start with counter $d+1$ and store the quotient $\frac{\varphi(\alpha)[d+1]+\varphi(\beta)[d+1]+\varphi(\gamma)[d+1]}{\varphi(\alpha)[d+1]}$ as the rational number t. As the numerator and denominator will be at most doubly exponential, we can store the number in exponential space. For the remaining entries, we only perform the required comparison. As noted above, they can be executed in polynomial time. If a comparison fails, we reject. If we have not rejected up to now, we have determined an inseparability flower while using space $|\mathcal{V}|^{\mathcal{O}((d+n)^2)} \leq 2^{|\mathcal{V}|\cdot\mathcal{O}((d+n)^2)} = 2^{poly(|\mathcal{V}_1|+|\mathcal{V}_2|)}$.

It remains to check whether this flower is coverable. There are two challenges. First, since cf is extended, we have a combination of a simultaneous unboundedness and a coverability problem. Second, the non-ω-entries in cf may be doubly exponentially large. We reduce the problem to simultaneous unboundedness, which is in EXPSPACE as shown by Demri [12, Theorem 4.6(I)]. The reduction uses a simple gadget that subtracts the counter valuation to be covered and, if succcessful, makes a new target counter j unbounded. In the end we check simultaneous unboundedness of $\omega(cf) \cup \{j\}$. To handle the large values, we utilize Lipton's construction [23], which allows a VASS to simulate EXPSPACE-computations. As a remark, simultaneous unboundedness cannot be expressed by a polynomial-sized formula in Yen's logic [33], a fact that was first observed in [12], which means we cannot just invoke the bounds in [1].

For the PSPACE upper bound, we merely need to observe that in fixed dimension, all our bounds on counter values become singly exponential. Moreover, for the gadget that subtracts counter values, we do not need the Lipton construction, as we can subtract exponentially bounded values directly using transitions. Finally, in fixed dimension, the simultaneous unboundedness check is also possible in PSPACE, as shown by Demri [12, Theorem 4.6(II)].

The remainder of the paper proves Theorem 5.4. SNLS help us deal with Constraint (iv).

6 Proof of Theorem 5.4

In this section we fix an n-visible lcVASS $\mathcal{V} = (Q, q_0, \Sigma_n, T, F)$. Note that since $\mathcal{V}$ is label-compressed, the effect on the external counters $\varphi(\rho)$ for a path ρ is also compressed. This matches the effect $\delta(\rho)$, which is anyway encoded in binary.

The proof is Rackoff-like, and we explain the analogy as we proceed. Like in Rackoff's upper bound for coverability [28], we reason over all (extended) configurations. Unlike Rackoff, however, we do not look at short covering sequences from a given configuration, but rather at small flowers rooted in said configuration. To bound the size while maximizing over all configurations, we measure each flower's size without considering the configuration it is rooted in. We therefore define the *flower bound*

$$\mathcal{B}_{\mathcal{V}} \;=\; \max_{cf} \min\{|\sigma| + |❀| \mid (cf.\sigma, ❀) \text{ is an inseparability flower}\}\ .$$

Theorem 5.4 is an immediate consequence of the following.

▶ **Lemma 6.1.** $\mathcal{B}_{\mathcal{V}} \leq 2^{|\mathcal{V}|^{\mathcal{O}((d+n)^2)}}$.

Step I: From length bounds to flower size bounds. Before we prove Lemma 6.1, let us see how it implies Theorem 5.4. Notice that Lemma 6.1 makes a statement about the lengths of the runs σ, α, β, γ, whereas Theorem 5.4 also promises a small starting configuration cf. Thus, it remains to construct a small starting configuration.

Proof of Theorem 5.4. To begin, assume $—❀ = (cf.\sigma, ❀)$ is coverable. By Lemma 6.1, there is another flower $(cf.\sigma', ❀')$ that is rooted in the same extended configuration and satisfies $|\sigma'| + |❀'| \leq \mathcal{B}_\mathcal{V}$. The extended configuration cf may not obey the desired bound. Thus, we replace it by cf' defined by $\omega(cf') = \omega(cf)$ and $cf'[j] = \min\{cf[j], \mathcal{B}_\mathcal{V} \cdot 2^{|\mathcal{V}|}\}$ for all $j \in I \setminus \omega(cf)$. Here, we keep the ω-entries, and for each non-ω-counter, we take the value of cf unless it is larger than $\mathcal{B}_\mathcal{V} \cdot 2^{|\mathcal{V}|}$, in which case we truncate to this value. We now claim that $—❀' = (cf'.\sigma', ❀')$ is still a coverable inseparability flower, and its size is at most doubly exponential (satisfying the desired bound). Coverability is immediate by $cf' \leq cf$. We also have $\|cf'\|_1 \leq 2^{2^{\mathcal{O}(|\mathcal{V}|)}}$. To be an inseparability flower, we have to check that the counters whose values we truncated when moving from cf to cf' remain non-negative while executing σ, and also $\sigma.\rho$ with $\rho = \alpha, \beta, \gamma$. This, however, is clear by the fact that $\sigma.\rho$ has length at most $\mathcal{B}_\mathcal{V}$, and each transition can subtract at most $2^{|\mathcal{V}|}$ tokens. As we have this budget available in cf', the run remains non-negative. ◀

In the remainder of the section, we prove Lemma 6.1. Similar to Rackoff's proof for coverability, we generalize the notion of flowers to admit negative counter values. Then we use an induction on the number of non-negative counters to establish a bound on the length of shortest generalized flowers. Let $i \in [0, d+n]$, $❀ = (q_f, I, \alpha, \beta, \gamma)$ a bloom, and $\Omega = [1, d+n] \setminus I$. We call an extended run $cf.\sigma$ of $\mathcal{V} \times \mathcal{D}_n$ an *i-stem for* $❀$, if it ends in an extended configuration cf' with $cf' = (q_f, \boldsymbol{m}')$ for some $\boldsymbol{m}'$, $\Omega \subseteq \omega(cf)$, and the counters in $([1, i] \cap I) \setminus \omega(cf)$ remain non-negative when executing σ from cf resp. α, β, and γ from cf' in $\mathcal{V} \times \mathcal{D}_n$. Here, we use $[1, 0] = \emptyset$. We call a pair $(cf.\sigma, ❀)$ consisting of an i-stem and an inseparability bloom an *i-inseparability flower*. Note that a $(d+n)$-inseparability flower is an inseparability flower as in Definition 5.2. We say the flower is *b-bounded* with $b \in \mathbb{N} \setminus \{0\}$, if the counters in $([1, i] \cap I) \setminus \omega(cf)$ are bounded by b along $cf.\sigma.\rho$ for all $\rho = \alpha, \beta, \gamma$. We wish to establish an estimate on the following function f, note that $f(d+n) = \mathcal{B}_\mathcal{V}$:

$$f(i) = \max_{cf} \min\{|\sigma| + |❀| \mid (cf.\sigma, ❀) \text{ is an } i\text{-inseparability flower}\} .$$

Step II: From value bounds to length bounds. Similar to Rackoff's proof, we will bound $f(i+1)$ in terms of $f(i)$, which will then yield the bound on $\mathcal{B}_\mathcal{V} = f(n+d)$. However, there is a key difference to Rackoff's proof: When constructing runs that, in the first $i+1$ coordinates stay non-negative, Rackoff argues that if such a run only uses counter values in $[0, b]$ on the first $i+1$ coordinates, then there is such a run of length at most $|Q| \cdot (b+1)^{i+1}$: whenever a combination of values in $[0, b]$ (and a control state) repeats, we can cut out the infix in between. Hence, *value bounds yield length bounds*. Our setting requires a different argument: Simply cutting out infixes in α, β, γ might spoil the properties (ii), (iii), and (iv) of i-inseparability flowers. Instead, we use our results on SNLS to obtain length bounds from value bounds. The technique is similar to some other generalizations of Rackoff's result by Yen [33, Lemma 3.5] Habermehl [14, Lemma 3.2] and Atig and Habermehl [1, Lemma 5]. However, the following proof also needs to work with non-linear constraints (which are also not present in Demri's extension of Rackoff's result [12]).

▶ **Lemma 6.2.** *Let $—❀ = (cf.\sigma, ❀)$ be a b-bounded i-inseparability flower. Then there is an i-inseparability flower $—❀' = (cf.\sigma', ❀')$ with $|\sigma'| + |❀'| \leq (2^{|\mathcal{V}|} \cdot |Q| \cdot b)^{\mathcal{O}((d+n)^6)}$.*

Proof. Let $—❀ = (cf.\sigma, ❀)$ be a b-bounded i-inseparability flower with $❀ = (q_f, I, \alpha, \beta, \gamma)$. Let Ω be the complement of I. Let $\widehat{cf}$ be the extended configuration reached by $cf.\sigma$.

Let $D = ([1, i] \cap I) \setminus \omega(cf)$ be the counters we wish to keep non-negative. By a *D-loop*, we mean an extended run $cf_1.\tau.cf_2$ of $\mathcal{V} \times \mathcal{D}_n$ where τ is a loop, $cf_1[D] = cf_2[D]$, and $cf_1[\Omega] = cf_2[\Omega]$. The D-loop is called *irreducible* if it does not contain further D-loops. By

the pigeonhole principle, any run that is longer than $u_{len} = |Q| \cdot (b+1)^{|D \cup \Omega|}$ contains a D-loop. This means the length of an irreducible D-loop is at most u_{len}. Moreover, these loops can have at most $2 \cdot 2^{|\mathcal{V}|} \cdot u_{len}$ distinct effects on each counter, where the leading 2 is for the distinction between positive and negative values. This analysis yields an upper bound of $u_{num} = (2^{|\mathcal{V}|+1} \cdot u_{len})^{d+n}$ on the number of irreducible D-loops with distinct effects. We show how to construct a flower as promised in the lemma.

We begin by cutting out irreducible D-loops from σ, yielding σ' of length at most u_{len}. The extended configuration cf' reached by $cf.\sigma'$ coincides with $\widehat{cf}$ on the ω-entries and on the non-ω-entries for D.

Next, we want to shorten $\rho = \alpha, \beta, \gamma$. To this end, we first decompose them into irreducible D-loops. We assume all such D-loops have distinct effects, otherwise we pick a representative. With the previous analysis, the result is $\rho_0, \rho_1, \ldots, \rho_v$ with $v \leq u_{num}$. Here, ρ_0 is the loop on the final state into which the irreducible D-loops $\rho_1, \ldots, \rho_v$ are inserted. Note that ρ_0 is not necessarily a D-loop. What we know, however, is that ρ_0 has a non-negative effect on the counters in I, due to Condition (i) in the definition of inseparability blooms. Since we want to be able to insert all the D-loops directly into ρ_0, the latter should still contain the same starting configurations for such loops as ρ, at least when only considering the counters in $D \cup \Omega$. Therefore we cannot guarantee that ρ_0 contains no D-loops, because cutting all of them out might reduce the number of such configurations visited by ρ_0. However, ρ_0 still has length at most u_{len}^2 by the following argument. If we mark in ρ_0 the first occurrence of each element from $Q \times [0,b]^{D \cup \Omega}$, then each infix leading from one such marker to the next has at most length u_{len}, because longer infixes still contain a D-loop that has no marked configuration and can therefore be removed. Since there are at most $|Q \times [0,b]^{D \cup \Omega}| = |Q| \cdot (b+1)^{|D \cup \Omega|} = u_{len}$ markers, we obtain the stated length bound of u_{len}^2 for ρ_0.

A vector $\boldsymbol{x} \in \mathbb{N}^v$ with $\boldsymbol{x}[0] > 0$ and $\boldsymbol{x}[j] \geq 0$ for $1 \leq j \leq v$ can now be turned into a run $\rho_{\boldsymbol{x}}$ by glueing together $\boldsymbol{x}[j]$-many instances of ρ_j. Note that also the base loop ρ_0 may be repeated. As the order of the transitions does not influence the effect of the run, $\delta\varphi(\rho) = \delta\varphi(\rho_{\boldsymbol{y}})$ holds, where $\boldsymbol{y} = \Psi(\rho)$ is the so-called Parikh vector that counts the occurrences of irreducible loops in ρ. As a consequence, we can directly define the effect on the vector $\boldsymbol{x}$, namely $\delta\varphi(\boldsymbol{x}) = \sum_{i=0}^{v} \boldsymbol{x}[j] \cdot \delta\varphi(\rho_j) \in \mathbb{Z}^{d+n}$.

With irreducible D-loops at hand, we can formulate the search for a small bloom as an SNLS. We define the vectors $\boldsymbol{x}_\alpha, \boldsymbol{x}_\beta, \boldsymbol{x}_\gamma$, and the following constraints:

$$\begin{aligned}
\boldsymbol{x}_\alpha[0], \boldsymbol{x}_\beta[0], \boldsymbol{x}_\gamma[0] &\geq 1 & \delta(\boldsymbol{x}_\alpha + \boldsymbol{x}_\beta + \boldsymbol{x}_\gamma)[1,d] &\geq 0 \\
\boldsymbol{x}_\alpha, \boldsymbol{x}_\beta, \boldsymbol{x}_\gamma &\geq 0 & \varphi(\boldsymbol{x}_\alpha + \boldsymbol{x}_\beta)[d+1, d+n] &\geq 0 \\
\delta\varphi(\boldsymbol{x}_\alpha)[I], \delta\varphi(\boldsymbol{x}_\beta)[I], \delta\varphi(\boldsymbol{x}_\gamma)[I] &\geq 0 & (\varphi(\boldsymbol{x}_\alpha + \boldsymbol{x}_\beta + \boldsymbol{x}_\gamma) - t \cdot \varphi(\boldsymbol{x}_\alpha))[d+1, d+n] &= 0
\end{aligned}$$

The constraints on the left say that we repeat the base loops at least once, and the remaining loops a non-negative number of times. The last constraint is Condition (i) in the definition of blooms. The constraints on the right correspond to the Conditions (ii) to (iv).

It is readily checked that $(\Psi(\alpha), \Psi(\beta), \Psi(\gamma))$ solves the SNLS. This means Theorem 3.2 applies and yields a small rational solution. To turn it into an integer solution, we observe that our SNLS is monotonic in the sense that if $(\boldsymbol{x}_\alpha, \boldsymbol{x}_\beta, \boldsymbol{x}_\gamma)$ is a solution, so is $(k\boldsymbol{x}_\alpha, k\boldsymbol{x}_\beta, k\boldsymbol{x}_\gamma)$ for any $k \geq 1$. We multiply the rational solution by the common denominator to obtain the integer solution $(\boldsymbol{x}'_\alpha, \boldsymbol{x}'_\beta, \boldsymbol{x}'_\gamma)$ with the associated loops α', β', γ'.

We argue that $\multimap^{\text{❀}'} = (cf.\sigma', \text{❀}')$ with $\text{❀}' = (q_f, I, \alpha', \beta', \gamma')$ is an i-inseparability flower. Remember that cf' is the configuration reached by $cf.\sigma'$. As cf and I have not changed, the only thing we have to show is the non-negativity for the counters in D. For σ', this follows from the non-negativity of σ, and the fact that we only cut-out D-loops. For $\rho' = \alpha', \beta', \gamma'$,

which are executed from cf', we argue as follows. The base loop ρ_0 has the same effect on the D-counters as ρ, because we obtained ρ_0 by cutting-out D-loops from ρ. The effect of ρ on even the entire set I is non-negative due to Condition (i) for blooms. This means if one repetition of ρ_0 stays non-negative from cf', arbitrarily many repetitions will. The one repetition stays non-negative, because ρ was non-negative from $\widehat{cf}$ (the extended configuration reached by $cf.\sigma$), and cf' coincides with $\widehat{cf}$ on D. Since the D-loops that we glue into ρ_0 come with a valuation of the counters in D, and this valuation keeps them non-negative in ρ, the entire ρ' stays non-negative.

We analyze the complexity of our system $\mathcal{S}$. It has at most $row(\mathcal{S}) \in \mathcal{O}(d+n)$ many rows, and note that the non-negativity constraints do not count towards the rows. There are at most $col(\mathcal{S}) \in \mathcal{O}(u_{num})$ many columns. The degree is $\deg(\mathcal{S}) = 1$. The maximal coefficient is bounded from above by the largest possible loop effect, $\max(\mathcal{S}) \leq 2^{|\mathcal{V}|} \cdot u_{len}^2$. Then, Theorem 3.2 gives us rational solutions $\boldsymbol{x}_\alpha, \boldsymbol{x}_\beta, \boldsymbol{x}_\gamma$ of the form $\boldsymbol{x}_\rho[i] = \frac{a}{K}$, meaning K is the common denominator of all entries, with

$$\begin{aligned}
\|\boldsymbol{x}_\rho\|_1 &\in (col(\mathcal{S}) \cdot \deg(\mathcal{S}) \cdot \max(\mathcal{S}))^{\mathcal{O}(\deg(\mathcal{S})^2 \cdot row(\mathcal{S})^4)} \\
&= (\mathcal{O}(u_{num}) \cdot 2^{|\mathcal{V}|} \cdot u_{len}^2)^{\mathcal{O}((d+n)^4)} \\
&= (\mathcal{O}((2^{|\mathcal{V}|+1} \cdot u_{len})^{d+n}) \cdot 2^{|\mathcal{V}|} \cdot u_{len}^2)^{\mathcal{O}((d+n)^4)} \\
&= (2^{|\mathcal{V}|} \cdot u_{len})^{\mathcal{O}((d+n)^5)} = \ (2^{|\mathcal{V}|} \cdot |Q| \cdot (b+1))^{\mathcal{O}((d+n)^6)} \\
&= (2^{|\mathcal{V}|} \cdot |Q| \cdot b)^{\mathcal{O}((d+n)^6)} \ .
\end{aligned}$$

We already argued that the integer vectors $K\boldsymbol{x}_\alpha$, $K\boldsymbol{x}_\beta$, $K\boldsymbol{x}_\gamma$ are also solutions with runs α', β', γ'. Since each entry of these vectors is smaller than $(2^{|\mathcal{V}|} \cdot |Q| \cdot b)^{\mathcal{O}((d+n)^6)}$, and we have at most $u_{num} = (2^{|\mathcal{V}|} \cdot |Q| \cdot (b+1))^{\mathcal{O}((d+n)^2)}$ many loops with maximal size $u_{len}^2 = |Q|^2 \cdot (b+1)^{2(d+n)}$, we get

$$\begin{aligned}
|\rho'| &\leq (2^{|\mathcal{V}|} \cdot |Q| \cdot b)^{\mathcal{O}((d+n)^6)} \cdot (2^{|\mathcal{V}|} \cdot |Q| \cdot (b+1))^{\mathcal{O}((d+n)^2)} \cdot |Q|^2 \cdot (b+1)^{2(d+n)} \\
&= (2^{|\mathcal{V}|} \cdot |Q| \cdot b)^{\mathcal{O}((d+n)^6)} \ .
\end{aligned}$$

◀

Step III: Rackoff-style induction. We now give the bound on $f(i)$ that we need to prove Lemma 6.1. In the base case, no counter has to remain non-negative and so we have a 1-bounded 0-inseparability flower. We employ the bound from Lemma 6.2.

▶ **Lemma 6.3.** $f(0) = (2^{|\mathcal{V}|}|Q|)^{\mathcal{O}((d+n)^6)}$.

In the induction step, and as in Rackoff's result, the bound takes the form of a recurrence.

▶ **Lemma 6.4.** $f(i+1) \leq (2^{|\mathcal{V}|} \cdot f(i))^{\mathcal{O}((d+n)^6)}$.

Proof. Consider an $(i+1)$-inseparability flower —❀ $= (cf.\sigma, ❀)$ with ❀ $= (q_f, I, \alpha, \beta, \gamma)$ and $\Omega = [1, d+n] \setminus I$. Let $r(i) = 2^{|\mathcal{V}|} \cdot f(i)$ serve as an abbreviation. We proceed by a case distinction. If —❀ is $r(i)$-bounded, then Lemma 6.2 provides another $(i+1)$-inseparability flower $(cf.\sigma', ❀')$ with $|\sigma'| + |❀'| \leq (2^{|\mathcal{V}|} \cdot |Q| \cdot r(i))^{\mathcal{O}((d+n)^6)} = (2^{|\mathcal{V}|} \cdot f(i))^{\mathcal{O}((d+n)^6)}$. This satisfies the bound stated in the lemma.

If —❀ is not $r(i)$-bounded, then $\sigma.\rho$ with $\rho = \alpha, \beta, \gamma$ exceeds $r(i)$. We identify the first moment when this happens, say after ρ_1 and for the $(i+1)$-th counter. The case where the run exceeds the bound already in σ is simpler. The run decomposes into

$$cf \xrightarrow{\sigma} cf_1 \xrightarrow{\rho_1} cf' \xrightarrow{\rho_2} cf_2 \ .$$

We argue that also $(cf'.\rho_2, ❀)$ is an $(i+1)$-flower, which means $cf'.\rho_2$ is an $(i+1)$-stem for ❀. Since ρ is a loop, it returns to q_f. We have $\omega(cf) = \omega(cf')$ and since $\Omega \subseteq \omega(cf)$, we get $\Omega \subseteq \omega(cf')$. As ❀ is a bloom, ρ has a non-negative effect on the counters in I. This means $cf_1[I] \leq cf_2[I]$, and so the counters in $([1, i+1] \cap I) \setminus \omega(cf')$ remain non-negative when executing α, β, γ from cf_2 as they did from cf_1. Also ρ_2 remains non-negative from cf', because ρ remained non-negative from cf_1.

Since $(cf'.\rho_2, ❀)$ is an $(i+1)$-flower, it is an i-flower. The induction hypothesis yields another i-flower $—❀' = (cf'.\sigma', ❀')$ with $|\sigma'| + |❀'| \leq f(i)$. Let $❀' = (q'_f, I', \alpha', \beta', \gamma')$ and let Ω' be the complement of I'.

We argue that $—❀'$ is actually an $(i+1)$-flower. If the counter $i+1$ that exceeds the bound $r(i)$ does not belong to I', there is nothing to show. Otherwise, even $\sigma'.\alpha'.\beta'.\gamma'$ in succession could subtract at most $f(i) \cdot 2^{|\mathcal{V}|} = r(i)$ tokens from counter $i+1$. Since this counter carries more than $r(i)$ tokens, this leaves us with a positive balance.

We show that also $(cf.\sigma.\rho_1.\sigma', ❀')$ is an $(i+1)$-flower, meaning $cf.\sigma.\rho_1.\sigma'$ is an $(i+1)$-stem for ❀'. We have $\omega(cf) = \omega(cf')$ and so $\Omega' \subseteq \omega(cf')$ implies $\Omega' \subseteq \omega(cf)$. It remains to show that the counters in $([1, i+1] \cap I') \setminus \omega(cf)$ remain non-negative. Consider the prefix $\sigma.\rho_1$ executed from cf. For the counters that also belong to I, non-negativity holds as $—❀$ is an $(i+1)$-flower. Assume there was a counter in $([1, i+1] \cap I') \setminus \omega(cf)$ that did not belong to I. Then it belonged to Ω. But as $\Omega \subseteq \omega(cf)$, we had a contradiction. For the suffix σ' executed from cf', and for α', β', γ', non-negativity holds as $—❀'$ is an $(i+1)$-flower.

To estimate the size of the newly constructed flower rooted in cf, we assume $\sigma.\rho_1$ does not repeat configurations on the first $(i+1)$-counters. If it does, we cut out the infix and adapt the values of the counters that are allowed to fall below zero. Then the length of $\sigma.\rho_1$ is bounded by $|Q| \cdot r(i)^{i+1}$, and we have

$$|\sigma.\rho_1.\sigma'| + |❀'| \leq |Q| \cdot r(i)^{i+1} + f(i) \leq \ (2^{|\mathcal{V}|} \cdot f(i))^{\mathcal{O}(i+1)} \leq (2^{|\mathcal{V}|} \cdot f(i))^{\mathcal{O}((d+n)^6)} \ . \qquad \blacktriangleleft$$

It remains to solve the recurrence. Let $a = 2^{|\mathcal{V}|}$ and $b = \mathcal{O}((d+n)^6)$. We have

$$f(d+n) = (a \dots (a \cdot f(0))^b \dots)^b \leq (a^{d+n} \cdot f(0))^{b^{d+n}} \ .$$

Since $f(0) = (2^{|\mathcal{V}|} \cdot |Q|)^{\mathcal{O}((d+n)^6)}$, we obtain the promised $\mathcal{B}_\mathcal{V} = f(d+n) \leq 2^{|\mathcal{V}|^{\mathcal{O}((d+n)^2)}}$.

References

1 M. F. Atig and P. Habermehl. On Yen's path logic for Petri nets. *Int. J. Found. Comput. Sci.*, 22(4):783–799, 2011. doi:10.1142/S0129054111008428.

2 P. Baumann, R. Meyer, and G. Zetzsche. Regular separability in Büchi VASS. In *Proc. STACS*, volume 254 of *LIPIcs*, pages 9:1–9:19. Schloss Dagstuhl, 2023. doi:10.4230/LIPICS.STACS.2023.9.

3 Pascal Baumann, Roland Meyer, and Georg Zetzsche. Regular separability in Büchi VASS. *CoRR*, abs/2301.11242, 2023. doi:10.48550/arXiv.2301.11242.

4 M. Blockelet and S. Schmitz. Model checking coverability graphs of vector addition systems. In *Proc. MFCS*, volume 6907 of *LNCS*, pages 108–119. Springer, 2011. doi:10.1007/978-3-642-22993-0_13.

5 L. Clemente, W. Czerwiński, S. Lasota, and C. Paperman. Regular separability of parikh automata. In *Proc. ICALP*, volume 80 of *LIPIcs*, pages 117:1–117:13. Schloss Dagstuhl, 2017. doi:10.4230/LIPICS.ICALP.2017.117.

6 L. Clemente, W. Czerwiński, S. Lasota, and C. Paperman. Separability of reachability sets of vector addition systems. In H. Vollmer and B. Vallée, editors, *Proc. STACS*, volume 66 of *LIPIcs*, pages 24:1–24:14. Schloss Dagstuhl, 2017. doi:10.4230/LIPICS.STACS.2017.24.

7 W. Czerwiński and P. Hofman. Language inclusion for boundedly-ambiguous vector addition systems is decidable. In *Proc. CONCUR*, volume 243 of *LIPIcs*, pages 16:1–16:22. Schloss Dagstuhl, 2022. doi:10.4230/LIPICS.CONCUR.2022.16.

8 W. Czerwiński, P. Hofman, and G. Zetzsche. Unboundedness problems for languages of vector addition systems. In *Proc. ICALP*, volume 107 of *LIPIcs*, pages 119:1–119:15. Schloss Dagstuhl, 2018. doi:10.4230/LIPICS.ICALP.2018.119.

9 W. Czerwiński and S. Lasota. Regular separability of one counter automata. *Log. Methods Comput. Sci.*, 15(2), 2019. doi:10.23638/LMCS-15(2:20)2019.

10 W. Czerwiński, S. Lasota, R. Meyer, S. Muskalla, K. N. Kumar, and P. Saivasan. Regular separability of well-structured transition systems. In *Proc. CONCUR*, volume 118 of *LIPIcs*, pages 35:1–35:18. Schloss Dagstuhl, 2018. doi:10.4230/LIPICS.CONCUR.2018.35.

11 W. Czerwiński and G. Zetzsche. An approach to regular separability in vector addition systems. In *Proc. LICS*, pages 341–354. ACM, 2020. doi:10.1145/3373718.3394776.

12 S. Demri. On selective unboundedness of VASS. *JCSS*, 79(5):689–713, 2013. doi:10.1016/J.JCSS.2013.01.014.

13 E. M. Gurari and O. H. Ibarra. An NP-complete number-theoretic problem. *J. ACM*, 26(3):567–581, 1979. doi:10.1145/322139.322152.

14 P. Habermehl. On the complexity of the linear-time μ-calculus for Petri Nets. In *Proc. ICATPN*, volume 1248 of *LNCS*, pages 102–116. Springer, 1997. doi:10.1007/3-540-63139-9_32.

15 E. Keskin and R. Meyer. Separability and non-determinizability of WSTS. In *Proc. CONCUR*, volume 279 of *LIPIcs*, pages 8:1–8:17. Schloss Dagstuhl, 2023. doi:10.4230/LIPICS.CONCUR.2023.8.

16 E. Keskin and R. Meyer. On the separability problem of VASS reachability languages. In *To appear in Proc. of LICS*, 2024.

17 C. Köcher and G. Zetzsche. Regular separators for VASS coverability languages. In *Proc. FSTTCS*, volume 284 of *LIPIcs*, pages 15:1–15:19. Schloss Dagstuhl, 2023. doi:10.4230/LIPICS.FSTTCS.2023.15.

18 S. R. Kosaraju. Decidability of reachability in vector addition systems (preliminary version). In *Proc. STOC*, pages 267–281. ACM, 1982. doi:10.1145/800070.802201.

19 J.-L. Lambert. A structure to decide reachability in Petri nets. *Theor. Comput. Sci.*, 99(1):79–104, 1992. doi:10.1016/0304-3975(92)90173-D.

20 S. Lang. *Algebra, Rev. 3rd Ed.* Springer, New York, 2002.

21 J. Leroux. Vector addition system reachability problem: a short self-contained proof. In *Proc. POPL*, pages 307–316. ACM, 2011. doi:10.1145/1926385.1926421.

22 J. Leroux and S. Schmitz. Demystifying reachability in vector addition systems. In *Proc. LICS*, pages 56–67. IEEE Computer Society, 2015. doi:10.1109/LICS.2015.16.

23 R. J. Lipton. The reachability problem requires exponential space. Technical Report 63, Yale University, 1976.

24 Angus Macintyre, Kenneth McKenna, and Lou van den Dries. Elimination of quantifiers in algebraic structures. *Advances in Mathematics*, 47(1):74–87, 1983. doi:10.1016/0001-8708(83)90055-5.

25 D. Marker. *Model Theory: An Introduction.* Springer, New York, 2002.

26 E. W. Mayr. An algorithm for the general Petri net reachability problem. *SIAM J. Comput.*, 13(3):441–460, 1984. doi:10.1137/0213029.

27 B. Mishra. *Algorithmic Algebra.* Texts and Monographs in Computer Science. Springer, 1993. doi:10.1007/978-1-4612-4344-1.

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

29 W. J. Savitch. Relationships between nondeterministic and deterministic tape complexities. *JCSS*, 4(2):177–192, 1970. doi:10.1016/S0022-0000(70)80006-X.

30 A. Schrijver. *Theory of linear and integer programming.* John Wiley & Sons, 1986.

31 R. S. Thinniyam and G. Zetzsche. Regular separability and intersection emptiness are independent problems. In *Proc. FSTTCS*, volume 150 of *LIPIcs*, pages 51:1–51:15. Schloss Dagstuhl, 2019. doi:10.4230/LIPICS.FSTTCS.2019.51.

32 V. Weispfenning. The complexity of almost linear diophantine problems. *J. Symb. Comput.*, 10(5):395–404, 1990. doi:10.1016/S0747-7171(08)80051-X.

33 H.-C. Yen. A unified approach for deciding the existence of certain Petri net paths. *Information and Computation*, 96(1):119–137, 1992. doi:10.1016/0890-5401(92)90059-O.